Stepwise versus globally optimal search in children and adults

Björn Meder^{a,b,c,1,*}, Jonathan D. Nelson^{a,d}, Matt Jones^e, Azzurra Ruggeri^{a,f}

^a*MPRG iSearch, Max Planck Institute for Human Development, Berlin, Germany*

^b*Center for Adaptive Rationality, Max Planck Institute for Human Development, Berlin, Germany*

^c*University of Erfurt, Germany*

^d*University of Surrey, United Kingdom*

^e*University of Colorado Boulder, United States*

^f*Technical University of Munich, Germany*

Abstract

How do children and adults search for information when stepwise-optimal strategies fail to identify the most efficient query? The value of questions is often measured in terms of stepwise information gain (expected reduction of entropy on the next time step) or other stepwise-optimal methods. However, such myopic models are not guaranteed to identify the most efficient sequence of questions, that is, the shortest path to the solution. In two experiments we contrast stepwise methods with globally optimal strategies and study how younger children (around age 8, $N=52$), older children (around age 10, $N=99$), and adults ($N=101$) search in a 20-questions game where planning ahead is required to identify the most efficient first question. Children searched as efficiently as adults, but also as myopically. Both children and adults tended to rely on heuristic stepwise-optimal strategies, focusing primarily on questions' implications for the next time step, rather than planning ahead.

Keywords: Sequential search, information gain, stepwise-optimal methods, 20-questions game, split-half heuristic, entropy, search efficiency, planning

1. Introduction

From children actively learning about the world, to the control of eye movements in visual perception, to a doctor running tests for medical diagnosis—humans are information foragers. Any comprehensive account of cognition and behavior must

*Corresponding author

Email address: meder@mpib-berlin.mpg.de, bmeder@posteo.de (Björn Meder)

¹Data and R code available at <https://osf.io/cq48j/>.

therefore explain not only how humans learn from observed data, but also how they actively search for relevant information.

Different metrics for quantifying the value of *queries* (such as questions, medical tests, or eye movements) have been suggested in statistics, computer science, philosophy, and psychology (for reviews, see Coenen et al., 2018; Crupi et al., 2018; Nelson, 2005; Settles, 2010). These metrics measure, for instance, how much uncertainty is reduced (Lindley, 1956), how much classification accuracy is improved (Baron, 1985), or how much a piece of information changes beliefs (Wells & Lindsay, 1980). They serve as candidate normative or descriptive models in different fields, including developmental psychology (Nelson et al., 2014; Ruggeri et al., 2016, 2017), vision research (Najemnik & Geisler, 2005; Nelson & Cottrell, 2007), and higher level cognition (Bramley et al., 2015; Markant & Gureckis, 2014; Meder & Nelson, 2012; Nelson et al., 2010; Oaksford & Chater, 1994; Steyvers et al., 2003; Wu et al., 2017).

Typically, these models are implemented in a *stepwise-optimal* way, meaning that they consider a query’s implications only for the immediate next time step. Because they disregard future queries and do not plan ahead they are also referred to as *greedy* or *myopic* models. Obtaining information in a stepwise-optimal fashion is often an efficient approach to information acquisition, but it is not generally guaranteed to identify the most efficient strategy when multiple queries can be conducted (Hyafil & Rivest, 1976; Nelson et al., 2018). This has important theoretical implications for both computational-level and mechanistic analyses of human search behavior. At the computational level, almost all psychological research to date has tacitly assumed that sequential search is governed by the same normative principles that govern one-shot scenarios, in which only a single piece of information can be obtained (but see Bramley et al., 2015; Meier & Blair, 2013; Nelson et al., 2018). At the mechanistic level, stepwise strategies engage a fundamentally different form of cognitive processing than methods that are sensitive to long-run efficiency, the former involving valuation of individual queries and the latter requiring explicit planning to evaluate alternative decision trees.

The focus on stepwise-optimal strategies for information acquisition sharply contrasts with the literature on reward-based tasks, where the tension between immediate and long-term outcomes has long been acknowledged and investigated from both computational and psychological perspectives (e.g., Bellman, 1957; Newell & Simon, 1959; Sutton & Barto, 1998; Watkins & Dayan, 1992). One critical issue contributing to this blind spot in the psychological literature on information acquisition is the lack of a precise characterization of different search task environments and the resulting implications for the performance of alternative search strategies. From a computational perspective, long-run considerations are critical to characterize and

understand the optimality conditions of different methods for question selection. From a psychological perspective, sequential search scenarios are particularly critical to assess the *efficiency* of human search behavior and to what extent people plan ahead when deciding what information to acquire.

In this paper, we directly contrast stepwise strategies with more globally efficient strategies in sequential search situations. We first provide a computational analysis of different search strategies in the 20-questions game, a search task widely used in developmental and cognitive psychology. In this search task, the goal is to identify an unknown target item by asking as few yes–no questions as possible. Our analysis specifies and illustrates the conditions under which stepwise-optimal strategies can be distinctly suboptimal when the goal of the searcher is to minimize the expected number of questions needed to identify the target item. We mathematically characterize the relationships between different models of the value of information, such as information gain (entropy reduction), and strategies for identifying the target item with the minimum total number of queries. We contribute code that can be used by other researchers to compute each question’s efficiency in the 20-questions game and to identify environments in which stepwise methods and efficiency considerations entail systematically different search behavior. We then present two studies investigating children’s and adults’ sequential search in a version of the 20-questions game in which it is necessary to plan at least two steps ahead to determine which of two stepwise-optimal queries is more efficient in the long run. In this sense, we explore from a developmental perspective to what extent searchers plan ahead to achieve high search efficiency by considering what questions will be available later on.

2. Theoretical background

Models of the informational value of queries are also known as *Optimal Experimental Design* (OED) theories (Nelson, 2005). They are optimal in the sense that they maximize some information measure by evaluating each query’s possible outcomes and their immediate implications for the hypotheses considered. For instance, *expected information gain* (Lindley, 1956) values questions in accordance with the expected reduction in uncertainty, as measured by Shannon (1948) entropy. OED theories were initially envisioned for situations in which just one query could be conducted, but they have since been applied in a stepwise fashion to situations where multiple queries can be conducted in series. A *stepwise-optimal* query is one that offers the maximum expected informational value (e.g., maximum expected information gain) in the next time step, without regard for what subsequent queries could be conducted.

However, crucially, a stepwise-optimal query is not necessarily the best first query when multiple queries can be conducted, because the feedback on one query affects and constrains the informativeness of subsequent queries. Sequential search is particularly interesting from a computational standpoint, because the problem of identifying the globally optimal strategy is generally intractable (Hyafil & Rivest, 1976; Nelson et al., 2018). From a psychological perspective, sequential search relates to many real-world situations (e.g., medical diagnosis involving multiple tests) and it is methodologically critical for differentiating between competing descriptive and normative models of human search.

2.1. Stepwise-optimal search in the 20-questions game

Figure 1a shows a variant of the 20-questions game. In this search task, the goal is to identify a randomly chosen target item (in this case, a *monster*) with as few yes–no questions as possible (e.g., “Is the monster green?” or “Is the monster round?”). If you could query one of the features possessed by multiple items—color, shape or pattern—or query an individual monster (Figure 1b), what would you ask?

Figure 1. A 20-questions game where stepwise-optimal information gain fails to identify the most efficient first question. a) Item set (identifiers M1–M12 not shown during experiments). b) Possible first queries, targeting features (color, shape, pattern) shared by multiple items or present in individual items.

Formally, this search task has the following characteristics: (i) uniform prior probabilities over the hypotheses (i.e., all monsters are equally likely to be the target), (ii) only binary (yes–no) questions can be asked, and (iii) deterministic likelihoods (i.e., each item either does or does not have a given feature). Early studies of the 20-questions game primarily focused on a qualitative distinction between different types of questions, such as *hypothesis-scanning* questions, which target a single item, and *constraint-seeking* questions, which pertain to features shared by multiple items (Denney & Denney, 1973; Mosher & Hornsby, 1966; Siegler, 1977; Thornton,

1982). For instance, for the stimuli shown in Figure 1a, a hypothesis-scanning question would be “Does the target monster have a crown?” whereas an example of a constraint-seeking question is “Is the target monster green?” Because this qualitative distinction does not fully capture the varying usefulness of questions, recent research has used stepwise information gain as a graded measure for questions’ informativeness and as a benchmark for evaluating human search behavior (Kachergis et al., 2017; Nelson et al., 2014; Ruggeri et al., 2016, 2017).

The idea behind information gain is that reduction in uncertainty, measured via Shannon (1948) entropy, indicates the amount of information gained about the true hypothesis (the unknown target item). The information gain of a question is the entropy in the hypothesis space (i.e., the set of items considered and the associated probability distribution) *before* asking that question minus the expected entropy *after* asking that question (for related metrics based on different entropy measures, see Crupi et al., 2018). In the 20-questions game, the assumptions of uniform prior probabilities and deterministic likelihoods imply all items that are consistent with the answers received thus far are equally likely to be the target. Thus the prior entropy at each step is $\log n$ where n is the number of remaining items.² For a given query Q , let n_{yes} denote the number of remaining items that have the queried feature value (e.g., are green), and let n_{no} denote the number that lack that feature value (e.g., are blue). The expected information gain of a question Q is then given by

$$\text{IG}(Q) = \log n - \left[\frac{n_{\text{yes}}}{n} \log n_{\text{yes}} + \frac{n_{\text{no}}}{n} \log n_{\text{no}} \right] \quad (1)$$

where IG is the expected information gain and the term in the brackets is the expected entropy after asking the question. Selecting questions by maximizing information gain in accordance with Equation 1 reduces uncertainty in a stepwise-optimal fashion (for an analysis of multistep methods see Nelson et al., 2018).

2.2. Stepwise-optimal search and the split-half heuristic

Equation 1 shows that stepwise information gain in the 20-questions game is directly related to the split a question induces in the hypothesis space (Navarro & Perfors, 2011; Nelson et al., 2014). Figure 2a shows the expected information gain of a question as a function of the split it induces, expressed here as the proportion of items

²The choice of base for the logarithm is arbitrary; we use base 2 throughout this paper, in which case the unit is the bit.

possessing the queried feature. The information gain function has its maximum when exactly half of the items possess the targeted feature, and it monotonically decreases in both directions as the distribution becomes less even. Accordingly, we define the *splithaleness* of a question as $\min(n_{\text{yes}}/n, n_{\text{no}}/n)$. This measure ranges from 0 to $\frac{1}{2}$ and is monotonically related to stepwise information gain for any fixed number of items n . Given the monster features in Figure 1, color and shape induce the most balanced possible split (3:9, *splithaleness* = .25) and thus tie for the greatest stepwise information gain (0.811 bits). Asking about the pattern and inquiring about an individual monster induce 2:10 and 1:11 splits (*splithaleness* = .17 and .08), respectively, and therefore have lower stepwise information gain (0.65 and 0.414 bits).

The close relationship between *splithaleness* and stepwise information gain provides a rational basis for the *split-half heuristic* (Nelson et al., 2014), which characterizes much of older children’s and adults’ behavior when selecting among given questions in the 20-questions game (Eimas, 1970; Denney & Denney, 1973; Nelson et al., 2014; Siegler, 1977; Thornton, 1982). The split-half heuristic queries the feature whose distribution comes as close as possible to a 50:50 split. By selecting the query with the greatest *splithaleness*, the heuristic maximizes stepwise information gain without explicitly calculating the underlying quantities in the information gain equations (i.e., prior and posterior entropy). The split-half heuristic also maximizes several other entropy-based measures (Crupi et al., 2018) and OED models (Nelson et al., 2018).

Figure 2. Relationships among models and simulation of their performance. *a)* Relation between the proportion of items possessing a feature and its stepwise expected information gain. *b)* Efficiency of different strategies in the environment of Figure 1 (see also Table A1). *c)* Average information gain per question asked. Box plots show interquartile range with mean (diamond) and range (whiskers). Strategy performance based on 10^5 runs for each of the 12 target items.

2.3. The question of efficiency: Stepwise-optimal versus globally optimal strategies

Selecting questions in accordance with stepwise information gain is not necessarily the most *efficient* strategy in the sense of minimizing the expected total number of queries. From a computational standpoint, the critical issue is whether there are limitations on the questions that can be asked or whether arbitrary queries are possible (e.g., only questions pertaining to a prespecified set of features or also questions pertaining to arbitrary subsets, such as “Is it one of the three monsters M1, M2, or M3?”). If arbitrary questions are allowed then one can always create a maximally informative split at each stage of the search process, in which case stepwise information gain is also the most efficient strategy (Huffman, 1952). However, if arbitrary queries are not possible, stepwise-optimal methods can be distinctly suboptimal in the long run. Due to the limitations on queries (e.g., available features), it might not always be possible to ask a question that splits the remaining hypotheses exactly or close to 50:50. Under these circumstances, choosing efficient queries requires planning ahead by considering what questions would be available later on, given each possible answer to the current question.

Consider again Figure 1. Both color and shape induce a 3:9 split and therefore have the highest expected information gain (0.811 bits). However, one of them is uniquely optimal with respect to expected number of questions needed to definitively identify the target, because it allows for more informative queries on the second step. When querying the color feature, the most likely outcome is that the monster is blue, leaving nine candidate hypotheses. In this case, the most informative follow-up question queries the pattern feature, which induces a 2:7 split. By contrast, when querying the shape feature, the most likely outcome is that nine round monsters remain. In this case, color and pattern both pick out a single item (i.e., there are exactly one remaining green monster and one remaining spotted monster), and thus the searcher can only ask a follow-up question that queries individual monsters’ features.³

Computationally, determining the most efficient sequence of queries corresponds to finding the binary question tree with the shortest expected path length—the globally optimal solution that minimizes the expected number of queries. We define the *expected path length* for each question Q as the expected total number of questions needed to identify the target when a searcher begins with this question and chooses globally optimally on subsequent steps. As mentioned above, this problem is gener-

³If only three monsters remain after the first query, only hypothesis-scanning questions inducing a 1:2 split can be asked, and all questions have identical usefulness, with the expected number of additional questions needed being $1\frac{2}{3}$. For a formal proof see Nelson et al. (2018).

Table 1. *Initial Split, Expected Stepwise Information Gain, Expected Path Length (Expected Number of Questions) and Reduction Thereof for Features in Figure 1a*

Question	Feature	Stepwise optimality metrics		Global optimality metrics	
		Initial split	Information gain	Path length	Path length reduction
Q1	Color	3:9	0.811	4.583	1.000
Q2	Shape	3:9	0.811	5.083	0.500
Q3	Pattern	2:10	0.650	4.667	0.917
Q4–Q10	Monster M1–M7	1:11	0.414	4.833	0.750
Q11–Q15	Monster M8–M12	1:11	0.414	5.250	0.333
Q4–Q15	Random monster M1–M12	1:11	0.414	5.007	0.576

Note. Questions pertaining to individual monsters have different path lengths; “Random monster” gives the weighted mean across the 12 items (see Table A1 in Appendix A for details). In the behavioral experiments, assignment of queries Q1–Q3 to physical features of color, shape, and pattern were counterbalanced across subjects.

ally intractable (Hyafil & Rivest, 1976), because the number of possible trees grows superexponentially with the number of available queries. For a moderate number of queries and items, the optimal tree and expected path length of each question can be determined through dynamic programming (Nelson et al., 2018).

Given each question’s expected path length, a given question Q ’s efficiency can be defined in terms of its *expected path length reduction*, which is the reduction in the expected number of remaining questions after asking Q , compared to the expected total number of questions when starting with the globally optimal question (see Appendix B for details). By definition, the expected path length reduction for all questions ranges between 0 and 1, and the expected path length reduction for the globally optimal first question equals 1.

Table 1 contrasts questions’ usefulness for the 20-questions game in Figure 1a in terms of stepwise information gain, expected path length (expected number of questions) and reduction thereof, illustrating how the stepwise and global optimality metrics can diverge. First, the ordering given by expected path length is generally inconsistent with the information gain ordering. Second, while the color and shape questions tie for maximal information gain, the color feature is uniquely optimal because it has a lower expected path length. Surprisingly, in terms of expected path length the shape query is even worse than selecting one of the 12 individual monsters at random and asking about it. Third, this also illustrates that the common assumption that constraint-seeking questions are superior to hypothesis-scanning questions does not always hold (for related issues in the case of unequal priors see Nelson et al., 2018, and Ruggeri et al., 2017).

Figures 2b and 2c illustrate the divergence between stepwise-optimal and globally optimal strategies in sequential search, assessed via simulation of each strategy’s

behavior. The baseline is provided by a *random* strategy, which chooses in each time step with equal probability among all informative questions (i.e., excluding redundant questions targeting features possessed by all or no item, given what has been learned up to that point). The *stepwise-optimal* strategy chooses at each step the query with the highest stepwise information gain (ties are broken randomly), mimicking a learner who uses the split-half heuristic. The *globally optimal* strategy selects questions in accordance with their expected path length reduction (ties are again broken randomly); that is, it directly minimizes the objective function (expected number of queries). We used the environment shown in Figure 1a (see also Table A1), where in each round a random target was chosen and the strategy selected questions until it had been identified. Figure 2b shows the efficiency of the three strategies, illustrating how stepwise information gain falls short of the globally optimal strategy. Note that the globally optimal strategy also achieves higher average information gain (averaged over all steps in the game) than the stepwise-optimal strategy, which greedily maximizes this quantity in each step (Figure 2c).

2.4. Summary

In the 20-questions game, applying the split-half heuristic (i.e., querying a feature whose distribution comes as close as possible to a 50:50 split) enables searchers to identify the question with the highest stepwise information gain. However, reducing uncertainty in a stepwise-optimal fashion does not necessarily minimize the expected number of questions needed to identify the true target item. Our analyses illustrate that stepwise information gain can be distinctly suboptimal in the long run, as stepwise-optimal methods only consider a query’s implications for the immediate next time step and do not take into account how a query’s outcome can affect the informativeness of subsequent queries. This limitation is a function of both the search strategy and of the task structure. Indeed, if it is possible within the task to ask arbitrary questions (i.e., questions targeting any arbitrary subset of items), then maximizing stepwise information gain is also the most efficient strategy, because at each step of the inquiry a maximally informative split can be created. However, often the available features and their distributions limit the available queries, therefore potentially introducing an efficiency gap between stepwise-optimal and globally optimal strategies.

Our analyses show that task environments in which stepwise and globally optimal strategies make diverging predictions can be particularly informative regarding the efficiency of human search. Moreover, quantifying questions’ usefulness in terms of the expected reduction in path length provides a direct means to evaluate them with respect to the objective function that searchers are asked to minimize, namely solving

the game by asking as few questions as possible.

3. Goals and scope of experiments

OED models such as stepwise information gain have been applied to a variety of tasks in cognitive science investigating adult subjects' behavior, including perceptual tasks (Najemnik & Geisler, 2005), categorization (Meder & Nelson, 2012; Nelson et al., 2010; Wu et al., 2017), causal learning (Bramley et al., 2015; Steyvers et al., 2003), and hypothesis testing (Coenen et al., 2017; Navarro & Perfors, 2011; Oaksford & Chater, 1994, 1996). Similarly, recent developmental studies investigating children's search behavior in the 20-questions game from a computational perspective have focused exclusively on stepwise-optimal methods (Kachergis et al., 2017; Nelson et al., 2014; Ruggeri & Lombrozo, 2015; Ruggeri et al., 2017). Although this research has advanced our understanding of the behavioral and statistical principles underlying human information search, there has been virtually no consideration of scenarios in which stepwise methods are suboptimal with respect to search efficiency.

We conducted two behavioral experiments to investigate whether and to what extent children and adults plan ahead in sequential search, using a 20-questions game with the stimuli shown in Figure 1. This task environment directly contrasts stepwise methods with globally optimal strategies, addressing the question of whether searchers would consider possible subsequent queries in order to determine which of two stepwise-optimal queries is more efficient (Table 1). By comparing the behavior of children of different age groups with that of adults we also aimed to tap into the development of planning ahead and strategy use in sequential information acquisition.

What is known about people's capability to search efficiently when stepwise-optimal and more globally optimal methods make diverging predictions about queries' usefulness? Although the conflict between short- and long-run optimality has been recognized for decades in other domains, such as reward-based learning and problem solving (Bellman, 1957; Newell & Simon, 1959), there is a lack of research on this question in the context of pure information-acquisition tasks. Researchers have either used one-shot search tasks, in which only a single piece of information can be obtained (e.g., Nelson et al., 2010; Skov & Sherman, 1986; Slowiaczek et al., 1992; Trope & Bassok, 1982), or have focused on stepwise-optimal methods to evaluate human search behavior when multiple pieces of information can be obtained (Oaksford & Chater, 1994, 1996; Markant & Gureckis, 2014; Nelson et al., 2014; Najemnik & Geisler, 2005; Ruggeri & Lombrozo, 2015). A notable exception is the work by Meier and Blair (2013), who used a probabilistic multiple-cue categorization task to

assess adults' ability to learn the most efficient sequence of queries. They used an experience-based search paradigm comprising several hundred trials, where in each trial searchers could view up to three features of a stimulus before making a classification decision. The task environment was designed such that stepwise-optimal strategies required $2\frac{1}{6}$ queries on average, whereas the globally optimal strategy required exactly 2 queries. The key finding was that, over the course of learning, searchers acquired a systematic preference to first query the feature that ultimately led to higher efficiency, even though it was not the most informative feature according to stepwise-optimal methods. The learning curves show that searchers developed this preference only after extensive experience with the search problem, based on dozens or even hundreds of learning trials (see Supplementary Material of Meier & Blair, 2013). These findings demonstrate that adult subjects can harness their learning experience to identify the most efficient sequence of queries, even when stepwise-optimal methods fail to identify the best first question.

An important question that the literature to date has not addressed is whether searchers can identify an efficient series of queries in a sequential search task performed only once. This situation strongly differs from experience-based paradigms, in which searchers can repeatedly interact with the problem to learn about the best search strategy. This is a crucial distinction, because in experience-based tasks searchers could learn the most efficient query through processes other than explicit planning, for instance by reinforcement learning mechanisms (Sutton & Barto, 1998). In this case, the number of queries serves as loss function that the searcher seeks to minimize and the learned reward values determine which feature is queried first, without the searcher ever explicitly planning ahead or even necessarily recognizing why that feature is most efficient (see General Discussion for details). In the experiments we report, the search task was performed only once, and therefore to solve the problem efficiently searchers had to plan ahead by considering not only the immediate implications for the hypotheses considered but also the consequences for subsequent queries. This situation aligns closely with the psychological literature on planning and problem solving, where planning is often conceptualized as a search through a space of interconnected states, traversing from an initial state to a desired goal state by applying a sequence of operations that connect the intermediate states (Morris & Ward, 2004; Newell & Simon, 1972). In the 20-questions game used here, the initial state corresponded to not knowing which of the n monsters was the true target, the goal state was knowing with certainty the true target monster, and the operations that transformed the problem states were the questions that could be asked.

There is a rich literature showing that healthy adults can plan ahead in tasks with no uncertainty (e.g., Hayes-Roth & Hayes-Roth, 1979; Newell & Simon, 1972;

Shallice, 1982), such as the Tower of Hanoi and London tasks. Developmental studies have also focused on tasks with deterministic actions, such as Tower of Hanoi and London problems (Bull et al., 2004; Klahr & Robinson, 1981) or navigation and maze tasks (Gardner & Rogoff, 1990; Völter & Call, 2014). For instance, using a maze navigation task, Völter and Call (2014) found that 5-year-olds planned up to two steps ahead when selecting the entry point for the maze, whereas 4-year-olds tended to plan only one step ahead. This research has helped map developmental trajectories in planning. However, action selection and information foraging are psychologically distinct processes, and planning a sequence of actions in a task with no uncertainty is quite different from planning a tree of queries where the goal is to resolve uncertainty. In fact, computational analyses of when humans and other animals do or do not plan ahead suggest that task uncertainty is a primary factor discouraging planning (Daw et al., 2005).

A key goal of the present research was therefore to investigate children’s and adults’ planning in a task of pure information acquisition where outcomes are probabilistic. The 20-questions game provides an ideal experimental paradigm for addressing this question from a developmental perspective because the statistical structure of the task environment is much simpler than in probabilistic classification tasks with unequal priors and probabilistic likelihoods (e.g., the proportion of items possessing a particular feature value can be directly observed). This enables a direct test of how well children can plan ahead in a more complex probabilistic information search task, and how their behavior compares to that of adult searchers.

4. Experiment 1

4.1. *Participants and design*

Children ($N = 47$) and adults ($N = 50$) were recruited through the subject pool of the Max Planck Institute for Human Development in Berlin and tested individually in the laboratory. The study was approved by the Ethics Committee of the Max Planck Institute for Human Development and written consent was obtained from all participants (in the case of children from their legal guardian). Following previous research, we targeted children of about 10 years of age, an age that has been shown to be a critical stepping stone for question asking and information search (Mosher & Hornsby, 1966; Ruggeri & Feufel, 2015; Ruggeri & Lombrozo, 2015; Ruggeri et al., 2016). Children and adults were tested individually in separate sessions, which were videotaped and coded later. Families received a flat fee of €10 for participating and an average performance-based bonus of €2.79. Two children were excluded from the analyses because they indicated having accidentally seen the target item; one adult

was excluded because of missing consent. Table 2 shows participant demographics; all N s in figures and tables refer to the sample size after excluding subjects.

Table 2. *Participant demographics*

	Age group	N	Female	Age (years)		
				Mean	Median	SD
Exp. 1	10-year-olds	45	47%	10.40	10.33	1.24
Exp. 1	Adults	49	69%	43.94	44.92	5.87
Exp. 2	8-year-olds	52	48%	7.94	8.08	0.58
Exp. 2	10-year-olds	54	46%	10.40	10.25	0.80
Exp. 2	Adults	52	58%	35.01	35.12	11.32
Total	Children	151	47%	9.55	9.67	1.47
Total	Adults	101	63%	39.3	40.08	10.12

4.2. Task environment

We used computer simulations to identify a set of items and feature distributions in which stepwise information gain and equivalent strategies such as the split-half heuristic fail to identify the most efficient first query. More specifically, the goal of these simulations was to generate an environment in which two questions tie as most splithalfy, although only one of those questions is globally optimal, and there is a large difference between those two questions in their path length reduction. Two important analytical constraints guided the design of the task environment: If there are fewer than three constraint questions (i.e., features possessed by more than one item), or if there are fewer than eight items (hypotheses), then both stepwise information gain and the split-half heuristic invariably select a first question with maximal expected path length reduction (or, equivalently, minimal expected path length). Both constraints can be proven analytically (Nelson et al., 2018, Theorems 5 and 6), and they provide a lower bound on the size of any environment that contrasts stepwise and long-run optimality. Since our goal was to run the experiments with children of different ages, an additional pragmatic constraint was to keep the problem tractable. We therefore searched for environments with small numbers of items (up to about 18, the fewer the better) and small numbers of constraint features (up to about 6, the fewer the better).

The simulations identified an environment with 12 items and three constraint-seeking questions, in which two questions (Q1 and Q2) are equal in their splithalfiness

and stepwise information gain, but one of them is uniquely optimal in terms of efficiency: Q1 has an expected path length reduction of 1, whereas Q2 has an expected reduction of only 0.5 (Tables 1 and A1). The decision to have two questions with equal splithaleness was based on pilot studies indicating that adult subjects have a strong preference to ask the most splithalfy question even if that question is not the most efficient first query. We hoped that being presented with two stepwise-optimal questions with the same splithaleness would encourage participants to look ahead to break the tie. In addition, Q3 was designed to have a lower split than Q2 (2:10 vs. 3:9) but higher path length reduction (0.917 vs. 0.5), which enabled us to test whether considerations about efficiency can trump splithaleness. Accordingly, we chose as a dependent variable in the experiments which first question was selected, because it is a more precise measure than the actual number of questions required (which varies a great deal according to the luck of the draw in individual rounds of the game).

Given the analytical constraints stated above, this environment is the simplest possible for our purposes in terms of the number of available constraint questions. In terms of the number of items, the simulations showed that having fewer items, although mathematically possible, would come at the cost of a smaller difference in efficiency between stepwise and globally optimal questions. Furthermore, 12 items is toward the lower range of the number of stimuli used in developmental studies to date, suggesting that it should be doable for children. Note also that the difference in path length reduction between Q1 and Q2 is quite large. The maximum possible difference is 1, which occurs only when the inferior question is completely uninformative. The difference of 0.5 in the present design implies that Q2 is only half as useful as Q1 in terms of progress toward completing the task. The difference in the expected number of queries in our environment (0.5) is also three times larger than in the experience-based search task used by Meier and Blair (2013), where the difference was $1/6$ (with $2\frac{1}{6}$ queries for the stepwise-optimal strategy vs. 2 queries for the globally optimal strategy). Because Meier and Blair found an effect of efficiency in their task, the present design should be more than adequate to detect whether people are sensitive to long-run efficiency in a 20-questions task.

4.3. Procedure and materials

The materials consisted of decks of cards depicting 12 monsters (Figure 1a; Table A1 shows the canonical item–feature matrix underlying our experiments). We created six different sets of cards by counterbalancing the assignment of queries Q1 to Q3 to the physical features of color, shape, and pattern. For each subject, we used two identical decks of cards and a “magic machine” (operated by the experimenter)

that would produce music when two identical monster cards were placed inside. The experimenter demonstrated how the machine works by placing two identical monster cards in it (these monsters were created only for that purpose and were not part of the monster sets used for the game).

Each subject played one round of the 20-questions game (see Appendix D for the complete experimental script). At the beginning of the game, one deck was shuffled and randomly arranged in a 3×4 grid on a table. To ensure that subjects—especially children—knew the feature distributions, they were asked to sort the cards according to the induced split for each feature, in random order, and count the number of monsters possessing each feature value. This exercise also served the purpose of highlighting which questions would be available on the second step if that feature was queried.

Subsequently, the second deck of monster cards was shuffled three times by the experimenter, who then drew a random monster card and placed it in the machine without revealing it to the subject. Subjects were instructed to identify the matching monster in the first deck by asking as few yes–no questions as possible. To incentivize efficient search, subjects had to pay a €50-cent coin for each question, from an initial endowment of ten €50-cent coins. At each stage of the search process, the monsters ruled out by the previous question were removed from the grid by the experimenter. When only one monster remained, it was placed into the magic machine, which then activated.

For the first two queries, subjects could choose among three constraint-seeking questions (Q1, Q2, Q3; randomly assigned to the shape, color, and pattern features) or query an individual monster. We refer to this as the *question-selection phase*. Allowed queries were visualized using four question cards (Figure 1b). Importantly, the available questions strongly differed in terms of expected path length reduction, with one of the two 3:9 questions being uniquely optimal (Q1, see Table 1). From Step 3 onward, subjects could ask arbitrary questions (e.g., “Does the monster have something on his head?”). We refer to this as the *question-generation phase*. While our primary focus is on selection of the first question, this phase was included so that children could finish playing the game; additionally, it provides information on age-related differences in the capability to generate useful questions. Note that participants were not initially informed that from the third query onward, they would be allowed to ask arbitrary questions. This was done to motivate subjects to start the search process with the most efficient query.

4.4. Results

4.4.1. Model performance

Because searchers could ask arbitrary questions after the first two queries, we implemented corresponding variants of the random, stepwise-optimal, and globally optimal models. For the first two queries, the random strategy selects randomly among the allowed queries, the stepwise-optimal strategy chooses the maximum information gain query (ties are broken randomly), and the globally optimal strategy selects questions according to expected path length reduction under the assumption that only predefined queries will be allowed (reflecting subjects' lack of foreknowledge about the question-generation phase).

The subsequent question-generation phase was modeled by creating, at each step of the search, an arbitrary instance of each possible split given the remaining items (e.g., with eight items there are four possible unique and informative splits, namely 1:7, 2:6, 3:5, and 4:4). In the game, such questions could correspond to arbitrary subsets (e.g., "Is it the monster with the pipe or the monster with the wings?") or queries pertaining to higher level features (e.g., "Does the monster have something on its head?"). The random strategy selects among the possible splits with equal probability (ignoring the difference in number of possible instantiations of each split), whereas the stepwise-optimal and globally optimal strategies select the most splithalfy question. Note that because arbitrary questions are allowed in this phase, the stepwise-optimal strategy is also the globally optimal strategy, such that both models achieve maximum stepwise information gain and expected path length reduction during this phase.

Figure 3a shows the average number of items left at each stage of the search process for each model. The baseline is provided by the random strategy. Throughout the search process, the globally optimal strategy eliminates more items on average than the stepwise-optimal strategy. Strategies' overall efficiencies are shown in Figure 4a (horizontal lines). Note that the gap between the stepwise and global models is reduced because of the question-generation phase; the difference in path length reduction over the first two steps is greater. Performance in terms of stepwise information gain is shown in Figure 5. The globally optimal strategy achieves higher information gain than the stepwise-optimal strategy because it selects a more efficient question on the first step, which enables higher information gain queries subsequently.

4.4.2. Efficiency of search

Figure 3b shows how many items remained at each stage of the search process for the subjects. Comparison of the two age groups shows that on average children

searched as efficiently as adults. Figure 4a shows the number of questions needed by each group to identify the target item. Children searched as efficiently as adults; in fact, they needed slightly fewer questions than adults. The difference in the number of questions was not significant though ($M_{\text{children}} = 4.18, 95\% \text{ CI} = [3.69, 4.66]$ vs. $M_{\text{adults}} = 4.65, 95\% \text{ CI} = [4.19, 5.12], z = -1.47, p = .14$, Wilcoxon–Mann–Whitney test).

Figure 3. Efficiency of search process. For each time step the average number of items remaining is shown, with subjects (and model runs) who already identified the target treated as having one item remaining. Model performance is based on 10^5 simulations for each target item.

4.4.3. Selection of first question

Our primary dependent variable of interest was which question searchers asked at the beginning of the game. Figure 4b and c shows the selection proportions for the question on the first step.⁴ Choices for the first question differed from chance for both adults ($\chi^2(3, N = 49) = 25.5, p < .0001, w = .72$), and children ($\chi^2(3, N = 45) = 22.6, p < .0001, w = .71$), with little difference between age groups ($p = .09$, Fisher’s exact test).

No searcher chose a hypothesis-scanning questions as first query. Considering only searchers who initially asked one of the three constraint questions, both children’s ($\chi^2(2, N = 45) = 5.73, p = .06, w = .36$) and adults’ ($\chi^2(2, N = 49) = 6.9, p = .03, w = .38$) selections differed from choosing randomly among them. Both age groups showed a preference for the two most splithalfy questions (3:9 split, Q1 and Q2), which together accounted for 75.5% of choices in adults (24.5% Q1, 51% Q2)

⁴We treat the 12 hypothesis-scanning questions pertaining to individual monsters as one question *type*, rather than individual queries, such that there are four possible first queries.

and 82.2% in children (46.7% Q1, 35.6% Q2). For children this proportion differed from randomly choosing among the three constraint questions ($p = .03$, binomial test against 2/3 chance level), but not for adults ($p = .23$). However, there was no difference between age groups ($p = .46$, Fisher’s exact test).

Considering only subjects who started out with one of the two 3:9 questions (37 children and 37 adults), there was a marginal difference between age groups ($\chi^2(1, N = 74) = 3.5, p < .06, \phi = .22$). Children’s preference for the more efficient 3:9 question, Q1, did not differ from chance ($p = .51$, binomial test), but adults were more likely to select the less efficient 3:9 question, Q2 ($p = .05$).

These results indicate that children’s efficiency was at least on par with adults’, but generally both children and adults showed little to no sensitivity to planning ahead or to questions’ efficiency. Question selection seemed to be primarily determined by the induced split and stepwise information gain (for children) or by arbitrary selection among constraint questions (for adults).

Figure 4. Efficiency (a) and selection proportions for first question (b, c) in Experiment 1. a) Number of questions needed to identify the true target item. Horizontal lines indicate mean performance of models, box plots show interquartile range with median (line), mean (diamond), and range (whiskers). b) and c) First questions chosen by children and adults. Questions ordered by path length reduction, with the expected reduction of each question shown in square brackets. Numbers on top of bars show induced splits.

4.4.4. Information gain analysis

We computed for each subject the mean stepwise information gain achieved in the question-selection phase (initial two questions, Figure 5a), in the question-generation phase (all but the first two questions, Figure 5b), and across all questions (Figure 5c). There were no differences between groups for the first two questions ($M_{\text{children}} = 0.72, 95\% \text{ CI} = [0.7, 0.75]$ vs. $M_{\text{adults}} = 0.73, 95\% \text{ CI} = [0.7, 0.75], z = 0.21, p = .21$;

here and in the following Wilcoxon-Mann-Whitney test). However, in the question-generation phase (Figure 5b), where stepwise information gain is also the most efficient strategy, adults tended to generate more informative questions than children ($M_{\text{children}} = 0.86, 95\% \text{ CI} = [0.81, 0.91]$ vs. $M_{\text{adults}} = 0.92, 95\% \text{ CI} = [0.88, 0.95]$, $z = -1.83$, $p = .07$, $r = .19$).

The developmental difference in question generation is consistent with recent studies showing that children can *select* the most informative of two questions already at around age 5 (Ruggeri et al., 2017) but still have difficulties *generating* informative questions at around age 10 (Ruggeri & Feufel, 2015; Ruggeri & Lombrozo, 2015; Ruggeri et al., 2017). Across all questions, adults also achieved higher information gain than children ($M_{\text{children}} = 0.79, 95\% \text{ CI} = [0.76, 0.82]$ vs. $M_{\text{adults}} = 0.83, 95\% \text{ CI} = [0.81, 0.85]$, $z = -2.33$, $p = .02$, $r = .24$).

Figure 5. Stepwise information gain in Experiment 1. *a*) Mean information gain in the question-selection phase (first two questions). *b*) Mean information gain in the question-generation phase (all but first two questions). *c*) Mean information gain of all questions. Each dot is one subject; box plots show interquartile range with median (line), mean (diamond), and range (whiskers). Horizontal lines show average information gain of random, stepwise-optimal, and globally optimal strategy.

5. Experiment 2

Three key findings emerged from Experiment 1. First, 10-year-olds searched as efficiently as adult subjects. Second, regardless of age, question selection seemed to be primarily driven by the induced split; even in our task where two questions tied for the highest initial split, searchers did not break the tie in favor of path length reduction. Third, in the question-generation phase, adults performed better than children, consistent with related findings from the developmental literature (Ruggeri & Feufel, 2015; Ruggeri & Lombrozo, 2015; Ruggeri et al., 2017).

The purpose of Experiment 2 was twofold. First, we introduced another test phase to assess more explicitly and directly whether children and adults are sensitive to questions’ efficiency. The game play was as in Experiment 1. However, prior to actually playing the game, subjects were presented with only two queries at a time (e.g., Q1 and Q2, both of which induce a 3:9 split, but Q1 has twice the expected path length reduction) and were asked to judge which one was better. Second, we additionally recruited children around age 8, to see whether even younger children would search as efficiently as older children and adults in our task. We did not aim to test even younger children though, because it has been found that children at this age struggle with generating useful questions (Legare et al., 2013; Ruggeri et al., 2018), and the current paradigm might have been too complicated for them.

5.1. *Participants and design*

We tested younger children around age 8 years ($N = 59$), older children around age 10 years ($N = 61$), and adults ($N = 53$). Table 2 provides detailed information on participant demographics. The Ethics Committee of the Max Planck Institute for Human Development approved the study and consent was obtained from all participants (in the case of children from their legal guardian). Subjects were recruited and tested at the Museum für Naturkunde (Natural History Museum) in Berlin and received an average performance-based payoff of €2.83. Note that it is not possible to know a child’s exact age when approaching them in a museum, so it was not always possible to target children who were close to a very specific age (although we did request and record exact ages once they were recruited). We address this issue below by conducting an aggregate analysis that treats age as a continuous variable (Section 6). All subjects were tested individually; sessions were videotaped and coded later. Fifteen subjects were excluded from the analyses: Two children were excluded because of missing consent, four children indicated having seen the target item, and eight children and one adult were excluded because the camera was not working properly.

5.2. *Procedure and materials*

The materials and procedure were largely identical to those of Experiment 1. As before, subjects first were introduced to the game and its goal, and then they sorted monster cards according to the splits induced by the four possible first questions (including querying an arbitrary individual monster). The only difference was that prior to actually playing the game, participants were presented with a series of pairwise choices among possible first questions.

For each comparison, the participant was asked to imagine that two children were playing the game, each of them selecting one of the two presented questions

to start with (see Appendix E for the script used). Queries were visualized using the corresponding question cards (Figure 1b). Subjects were asked to indicate who they thought would identify the true target item faster (e.g., the child starting with Q1 or the child starting with Q2). The primary purpose of this additional test was to directly pit critical queries against each other, namely Q1 versus Q2 (they tie for maximum information gain, but Q1 has higher path length reduction) and Q2 versus Q3 (Q2 has higher information gain but lower path length reduction than Q3). In addition, we hypothesized that presenting pairwise choices prior to playing the actual game might highlight questions' varying efficiencies.

Each subject answered all three pairwise comparisons among Q1, Q2, and Q3, as well as two pairwise comparisons between the most efficient query, Q1, and a randomly chosen hypothesis-scanning question (from the set Q4–Q10 and Q11–Q15, respectively, which differ in their efficiency; see Table 1). Comparisons were conducted in random order.

5.3. Results

5.3.1. Pairwise choices

Figure 6 shows participants' choices for the different pairs of queries. (Due to errors in data collection, 6 of 790 comparisons are missing.) There are two critical comparisons for which path length reduction and stepwise information gain make diverging predictions. The first comparison is Q1 versus Q2, which are both stepwise optimal but Q1 is globally optimal. None of the age groups differed from chance in their choices (all $ps > .21$, binomial tests); there was also no difference among the three age groups ($p = .17$, Fisher's exact test).

The other critical comparison concerns Q2 (split 3:9) versus Q3 (split 2:10), with the former being stepwise optimal but the latter being more efficient in terms of expected path length reduction. Adults and 10-year-olds had a slight preference for the less efficient but more splithalfy question, although the selection proportions did not differ from chance ($p = .13$ and $p = .07$, respectively, binomial tests). Eight-year-olds were indifferent ($p = .58$). Choice patterns did not differ between age groups though ($p = .13$, Fisher's exact test). Taken together, these results indicate that searchers were not sensitive to questions' long-run efficiency, even when only evaluating two queries at once.

A baseline measure of sensitivity is provided by the comparison of Q1 versus Q3, with the former having higher splithalfness and also being optimal with respect to path length reduction. Interestingly, only adults showed a clear preference for the better query ($p = .001$, binomial test), whereas children did not differ from chance (both $ps > .12$). Consequently, selection frequencies differed among age groups

Figure 6. Pairwise choices among questions in Experiment 2. For each pair, the more efficient query is shown on the bottom. IG = stepwise information gain, PLR = path length reduction.

($p = .002$, Fisher's exact test).

The final analysis concerns the two comparisons involving the globally optimal query, Q1, against a randomly chosen hypothesis-scanning question pertaining to an individual monster from the set Q4–Q10, and Q1 versus a randomly chosen monster from the set Q11–Q15. Regardless of age, subjects had a strong preference for the constraint-seeking question, Q1, over the hypothesis-scanning questions (all $ps < .01$, binomial tests). Consistent with the literature (Mosher & Hornsby, 1966; Ruggeri & Feufel, 2015; Ruggeri & Lombrozo, 2015), we also observed an age-related trend regarding the selection of hypothesis-scanning questions. Aggregating across the two comparisons, adult subjects endorsed querying an individual monster 4.8% of the time, whereas 10- and 8-year-olds did so 15.7% and 28.2% of the time, respectively. Accordingly, choice proportions differed among age groups for both comparisons ($p = .0003$ and $p = .03$, respectively, Fisher's exact test).

5.3.2. Efficiency of search

As in Experiment 1, overall children searched as efficiently as adults; in fact, 8-year-olds performed best in terms of mean and median number of questions (Figures 3b and 7a). The difference among age groups was not reliable though ($\chi^2 = 1.89$, $p = .39$, $df = 2$, Kruskal–Wallis test).

5.3.3. Selection of first question

Choices for the first question (Figure 7b–d) clearly differed from chance for adults ($\chi^2(3, N = 52) = 26.3$, $p < .0001$, $w = .71$) and 10-year-olds, ($\chi^2(3, N = 54) = 18.4$, $p = .0004$, $w = .58$); choices of 8-year-olds were more noisy ($\chi^2(3, N = 52) =$

7.2, $p = .06$, $w = .37$). Selection frequencies did not differ among age groups ($p = .16$, Fisher's exact test). In contrast to Experiment 1, where 10-year-olds but not adults showed a tendency to preferably select the most efficient query, this was not the case in Experiment 2; in the context of the other findings we therefore assume that this observation was just noise.

No adult chose a hypothesis-scanning question as first query, one of the 10-year-olds did, and five of the 8-year-olds did. Considering only searchers who started out with a constraint-seeking question, only adults' choices among Q1, Q2, and Q3 differed from chance ($\chi^2(2, N = 52) = 6.73$, $p = .03$, $w = .36$), whereas this was not the case for 8- and 10-year-olds ($p = .76$ and $p = .32$, respectively). The overall pattern did not differ between groups though ($p = .59$, Fisher's exact test).

Choice proportions show an age-related trend regarding question selection in accordance with their splithaleness. The rate of selecting one of the two 3:9-split questions was 61.5% for 8-year-olds (34.6% Q1, 26.9% Q2), 74.1% for 10-year-olds (40.7% Q1, 33.3% Q2), and 82.7% for adults (46.2% Q1, 36.5% Q2). Preference for a 3:9 split question over the other questions differed among age groups ($\chi^2(2, N = 158) = 5.9$, $p = .05$, $w = .19$). Comparing these choice proportions with randomly choosing among the three constraint questions showed that only adults differed from chance ($p = .01$, binomial test against 2/3 chance, vs. $p = .46$ and $p = .31$ for 8- and 10-year-olds, respectively).

The final analysis considers only searchers who initially asked one of the two 3:9 split questions (32 8-year-olds, 38 10-year-olds, and 42 adults). There were no differences in rate of selecting the more efficient question among the age groups ($p \approx 1$, Fisher's exact test), and choice of the more efficient 3:9 question did not differ from chance for any of the age groups (all $ps > .54$, binomial test), or when aggregating across all subjects ($p = .26$). These results indicate that regardless of age, searchers did not distinguish between more and less efficient queries that induce the same split.

5.3.4. Information gain analysis

Figure 8 shows stepwise information gain by age group and phase of the game. There were no differences among age groups in the question-selection phase ($\chi^2=0.69$, $p = .71$, $df = 2$, Kruskal-Wallis test). In the question-generation phase, adults asked slightly more informative questions than younger and older children, but differences among age groups were not significant ($\chi^2=2.96$, $p = .23$, $df = 2$). There were also no statistical differences when considering the average information gain of all questions ($\chi^2=4.66$, $p = .1$, $df = 2$).

Figure 7. Efficiency and selection proportions for first question in Experiment 2. *a)* Number of questions needed to identify the true target item. Each dot is one subject; box plots show interquartile range with median (line), mean (diamond), and range (whiskers). Horizontal lines show average performance of random, stepwise-optimal, and globally-optimal models. *b–d)* First questions chosen by 8-year-olds, 10-year-olds, and adults. Questions ordered by path length reduction, with the expected reduction of each question shown in square brackets. Numbers on top of bars show induced splits.

Figure 8. Information gain analysis of Experiment 2. *a)* Mean information gain in the question-selection phase (first two questions). *b)* Mean information gain in the question-generation phase (all but first two questions). *c)* Mean information gain of all questions. Each dot is one subject; box plots show interquartile range with median (line), mean (diamond), and range of data (whiskers). Horizontal lines show average information gain of random, stepwise-optimal, and globally optimal models.

6. Aggregate analyses

To further assess age-related differences in search we considered all subjects from both experiments (Table 2) and treated children’s age as a continuous variable. Figure 9a shows children’s search efficiency. There was no relation between children’s age and the number of questions needed to identify the target item ($r = .004$; Spearman rank correlation here and in the following). On average adults required slightly more

questions than children, but the difference was not significant ($M_{adults} = 4.47$, 95% CI = [4.17, 4.76]) vs. $M_{children} = 4.36$, 95% CI = [4.09, 4.64], $z = 0.73$, $p = .46$, Wilcoxon–Mann–Whitney test). There was no correlation between children’s age and the number of questions asked ($r = .004$).

Figure 9b plots the expected path length reduction of the first question as a function of age. Again children performed as well as adults ($M_{adults} = 0.77$, 95% CI = [0.72, 0.81] vs. $M_{children} = 0.80$, 95% CI = [0.76, 0.84], $z = 1.02$, $p = .31$), with no relation between children’s age and path length reduction ($r = .02$).

Figure 9c shows the mean information gain in the question-generation phase. For this analysis, we considered only games in which there were more than three items remaining at the beginning of this phase (65% of children and 68% of adults), as otherwise only hypothesis-testing questions could be asked.⁵ The mean information gain of children was 0.84 (95% CI = [0.81, 0.87]) and of adults 0.91 (95% CI = [0.88, 0.93]), showing that on average adults generated more informative questions than children ($z = 3.77$, $p = .0002$; Wilcoxon–Mann–Whitney test). There was also a relation between age and children’s ability to generate informative questions ($r = .27$, $p = .008$; the solid line in Figure 9c shows a linear regression).

Figure 9. Aggregate analyses combining Experiments 1 and 2. Each dot is one child, the dashed line shows the mean across all adult subjects, and r denotes Spearman rank correlation. *a*) Number of questions. *b*) Expected path length reduction of first question. *c*) Mean information gain in the question-generation phase; the solid line shows a linear regression with children’s age as predictor.

⁵If there are only two or three items, in each case all informative questions have identical information gain and path length reduction (because the only informative splits are 2:1 and 1:1, respectively) and therefore query choices are not indicative of people’s ability to generate good questions.

7. General discussion

The present research assessed to what extent people plan ahead in the 20-questions game. Children searched as efficiently, but also as myopically, as adults. Thus, regardless of age, searchers were not sensitive to achieving efficiency through planning ahead and considering what questions could be available subsequently. Instead, question selection in both children and adults seemed to be primarily driven by the splithaleness of questions and their corresponding stepwise information gain. In this sense, our results are the first suggesting that stepwise-optimal methods better account for human behavior, not only in situations where there is no conflict between stepwise-optimal question selection and long-run efficiency (Denney & Denney, 1973; Eimas, 1970; Nelson et al., 2014; Ruggeri & Lombrozo, 2015; Ruggeri et al., 2017; Siegler, 1977; Thornton, 1982), but also when they are *suboptimal* and fail to identify the most efficient first query.

7.1. *Experience-based search, reinforcement learning, and planning an efficient sequence of queries*

The ability to obtain information in the most efficient manner might strongly depend on learning and experience. The findings of Meier and Blair (2013) demonstrate that adult searchers can achieve high efficiency when given the opportunity to repeatedly interact with the task and observe the outcomes of their choices over hundreds of trials.

Experience-based learning could shape people’s search behavior in at least two different ways. At a metacognitive level, searchers might learn from experience the strategic value of planning ahead. In the same manner that a novice chess player needs to learn, through experience or instruction, that improving performance requires looking more than one step ahead, information foragers may need to learn that stepwise methods can sometimes be inefficient. For instance, Meier and Blair (2013, Supplemental Material) in their analysis of searchers’ learning curves, noted that “while it appears in the aggregate that a gradual shift to F1-first trials occurs [where F1 refers to the feature that was most efficient to query first], this generally reflects the tendency for participants to adopt primarily F1-first strategies at different points in time” (p. 7). The absence of a gradual shift, at the individual level, towards searching the most efficient feature may indicate metacognitive learning, for instance in the form of insight effects. In the present task, this would mean to realize, based on playing the experimental game repeatedly, that one of the two stepwise-optimal queries is more efficient because it allows for more informative questions on the second step.

At a task-specific level, experience-based learning might impact search strategies by enabling people to discover the long-run value of different questions through experience. For instance, if given the opportunity to play the present 20-questions game repeatedly, searchers might learn that Q1 is a better initial query than Q2 just by observing the follow-up questions available for each outcome and adapt their behavior accordingly. Just like the temporal-difference algorithm in reinforcement learning (Sutton, 1988; Watkins & Dayan, 1992; Witten, 1977), this sort of learning over repeated play could lead people to choose in accordance with long-run efficiency without explicitly planning ahead within any single run through the task.

An important area for future research is therefore to study the interplay between learning and search in situations in which people experience a task multiple times. The reinforcement learning framework (Sutton & Barto, 1998) offers several approaches for modeling experience-based learning and has been widely used to develop strategies and learning mechanisms to tackle the conflict between short- and long-run optimality. Although this line of research has been primarily concerned with reward-based decision making, extending it to pure information search tasks would be fairly straightforward. If the reward is defined as -1 on every step until the game is completed, then to maximize total reward (or, equivalently, to minimize loss) the agent needs to identify the target with the fewest number of queries. Alternatively, by using informational reward functions, it is possible to incorporate the quantities underpinning different OED models, such as entropy reduction in the case of information gain. Stepwise OED models can be implemented by valuing only immediate rewards, which in the reinforcement learning framework is achieved by setting the temporal discount factor to zero. Models with a longer planning horizon can be obtained by increasing the temporal discount parameter such that later steps are also valued (e.g., Butko & Movellan, 2010). Such research would connect hitherto separate literatures by integrating reinforcement learning mechanisms with formal models of the value of information, as well as highlighting connections to multistep OED models that have a longer planning horizon (Nelson et al., 2018).

7.2. *One step at a time: Normative issues and stepwise information acquisition*

Stepwise methods are widely used in different fields, including, but not limited to, psychology and cognitive science. For instance, stepwise (greedy) methods are an integral part of many important machine learning approaches, such as classification and regression trees (Breiman et al., 1984; Quinlan, 1986), Gaussian Process regression (Rasmussen & Williams, 2006; Wu et al., 2018), and active learning (Settles, 2010). One argument for relying on such approaches is that although they are not necessarily globally optimal, they in many cases yield good performance with low

computational costs.

To better understand the performance of stepwise information gain and the split-half heuristic in the 20-questions game, we ran a series of computer simulations. Using the methodology of Nelson et al. (2018), we generated a wide range of environments with varying numbers of possible questions and items. Specifically, we varied three factors: the number of items (12, 15, 18), the number of constraint questions (2, 3, 5, 7, 9), and the mean feature density (splithaleness) of the features targeted by the constraint questions (0.1, 0.2, 0.3, 0.4, 0.5; see Appendix C for details). For each environment, we tracked whether the first question with the highest information gain was also the most efficient first query in terms of expected path length reduction, and how much loss was incurred by stepwise information gain (absolute and relative increase in the expected number of questions relative to the globally optimal query).

Figure 10a shows how often stepwise optimal information gain identified the most efficient first query; Figure 10b shows the total costs incurred in terms of increase in the expected number of queries. Overall, the simulations led to three key findings. First, if there are only two constraint questions, stepwise information gain always identifies the most efficient query. In fact, it can be proven analytically that this result holds regardless of the number of items and feature densities (Nelson et al., 2018). Second, performance of the stepwise strategy decreases when the environment becomes more complex, that is, when the number of hypotheses (items) and constraint questions increases. Third, although the agreement between the stepwise and globally optimal strategies decreases to about 80% in some types of environments, the incurred costs are rather low. Aggregating across all environments, stepwise information gain chose the globally optimal query 93.8% of the time, with a mean expected increase of 0.005 ($SD = 0.026$) queries; the relative increase in path length was about 0.13% ($SD = 0.6$). These findings suggest that selecting queries in accordance with their splithaleness, thereby reducing uncertainty in a stepwise-optimal fashion, is a reasonable and robust strategy in a wide range of environments. If people pick up on this pattern in their natural environment, they may come to rely on that heuristic even in cases (such as those in the present experiments) where it performs worse. The computational complexity of identifying the globally optimal solution makes it infeasible to use computers to evaluate model performance in environments with large numbers of hypotheses. This same consideration supports humans' use of stepwise-optimal methods: they are computationally much simpler and are also applicable in situations in which the globally optimal solution is intractable.

The question of how efficient greedy methods such as stepwise information gain can be has been a topic of research in computer science, where the focus has been

Figure 10. Performance of stepwise information gain in different environments. *a)* Proportion of times the splithalfiest (maximum stepwise information gain) question is also the most efficient (globally optimal) first query. *b)* Cost incurred by selecting the first query stepwise optimally, defined as the absolute increase in the expected number of queries relative to the globally optimal query. Each tile shows the mean across 1,000 randomly generated environments. See Nelson et al. (2018) for a more comprehensive set of simulations.

on the mathematical construct of submodular functions (Nemhauser et al., 1978). This work has led to provable upper bounds on the number of questions required by a greedy strategy in information-acquisition tasks (Golovin & Krause, 2011). For the 20-questions game, these results imply that the expected number of questions required by stepwise information gain is no more than $\ln n + 1$ times that of the globally optimal strategy. However, these analytic results are asymptotic in the number of items (hypotheses), and the bounds for smaller environment sizes such as those considered here can be shown to be significantly tighter. For example, the expected path length of a tree using only hypothesis-scanning questions (i.e., targeting an individual item at each step) provides a tighter bound on the efficiency of stepwise strategies than do the submodularity results, when the number of items is sufficiently small. This is one advantage of the present approach of focusing on

subjects' selections of single questions (mainly on the first step), rather than on their full sequences of choices. The expected path length reduction of a question must lie between 0 and 1, giving a simple and universally interpretable measure of efficiency of each choice.

Another perspective on the relative performance of stepwise and planning-based methods comes from computational work on the distinction between model-free and model-based reinforcement learning. Daw et al. (2005) developed an influential dual-process model, wherein behavior can be controlled either by a system that explicitly plans ahead (model-based) or by a stepwise (model-free) system that considers only the next action together with cached values for its possible outcomes. In the domain of information search, an action corresponds to a query, and the outcome of an action is an updated belief state that incorporates the answer to the query. In the 20-questions game, a belief state is simply the set of remaining candidate items (i.e., the items that satisfy all answers received thus far). Absent the opportunity to learn from repeated experience, the most natural value to assign to such a state is the (inverse or negative) number of items. Thus model-free decision making maps onto the split-half heuristic: evaluate each question solely in terms of the number of items it could leave on the next step. In their normative analysis of factors that should encourage lookahead versus stepwise control, Daw et al. (2005) showed that stochasticity or uncertainty in the task environment reduces the effectiveness of planning ahead, due to the computational error that accumulates when searching through a branching tree of possible outcomes. The principle is similar (though not identical) to the bias-variance trade-off (Geman et al., 1992): Stepwise control is biased because it ignores considerations of future steps, but it is less noisy. By its nature, information search involves uncertainty and stochastic outcomes (i.e., the answer to a question is unknown before it is asked). Thus Daw et al.'s analysis as applied to information search would suggest that planning ahead incurs too much cost in computational noise as compared to methods such as stepwise OED models and the split-half heuristic.

From a normative perspective, acquiring information in a stepwise way could be particularly harmful when external payoffs apply. For instance, in medical diagnosis, tests are often expensive in terms of monetary costs, can incur harmful side effects, or are time sensitive, so that it is vital to acquire relevant information in an efficient manner. In such cases, multistep methods that plan more than one query ahead can aid search and inference, with the applicable costs and benefits outweighing the additional computational costs.

7.3. *Developmental trajectories: From 20-questions games to active learning*

Twenty-questions games have been widely used in the developmental literature on active learning and have helped researchers map developmental trajectories in children’s information-acquisition strategies (Herwig, 1982; Mosher & Hornsby, 1966; Nelson et al., 2014; Ruggeri & Feufel, 2015; Ruggeri & Lombrozo, 2015; Ruggeri et al., 2016). Only recently have developmental and cognitive researchers gone beyond examining the sometimes not too surprising developmental trends (e.g., that adults tend to generate more informative questions than younger children), by using cognitive modeling techniques to investigate the computational principles underlying the development of search and question asking. This deeper level of analysis has revealed that 5-year-olds can identify the most informative of two given questions (Ruggeri et al., 2017); that 7-year-olds can adapt their question-asking behavior in accordance with the prior probability of different hypotheses; and that adults do not necessarily adapt their search strategies more promptly than children do (Ruggeri & Lombrozo, 2015). Our present results contribute to this literature by showing that adults are not necessarily better than 7- and 10-year-old children at planning ahead when asking questions, thereby highlighting the value of cognitive modeling and computational analysis of the task environment for developmental theorizing.

Future research should place a stronger emphasis on investigating the development of planning capabilities in situations characterized by different types and degrees of uncertainty. The focus of research to date has been on deterministic tasks (e.g., Towers of London and Hanoi or maze-navigation problems). This contrasts with the probabilistic outcomes of many real-world situations, which require taking multiple possible outcomes into account. Another important set of issues, in the context of experience-based learning tasks where there is a conflict between stepwise and globally optimal methods, is how speed of learning varies as a function of age, whether learning takes place in the form of a gradual transition or is punctuated with insight effects, and how learning and generalization change across the lifespan (Schulz et al., 2018). Finally, future research should explore how to *teach* children and adults to become good information foragers. Previous attempts to improve children’s question-asking strategies, for example, by providing explicit instructions or examples of adults asking informative questions, have had only moderate success (e.g., Courage, 1989; Denney, 1972; Denney et al., 1973; Denney & Turner, 1979). The effects did not generalize to other sets of stimuli and tended to be short-lived. We believe that a computational analysis of children’s and adults’ question-asking strategies and information search behavior, along the lines of the work presented in this paper, could greatly inform the design of more effective, individualized active-learning training programs and interventions.

7.4. Concluding remarks

In recent years, a strong interest in human information acquisition and active learning has emerged across different fields, including developmental psychology (Bonawitz et al., 2014; Nelson et al., 2014; Ruggeri et al., 2016, 2017), visual perception (Najemnik & Geisler, 2005; Nelson & Cottrell, 2007), higher level cognition (Coenen et al., 2018; Gureckis & Markant, 2012; Nelson, 2005; Markant & Gureckis, 2014; Meder & Nelson, 2012), and cognitive neuroscience (Schulz & Gershman, 2019; Zajkowski et al., 2017). These studies have advanced understanding of the behavioral and computational principles that guide human information acquisition and have uncovered key relationships between more complex statistical models of the value of information and simple heuristics for information search.

However, research to date has almost exclusively relied on stepwise-optimal methods for quantifying questions’ usefulness and evaluating human search behavior. One exception is Bramley et al. (2015), who investigated human behavior in a causal reasoning task where subjects could actively intervene in a causal system to learn about its structure. Although the study was not designed to directly contrast stepwise with longer run planning-based strategies, results showed that models with two-step-ahead planning horizons did not predict human behavior as well as the corresponding stepwise models.

The present research was motivated by the observation that stepwise models can be inadequate in situations where multiple queries can be conducted (Nelson et al., 2018). Consideration of such situations in future research is critical for both the descriptive and normative analysis of human information acquisition. For instance, whereas Meier and Blair (2013) found that after extensive experience people were sensitive to long-run efficiency considerations in a probabilistic classification task, the present studies indicate that both children and adults have difficulties planning ahead in the seemingly simple task of the 20-questions game, at least in their first encounter with the environment. This discrepancy calls for more research to better understand and characterize the conditions under which searchers can plan an efficient sequence of queries, and the cognitive processes underlying sequential search in different situations.

Acknowledgments

This research was supported by grants ME 3717/2-2 to BM and NE 1713/1-2 to JDN, from the Deutsche Forschungsgemeinschaft (DFG) as part of the priority program "New Frameworks of Rationality" (SPP 1516), and grant FA9550-14-1-0318 to MJ from the Air Force Office of Scientific Research. We thank Andreas

Sommer, Lara Bertram, and Gregor Caregnato for collecting the data, and Henry Brighton, Vincenzo Crupi, Laura Martignon, Jan Woike, and Charley Wu for helpful feedback on this research. We are grateful to the Museum für Naturkunde in Berlin for supporting our research and providing space for running Experiment 2.

References

- Baron, J. (1985). *Rationality and intelligence*. New York, NY: Cambridge University Press.
- Bellman, R. (1957). *Dynamic programming*. Princeton, NJ: Princeton University Press.
- Bonawitz, E., Denison, S., Griffiths, T. L., & Gopnik, A. (2014). Probabilistic models, learning algorithms, and response variability: sampling in cognitive development. *Trends in Cognitive Sciences*, 18, 497–500.
- Bramley, N. R., Lagnado, D. A., & Speekenbrink, M. (2015). Conservative forgetful scholars: How people learn causal structure through sequences of interventions. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 41, 708–731.
- Breiman, L., Friedman, J. H., Olshen, R. A., & Stone, C. J. (1984). *Classification and regression trees*. Monterey, CA: Wadsworth and Brooks.
- Bull, R., Espy, K. A., & Senn, T. E. (2004). A comparison of performance on the Towers of London and Hanoi in young children. *Journal of Child Psychology and Psychiatry*, 45, 743–754.
- Butko, N. J., & Movellan, J. R. (2010). Infomax control of eye movements. *IEEE Transactions on Autonomous Mental Development*, 2, 91–107.
- Coenen, A., Bramley, N., Ruggeri, A., & Gureckis, T. M. (2017). Beliefs about sparsity affect causal experimentation. In *Proceedings of the 39th Annual Meeting of the Cognitive Science Society* (pp. 1788–1793). Austin, TX: Cognitive Science Society.
- Coenen, A., Nelson, J. D., & Gureckis, T. M. (2018). Asking the right questions about the psychology of human inquiry: Nine open challenges. *Psychonomic Bulletin & Review*, . Advance online publication.

- Courage, M. L. (1989). Children's inquiry strategies in referential communication and in the game of twenty questions. *Child Development*, 60, 877–886.
- Crupi, V., Nelson, J. D., Meder, B., Cevolani, G., & Tentori, K. (2018). Generalized information theory meets human cognition: Introducing a unified framework to model uncertainty and information search. *Cognitive Science*, 42, 1410–1456.
- Daw, N. D., Niv, Y., & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience*, 8, 1704–1711.
- Denney, D. R. (1972). Modeling and eliciting effects upon conceptual strategies. *Child Development*, 43, 810–823.
- Denney, D. R., & Denney, N. W. (1973). The use of classification for problem solving: A comparison of middle and old age. *Developmental Psychology*, 9, 275–278.
- Denney, D. R., Denney, N. W., & Ziobrowski, M. J. (1973). Alterations in the information-processing strategies of young children following observation of adult models. *Developmental Psychology*, 8, 202–208.
- Denney, N. W., & Turner, M. C. (1979). Facilitating cognitive performance in children: A comparison of strategy modeling and strategy modeling with overt self-verbalization. *Journal of Experimental Child Psychology*, 28, 119–131.
- Eimas, P. D. (1970). Information processing in problem solving as a function of developmental level and stimulus saliency. *Developmental Psychology*, 2, 224–229.
- Gardner, W., & Rogoff, B. (1990). Children's deliberateness of planning according to task circumstances. *Developmental Psychology*, 26, 480–487.
- Geman, S., Bienenstock, E., & Doursat, R. (1992). Neural networks and the bias/variance dilemma. *Neural Computation*, 4, 1–58.
- Golovin, D., & Krause, A. (2011). Adaptive submodularity: Theory and applications in active learning and stochastic optimization. *Journal of Artificial Intelligence Research*, 42, 427–486.
- Gureckis, T. M., & Markant, D. B. (2012). Self-directed learning: A cognitive and computational perspective. *Perspectives on Psychological Science*, 7, 464–481.
- Hayes-Roth, B., & Hayes-Roth, P. F. (1979). A cognitive model of planning. *Cognitive Science*, 3, 275–310.

- Herwig, J. E. (1982). Effects of age, stimuli, and category recognition factors in children's inquiry behavior. *Journal of Experimental Child Psychology*, 33, 196–206.
- Huffman, D. A. (1952). A method for the construction of minimum-redundancy codes. *Proceedings of the IRE*, 40, 1098–1101.
- Hyafil, L., & Rivest, R. L. (1976). Constructing optimal binary decision trees is NP-complete. *Information Processing Letters*, 5, 15–17.
- Kachergis, G., Rhodes, M., & Gureckis, T. (2017). Desirable difficulties during the development of active inquiry skills. *Cognition*, 166, 407–417.
- Klahr, D., & Robinson, M. (1981). Formal assessment of problem-solving and planning processes in preschool children. *Cognitive Psychology*, 13, 113–148.
- Legare, C. H., Mills, C. M., Souza, A. L., Plummer, L. E., & Yasskin, R. (2013). The use of questions as problem-solving strategies during early childhood. *Journal of Experimental Child Psychology*, 114, 63–76.
- Lindley, D. V. (1956). On a measure of the information provided by an experiment. *The Annals of Mathematical Statistics*, 27, 986–1005.
- Markant, D. B., & Gureckis, T. M. (2014). Is it better to select or to receive? Learning via active and passive hypothesis testing. *Journal of Experimental Psychology: General*, 143, 94–122.
- Meder, B., & Nelson, J. D. (2012). Information search with situation-specific reward functions. *Judgment and Decision Making*, 7, 119–148.
- Meier, K. M., & Blair, M. R. (2013). Waiting and weighting: Information sampling is a balance between efficiency and error-reduction. *Cognition*, 126, 319–325.
- Morris, R., & Ward, G. (2004). *The cognitive psychology of planning*. New York, NJ: Psychology Press.
- Mosher, E., & Hornsby, J. R. (1966). On asking questions. In J. Bruner, R. Oliver, & P. Greenfield (Eds.), *Studies in cognitive growth* (pp. 86–102). New York, NY: Wiley.
- Najemnik, J., & Geisler, W. S. (2005). Optimal eye movement strategies in visual search. *Nature*, 434, 387–391.

- Navarro, D. J., & Perfors, A. F. (2011). Hypothesis generation, sparse categories, and the positive test strategy. *Psychological Review*, 118, 120–134.
- Nelson, J., Meder, B., , & Jones, M. (2018). Towards a theory of heuristic and optimal planning for sequential information search. Manuscript submitted for publication.
- Nelson, J. D. (2005). Finding useful questions: On Bayesian diagnosticity, probability, impact, and information gain. *Psychological Review*, 112, 979–999.
- Nelson, J. D., & Cottrell, G. W. (2007). A probabilistic model of eye movements in concept formation. *Neurocomputing*, 70, 2256–2272.
- Nelson, J. D., Divjak, B., Gudmundsdottir, G., Martignon, L. F., & Meder, B. (2014). Children’s sequential information search is sensitive to environmental probabilities. *Cognition*, 130, 74–80.
- Nelson, J. D., McKenzie, C. R. M., Cottrell, G. W., & J., S. T. (2010). Experience matters: Information acquisition optimizes probability gain. *Psychological Science*, 21, 960–969.
- Nemhauser, G. L., Wolsey, L. A., & Fisher, M. L. (1978). An analysis of approximations for maximizing submodular set functionsi. *Mathematical Programming*, 14, 265–294.
- Newell, A., & Simon, H. (1959). *The simulation of human thought (RAND Corporation paper P-1734)*. Santa Monica, US: RAND Corporation.
- Newell, A., & Simon, H. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.
- Oaksford, M., & Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review*, 101, 608–631.
- Oaksford, M., & Chater, N. (1996). Rational explanation of the selection task. *Psychological Review*, 103, 381–391.
- Quinlan, J. R. (1986). Induction of decision trees. *Machine Learning*, 1, 81–106.
- Rasmussen, C., & Williams, C. (2006). *Gaussian Processes for machine learning*. Cambridge, MA: MIT Press.
- Ruggeri, A., & Feufel, M. (2015). How basic-level objects facilitate question-asking in a categorization task. *Frontiers in Psychology*, 6, 918.

- Ruggeri, A., & Lombrozo, T. (2015). Children adapt their questions to achieve efficient search. *Cognition*, *143*, 203–216.
- Ruggeri, A., Lombrozo, T., Griffiths, T. L., & Xu, F. (2016). Sources of developmental change in the efficiency of information search. *Developmental Psychology*, *52*, 2159–2173.
- Ruggeri, A., Sim, Z. L., & Xu, F. (2017). “Why is Toma late to school again?” Preschoolers identify the most informative questions. *Developmental Psychology*, *53*, 1620–1632.
- Ruggeri, A., Walker, C., Lombrozo, T., & Gopnik, A. (2018). Effects of scaffolding on the type and informativeness of preschoolers’ questions. Unpublished manuscript.
- Schulz, E., & Gershman, S. J. (2019). The algorithmic architecture of exploration in the human brain. *Current Opinion in Neurobiology*, *55*, 7–14.
- Schulz, E., Wu, C. M., Ruggeri, A., & Meder, B. (2018). Searching for rewards like a child means less generalization and more directed exploration. Manuscript submitted for publication.
- Settles, B. (2010). *Active learning literature survey*. Technical Report University of Wisconsin–Madison.
- Shallice, T. (1982). Specific impairments of planning. *Philosophical Transactions of the Royal Society of London B*, *298*, 199–209.
- Shannon, C. E. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, *27*, 379–423.
- Siegler, R. S. (1977). The twenty questions game as a form of problem solving. *Child Development*, *48*, 395–403.
- Skov, R. B., & Sherman, S. J. (1986). Information-gathering processes: Diagnosticity, hypothesis-confirmatory strategies, and perceived hypothesis confirmation. *Journal of Experimental Social Psychology*, *22*, 93–121.
- Slowiaczek, L. M., Klayman, J., Sherman, S. J., & Skov, R. B. (1992). Information selection and use in hypothesis testing: What is a good question, and what is a good answer? *Memory & Cognition*, *20*, 392–405.

- Steyvers, M., Tenenbaum, J. B., Wagenmakers, E.-J., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, *27*, 453–489.
- Sutton, R. S. (1988). Learning to predict by the method of temporal differences. *Machine Learning*, *3*, 9–44.
- Sutton, R. S., & Barto, A. G. (1998). *Introduction to reinforcement learning*. Cambridge, MA: MIT Press.
- Thornton, S. (1982). Challenging early competence: A process oriented analysis of children’s classifying. *Cognitive Science*, *6*, 77–100.
- Trope, Y., & Bassok, M. (1982). Confirmatory and diagnosing strategies in social information gathering. *Journal of Personality and Social Psychology*, *43*, 22–34.
- Völter, C. J., & Call, J. (2014). Younger apes and human children plan their moves in a maze task. *Cognition*, *130*, 186–203.
- Watkins, C. J. C. H., & Dayan, P. (1992). Q-learning. *Machine Learning*, *8*, 279–292.
- Wells, G. L., & Lindsay, R. C. (1980). On estimating the diagnosticity of eyewitness nonidentifications. *Psychological Bulletin*, *88*, 776–784.
- Witten, I. H. (1977). An adaptive optimal controller for discrete-time Markov environments. *Information and Control*, *34*, 286–295.
- Wu, C. M., Meder, B., Filimon, F., & Nelson, J. D. (2017). Asking better questions: How presentation formats influence information search. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *43*, 1274–1297.
- Wu, C. M., Schulz, E., Speekenbrink, M., Nelson, J. D., & Meder, B. (2018). Generalization guides human exploration in vast decision spaces. *Nature Human Behaviour*, *2*, 915–924.
- Zajkowski, W. K., Kossut, M., & Wilson, R. C. (2017). A causal role for right frontopolar cortex in directed, but not random, exploration. *eLife*, *6*, e27430.

Appendix A. Task environment used in the experiments

Table A1 shows the canonical item–feature matrix for the task used in the present experiments. Questions Q1 through Q3 are constraint questions pertaining to features possessed by multiple items (mapped randomly onto physical features color, shape, and pattern for each subject), and Q4 through Q15 to features possessed by only a single item (e.g., crown, pipe, mustache etc., which are individuating visual features representing unique identifiers M1–M12; see Figure 1a).

Table A1. *Task environment. Rows M1 through M12 Denote the Individual Items (see Figure 1a) Represented by a Binary Feature Vector. Columns Q1 Through Q15 Represent the Possible Queries*

	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Q9	Q10	Q11	Q12	Q13	Q14	Q15
	Color	Shape	Pattern	Ears	Bow tie	Yellow hair	Blue horn	Mustache	Pipe	Red horns	Crown	Beard	Wings	Viking	Black hair
M1	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0
M2	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0
M3	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0
M4	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0
M5	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0
M6	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0
M7	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0
M8	0	1	1	0	0	0	0	0	0	0	1	0	0	0	0
M9	1	0	0	0	0	0	0	0	0	0	0	1	0	0	0
M10	1	1	0	0	0	0	0	0	0	0	0	0	1	0	0
M11	0	0	1	0	0	0	0	0	0	0	0	0	0	1	0
M12	1	1	0	0	0	0	0	0	0	0	0	0	0	0	1
Split	3:9	3:9	2:10	1:11	1:11	1:11	1:11	1:11	1:11	1:11	1:11	1:11	1:11	1:11	1:11
IG	0.811	0.811	0.650	0.414	0.414	0.414	0.414	0.414	0.414	0.414	0.414	0.414	0.414	0.414	0.414
PL	4.583	5.083	4.667	4.833	4.833	4.833	4.833	4.833	4.833	4.833	5.250	5.250	5.250	5.250	5.250
PLR	1.000	0.500	0.917	0.750	0.750	0.750	0.750	0.750	0.750	0.750	0.333	0.333	0.333	0.333	0.333

Note. IG = stepwise information gain in bits; PL = expected path length (expected number of questions when starting with this feature and searching globally optimally thereafter); PLR = expected path length reduction when starting with this feature and searching globally optimally thereafter.

Appendix B. Expected path length and path length reduction

In the 20-questions game, the goal is to identify an unknown target item with as few binary (yes–no) questions as possible. The globally optimal solution for this problem can be defined as the binary question tree with the shortest expected path length. In any complete question tree, each item corresponds to a leaf node at the end of the path of questions that uniquely identifies that item. The expected path length of the tree is the mean path length over all items (assuming a uniform prior for the target). Generally, the problem of identifying the optimal tree is intractable (Hyafil & Rivest, 1976) because the number of possible trees grows superexponentially in the number of queries.

For a moderate number of queries and items the globally optimal solution can be identified through dynamic programming (Nelson et al., 2018). Specifically, one identifies the optimal next question and expected remaining path length for every possible subset of remaining items. This can be done recursively, where once the answer has been found for all subsets of n or fewer items, it can then be solved for all subsets of $n + 1$, continuing until one reaches the full set of items. Code for implementing this algorithm, which requires only an item–feature matrix as input, is available at <https://osf.io/cq48j/>.

Here we describe how optimal question trees can be used to define the expected reduction in path length for any question on the first step of the search task. Let $\text{PPL}(Q)$ denote the expected posterior path length of a given query Q , that is, how many questions are remaining (on average) after asking that question first and deciding globally optimally thereafter:

$$\text{PPL}(Q) = \mathbb{E}[\text{remaining questions} | \text{start with } Q, \text{ globally optimal thereafter}]. \quad (\text{B.1})$$

The PPL can be calculated by considering the possible outcomes of the query (i.e., affirmative and negative answers), the probabilities of these outcomes, and the expected number of remaining questions in each case if deciding optimally thereafter. First, consider Q1. If $Q1 = 1$ (i.e., the question yields an affirmative answer, with $p = 3/12$) then three items remain, and if $Q1 = 0$ (a negative answer, with $p = 9/12$) then nine items remain. In the former case, only hypothesis-scanning questions are possible, and on average $1/3 * 1 + 2/3 * 2 = 5/3$ more questions are needed to identify the target item. If nine items remain, the globally optimal next query is Q2, which for this subset has an expected path length of 4.222. Thus, $\text{PPL}(Q1) = 3/12 * 5/3 + 9/12 * 4.222 = 3.583$.

Now consider Q2 as the first query. If $Q2 = 1$, three items remain (with $p = 3/12$), and if $Q2 = 0$, nine items remain (with $p = 9/12$). In the former case, on average

5/3 questions are needed to identify the target. If nine items remain, however, also only hypothesis-scanning questions remain, all of which have path length 4.889. Accordingly, $\text{PPL}(Q2) = 3/12 * 5/3 + 9/12 * 4.889 = 4.083$.

We can also define $\text{PL}(Q)$ as the expected total path length of a given query Q :

$$\text{PL}(Q) = \mathbb{E}[\text{total questions} | \text{start with } Q, \text{ globally optimal thereafter}]. \quad (\text{B.2})$$

By definition, $\text{PL}(Q) = \text{PPL}(Q) + 1$. For instance, $\text{PL}(Q1) = 3.583 + 1 = 4.583$, meaning that on average 4.583 total questions are required to identify the target item with certainty, if starting with Q1 and deciding globally optimally thereafter (i.e., on each subsequent step choosing the query with the minimum PL based on the remaining items).

The expected path length *reduction* (PLR) of a given query Q can now be defined based on its expected posterior path length, $\text{PPL}(Q)$, relative to the expected path length of the globally optimal query:

$$\text{PLR}(Q) = \min_{Q' \in \mathbb{Q}} \text{PL}(Q') - \text{PPL}(Q) \quad (\text{B.3})$$

where $\mathbb{Q}$ denotes the set of all available queries. For instance, in the environment considered here, the most efficient first query is Q1, such that $\min_Q \text{PL}(Q) = \text{PL}(Q1) = 4.583$ (Table 1). Accordingly, $\text{PLR}(Q1) = \text{PL}(Q1) - \text{PPL}(Q1) = 4.583 - 3.583 = 1.0$, whereas $\text{PLR}(Q2) = \text{PL}(Q1) - \text{PPL}(Q2) = 4.583 - 4.083 = 0.5$.

Appendix C. Simulations for assessing performance of stepwise-optimal information gain and the split-half heuristic

To assess the performance of stepwise-optimal information gain and equivalent strategies like the split-half heuristic, we generated a large number of environments (each defined by an item–feature matrix; cf. Table A1). Nelson et al. (2018) report a more extensive variety of simulations, on which the present simulations are based. In addition to the constraint questions, each item had a unique feature, such that the total number of available questions in an environment is the number of items plus the number of constraint questions (e.g., if there are 15 items and 3 constraint questions, the total number of questions is 18). For each generated environment, we tracked whether the question at Step 1 with the highest stepwise information gain also had the shortest expected path length (or, equivalently, highest expected reduction in path length). If there was more than one maximum information gain query, one of them was picked at random. We also tracked how much loss was incurred by the query chosen by the stepwise-optimal strategy, where loss was defined as the difference in the expected number of questions relative to the globally optimal query,

$$\text{PL}(\text{globally optimal query}) - \text{PL}(\text{max IG query}) \quad (\text{C.1})$$

where PL denotes expected path length and IG denotes stepwise-optimal information gain. We also computed the relative increase of queries, defined as

$$1 - \frac{\text{PL}(\text{globally optimal query})}{\text{PL}(\text{max IG query})} \quad (\text{C.2})$$

For the simulations, we varied three factors: the number of items (12, 15, 18), the number of constraint questions (2, 3, 5, 7, 9), and the mean feature density (splithaleness) of the features. The density of each constraint feature was sampled from a $\text{Beta}(\alpha, \beta)$ distribution with an expectation corresponding to the desired mean feature density, where α and β denote the shape parameters of the distribution. The five beta distributions were $\text{Beta}(1,9)$, $\text{Beta}(2,8)$, $\text{Beta}(3,7)$, $\text{Beta}(4,6)$, and $\text{Beta}(5,5)$, which have means of 0.1, 0.2, 0.3, 0.4, and 0.5, respectively.

Sampling was done independently for each constraint feature. The sampled density was used to generate a binary vector of length n (where n = number of items) using Bernoulli sampling. Thus the actual density of a feature differed from the generating density (the value sampled from the beta distribution) according to binomial sampling error. If the sampled vector corresponded to an uninformative feature (i.e., a feature possessed by all or no item), we resampled (keeping the generating density) until an informative feature was created. For each of the $3 \times 5 \times 5 = 75$ parameter

combinations, we generated 1,000 environments (binary item $\times$ feature matrices) subject to the applicable constraints.

Appendix D. Script used in Experiment 1

Script used in Experiment 1. This script, as well as the script for Experiment 2, was translated from German. The instructions to the experimenter running the session is given in italicized text.

Hello [*name of participant*]! My name is [*name of experimenter*]. We are going to play a game together! In this game I have a very special machine [*point to machine*], see? This machine is super cool: when it's turned on, it plays music! I'll show you how it works.

The machine turns on when you put two identical monsters inside. Here are two monsters, which are completely identical, see? [*Show two additional identical red monsters.*] First, I put one of the monsters in the machine and nothing happens. But if I add the second, identical monster... the machine turns on! Look! Isn't that funny? So, here's a bunch of other monsters: look! [*Show both decks of cards*]. These two decks contain the same monsters. I'll show you: First, I take all monsters of the first group.

[*Take one deck and shuffle it three times. From the shuffled deck of cards, place monsters individually, step by step, in a 3×4 grid. Cards should be placed in front of the child or at least on the child's side of the table; to the right of the cards there should be enough space to later on in the game place the rejected cards. Start with the card on top, which is placed in the top left corner. The subsequent card goes to the right of it; fill the rows in this manner and start a new row when a row is full.*]

Now look at these monsters here [*show second deck of cards*]; these are exactly the same. [*Place monsters from second deck individually, step by step in a second 3×4 grid, so their order matches the one of the first grid. Afterward collect the cards and shuffle them three times.*] Now I randomly pick one of these monsters [*Pick one monster blindly*], look at it, and put it into the machine. [*Put monster into machine.*] This means that one of these monsters here [*point to cards*] looks exactly the same as the monster that I just put into the machine [*point to machine*]. The goal of this game is to find this identical monster, to turn on the machine!

To find the right monster, you can ask me questions. To these questions, I can only reply with "yes" or "no." And you also get ten 50-cent coins. [*Hand over 10 coins and place in front of the child, to the right of the cards.*] In this game, you have to pay one 50-cent coin for every question that you ask me. At the end of the game you can keep all the coins that you have left, when you have found the monster to turn on the machine. Do you like that? This means you want to ask as few questions as possible, to take home as much money as possible.

The questions that you can ask me are on these cards. [*Take deck of shuffled question cards and put them down one by one from left to right.*] Let's have a look at the questions that you can ask me.

[*randomized order of question cards (Q cards): cards are shuffled three times. Point to first/second/third/fourth/last Q card*] By using this question you can find out whether the monster [*is blue/green*] [*is round/square*] [*has dots/no dots*] [*is one particular monster*]. [*Every question has two versions; which version is used is chosen randomly. One of the Q cards is not available in the second round, as long as the first question is not for a particular monster.*]

COLOR: By using this [*next/last*] question, you can ask about the color of the monster. The monsters are [*upper color*] or [*lower color*]. The question therefore is whether the monster is [*green/blue, state upper color first*]. Lets order the monsters according to their color! [*Start with upper color on Q card*]: Lets put all [*green/blue*] monsters here [*point to right side of table*] and all [*blue/green*] monsters here [*point to left side of the table*]. Great! Now lets count together how many green/blue monsters there are: one, two, ... [*n green/blue*] monsters! [*Start counting: if child starts to count as well let them keep on counting independently*]. And now, let's count how many [*blue/green*] monsters we have: one, two, ... [*n blue/green*] monsters! [*Start counting: if child starts to count as well, let them keep on counting independently.*] So there are [*n green/blue*] monsters and [*n blue/green*] monsters! [*Shuffle monsters and place again randomly in grid; the ordering can be different from the previous ordering.*]

SHAPE: By using this [*next/last*] question, you can ask about the shape of the monster. The monsters are [*upper shape*] or [*lower shape*]. The question therefore is whether the monster is round/square [*state upper shape first*]. Let's order the monsters according to their shape! [*Start with upper shape on Q-card*] Let's put all [*round/square*] monsters here [*point to right side of the table*] and all [*square/round*] monsters here [*point to left side of the table*]. Great! Now let's count together how many [*round/square*] monsters there are: one, two, ... [*n round/square*] monsters! [*Start counting: if child starts to count as well let them keep on counting independently.*] And now, how many [*n square/round*] monsters we have: one, two, ... [*n square/round*] monsters! [*Start counting: if child starts to count as well let them keep on counting independently.*] So there are [*n round/square*] monsters and [*n round/square*] monsters! [*Shuffle monsters and place again randomly in grid; the ordering can be different from the previous ordering.*]

PATTERN: By using this [*next/last*] question, you can ask about the pattern of the monster. The monsters are [*upper pattern*] or [*lower pattern*]. The question therefore is whether the monster is [*dotted/not dotted*] [*state upper first*]. Let's order

the monsters according to their pattern! [*Start with upper pattern on Q-card*] Let's put all [*dotted/not dotted*] monsters here [*point to right side of table*] and all [*not dotted/dotted*] monsters here [*point to left side of the table*]. Great! Now let's count together how many [*dotted/not dotted*] monsters there are: one, two, ... [*n dotted/not dotted*] monsters! [*Start counting: if child starts to count as well, let them keep on counting independently.*] And now, how many [*not dotted/dotted*] monsters there are: one, two,... *n non dotted/dotted* monsters! [*Start counting, if child starts to count as well let them keep on counting independently.*] So there are [*n non dotted/dotted*] monsters and [*n non dotted/dotted*] monsters! [*Shuffle monsters and place again randomly in grid; the ordering can be different from the previous ordering.*]

SINGLE MONSTER: By using this [*next/last*] question, you can ask about a particular monster! When you ask this question, you can choose one monster and ask whether this monster can turn on the machine. Which of the monsters could you ask for, just as an example? [*Let child generate a question for a single monster.*] Great! Now let's take this monster out. [*Take monster out and put it to the right of the grid.*] We therefore have one monster here, and one, two, ... [*n*] monsters here! [*Start counting: if child starts to count as well, let them keep on counting independently.*] [*Shuffle monsters and place again randomly in grid; the order can diverge from previous order.*]

Remember, you want to find out which monster can turn on the machine by asking as few yes/no questions as possible! For every question that you ask, you have to give me one 50-cent coin. Do you understand the rules? [*Wait for confirmation. If positive: OK, if everything is clear, let's get started! If negative: Out of these (point to Q cards) you can choose one question you would like to ask. Try to use as few questions as possible to find the monster that looks like the one in the machine. Do you have any specific question regarding the game? If no and game understood: OK, if everything is clear, let's get started!*] [*Start playing the game.*]

So, which question would you like to ask me first? [*Point to Q cards. Let child choose a Q card and hand back one 50-cent coin. Hold hand open, so child can place the coin in it.*] [*If necessary: For every question, you have to give me one 50-cent coin. Can you please give me one?*] Take chosen card from the table: We can put this card away now. [*Except if it's the question for a particular monster, this question will be left in every round.*] Give feedback according to the randomly chosen "special" monster:

Well done! Well, [*yes/no it is/isn't feature*]. Alright, so we know now that it's not one of these monsters [*Remove monster(s) excluded by the question from the table.*]

So we have these monsters left. [*Point to monsters.*] [*After first round:*] Let's have a look at the questions that you could ask now. [*Point to Q cards.*]

[*After second question:*] Alright, we can put these questions away now. [*Put away remaining Q cards.*]

****These monsters are left now. From now on, you can ask every yes/no question that you want, even those that weren't on the cards before! For every question you still have to trade one 50-cent coin. What would you like to ask me? [*Give child time to generate question. Give feedback: A very good question!*] Well, [*yes/no, the monster is/isn't feature*]. Alright, so we know now that it's not one of these monsters. [*Remove monsters excluded by the question from the table.*]

[*Repeat from **** until only one monster is left.*]

[*If necessary, encourage: Are you having trouble thinking of a question? Come on, let's have a closer look at the monsters. What could you possibly ask? If still no question generation: Is there any specific feature about the monsters you could ask for? Or do you want to ask for a particular monster?*]

Yay! There's only one monster left! Shall we have a look and see whether it can turn on the machine? You can put it inside and try! [*Machine turns on.*] Wow, isn't that fun! And you even have [*n*] 50-cent coins left!

Thank you for playing this game with me! That was fun, wasn't it?

Appendix E. Script used in Experiment 2

The procedure in Experiment 2 was identical to that of Experiment 1, with the additional inclusion of the binary forced-choice paired-comparison task in which participants chose between two queries. As in Experiment 1, participants were first familiarized with the game and the materials (Appendix D). Then, before playing the game, the paired-comparison task was implemented as follows:

Before you can start selecting the questions you want to ask, I have some questions for you, OK?

[Order of paired-comparison tasks follows a predetermined random order. In each comparison, questions are presented at the same time, before the explanation is given.]

Comparison Q1 vs. Q2: Imagine that another child plays this game and selects this question first *[point to Q1 of the corresponding set]*. Another child also plays this game and selects this question first *[point to Q2 of the corresponding set]*. What do you think: Which of the two will find the monster that turns on the machine faster? *[Let the child select a query.]* Ah, that's an interesting idea!

Comparison Q1 vs. Q3: Imagine that another child plays this game and selects this question first *[point to Q1 of the corresponding set]*. Another child also plays this game and selects this question first *[point to Q3 of the corresponding set]*. What do you think: Which of the two will find the monster that turns on the machine faster? *[Let the child select a query.]* Ah, that's an interesting idea!

Comparison Q2 vs. Q3: Imagine that another child plays this game and selects this question first *[point to Q2 of the corresponding set]*. Another child also plays this game and selects this question first *[point to Q3 of the corresponding set]*. What do you think: Which of the two will find the monster that turns on the machine faster? *[Let the child select a query.]* Ah, that's an interesting idea!

Comparison Q1 vs. random monster M1–M7: Imagine that another child plays this game and selects this question first *[point to Q1 of the corresponding set]*. Another child also plays this game and selects this question first *[point to a monster card M1–M7 of the corresponding set, following randomization]*. What do you think: Which of the two will find the monster that turns on the machine faster? *[Let the child select a query.]* Ah, that's an interesting idea!

Comparison Q1 vs. random monster M9–M12: Imagine that another child plays this game and selects this question first *[point to Q1 of the corresponding set]*. Another child also plays this game and selects this question first *[point to a monster card M8–M12 of the corresponding set, following randomization]*. What do you think: Which of the two will find the monster that turns on the machine faster? *[Let the child select a query.]* Ah, that's an interesting idea!

Great, then its your turn now!

[*Playing the game follows the same procedure as in Experiment 1: see Appendix D.*]