

Brief article

Children's sequential information search is sensitive to environmental probabilities



Jonathan D. Nelson^{a,*}, Bojana Divjak^b, Gudny Gudmundsdottir^a, Laura F. Martignon^b, Björn Meder^{a,*}

^a Center for Adaptive Behavior and Cognition (ABC), Max Planck Institute for Human Development, Lentzeallee 94, 14195 Berlin, Germany

^b Ludwigsburg University of Education, Institute of Mathematics and Computing, Reuteallee 46, 71634 Ludwigsburg, Germany

ARTICLE INFO

Article history:

Received 27 April 2013

Revised 24 September 2013

Accepted 25 September 2013

Available online 1 November 2013

Keywords:

Information search

Optimal experimental design principles

Twenty-questions game

Optimality

Heuristics

Information gain

ABSTRACT

We investigated 4th-grade children's search strategies on sequential search tasks in which the goal is to identify an unknown target object by asking yes–no questions about its features. We used exhaustive search to identify the most efficient question strategies and evaluated the usefulness of children's questions accordingly. Results show that children have good intuitions regarding questions' usefulness and search adaptively, relative to the statistical structure of the task environment. Search was especially efficient in a task environment that was representative of real-world experiences. This suggests that children may use their knowledge of real-world environmental statistics to guide their search behavior. We also compared different related search tasks. We found positive transfer effects from first doing a number search task on a later person search task.

© 2013 The Authors. Published by Elsevier B.V. Open access under [CC BY license](#).

1. Introduction

Often inferences and decisions must be made before all relevant information can be obtained. In these situations, careful selection of questions to ask (or queries to make or experiments to conduct) is very important. Examples include a child asking a question to learn the meaning of a novel word, a scientist choosing an experiment to differentiate between competing hypotheses, or a person's visual system directing the eyes' gaze to informative parts of a visual scene.

How do children and adults search for information? Many studies investigating information search have used

variants of the “20-questions” game. In this game, the task is to identify an unknown target item by asking as few yes–no (binary) questions as possible. Much research has focused on the frequency of different kinds of questions in different age groups (Denney & Denney, 1973; Eimas, 1970; Mosher & Hornsby, 1966; Ruggeri & Feufel, submitted for publication; Thornton, 1982). Younger children tend to ask about specific objects (*hypothesis-testing* questions, e.g., “Is it Paul?”), or questions that, while phrased in terms of constraints, in fact pertain to individual objects (*pseudoconstraint* questions). An example would be asking “Does the person have a beard?” when there is only one person with a beard in the set. Older children tend to ask about properties that differentiate between subsets of multiple objects (*constraint* questions, e.g., “Is the person wearing a hat?”). Interestingly, one study found that elderly adults ($M_{\text{age}} = 83$) required 32 questions, whereas younger adults ($M_{\text{age}} = 38$) required only 18 questions on a related task (Denney & Denney, 1973).

We study information search in fourth-grade (8–10 year old) children, an age in which they begin to

* Corresponding authors. Tel.: +49 030 82406 658 (J.D. Nelson), +49 030 82406 239 (B. Meder).

E-mail addresses: nelson@mpib-berlin.mpg.de, jonathan.d.nelson@gmail.com (J.D. Nelson), meder@mpib-berlin.mpg.de (B. Meder).

imagine two or three steps ahead in problem solving (Sieglar & Stern, 1998) and playing games (Amit & Jan, 2006). They also develop skill at comparing simple proportions (Fischbein, Pampu, & Minzat, 1970; Martignon & Krauss, 2009; Tourniaire & Pulos, 1985). We investigate children's sensitivity to the varying usefulness of constraint questions in different environments.

1.1. Theoretical background and the Person Game

The goal in the *Person Game*, which we analyze mathematically and use in our experiment, is to identify an unknown target person by asking as few yes–no questions about the person's features as possible. This equates to finding the *question tree* (binary decision tree) that has the smallest expected total number of questions. A question tree specifies which question is asked first, and depending on the answer to that question, what question is asked next, and so on (Fig. 1). In the Person Game, the available questions correspond to 20 physical features of the cartoon faces. The possible people are equally probable a priori. Suppose that the question is whether the (unknown) person is wearing a hat. If the answer is “no”, all persons with hats can be eliminated; if the answer is “yes”, all persons without hats can be eliminated. For large

problems, such as person games with large numbers of people, it can be infeasible to use exhaustive search (which is NP-complete; Hyafil & Rivest, 1976) to identify the optimal question tree. We therefore also discuss stepwise *information gain* (Cover & Thomas, 1991; Lindley, 1956; Oaksford & Chater, 1996, 2003), a statistical model that is computationally simpler to implement. The highest-information-gain question is the question that, in the expectation after the question's answer is known, will lead to lowest expected posterior uncertainty (Shannon, 1948, entropy). Reduction in uncertainty is considered information about the true category. In sequential search tasks, stepwise (greedy) procedures are not in general optimal.

As descriptive models, information gain and related optimal experimental design ideas (like probability gain and impact) have been used to predict questions on a variety of tasks (Austerweil & Griffiths, 2011; Eimas, 1970; Markant & Gureckis, 2012; Meder & Nelson, 2012; Nelson, 2005, 2008; Nelson, McKenzie, Cottrell, & Sejnowski, 2010; Nelson, Tenenbaum, & Movellan, 2001; Oaksford & Chater, 1994; Oaksford & Chater, 2003), to predict human eye movements (Bicknell, 2011; Legge, Klitz, & Tjan, 1997; Meier & Blair, 2013; Najemnik & Geisler, 2005; Nelson & Cottrell, 2007; Walker Renninger, Coughlan, Verghese, & Malik, 2005), and to predict firing of individual neurons (Nakamura, 2006).

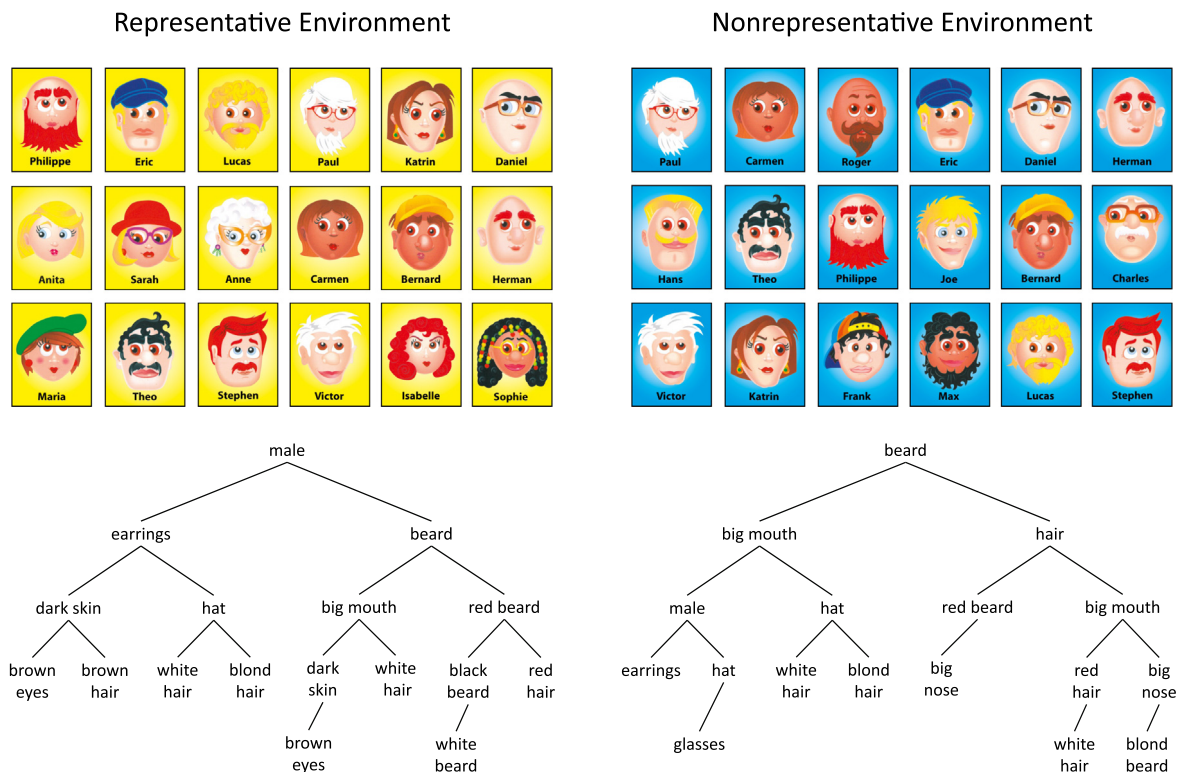


Fig. 1. Task environments and optimal decision trees in the Person Game. In both environments, stepwise information gain and the split-half heuristic identify the optimal question tree. In the Representative Environment (top left), Gender is the most informative first question; in the Nonrepresentative Environment (top right), Beard is the most informative first question. (The German word for beard, *Bart*, refers to various kinds of facial hair including a full beard, mustache, or chin-only beard.) The trees below the stimuli show the optimal (shortest expected path length) search trees for the Representative Environment (bottom left) and for the Nonrepresentative Environment (bottom right), as identified through exhaustive search. In the question trees, if the answer is “no” one takes the left branch; if the answer is “yes” one takes the right branch.

In the Person Game each of the n persons is equally probable in the beginning. Let n_{yes} denote the number of faces which have a particular feature, and n_{no} denote the number of faces that lack that feature. The information gain (IG) of a question Q about that feature is:

$$IG(Q) = \log_2 n - \left[\frac{n_{no}}{n} \log_2 n_{no} + \frac{n_{yes}}{n} \log_2 n_{yes} \right] \quad (1)$$

Information gain is defined in terms of a weighted average of logarithms. Are there simple strategies that could identify the highest-information-gain question? Consider the *split-half heuristic*. It finds a feature that comes closest to being possessed by half of the remaining individuals, and asks about that feature. Importantly, it can be proven that in the Person Game the split-half heuristic always chooses the highest-information-gain question (Navarro & Perfors, 2011). This finding contributes to a body of research showing that heuristic information-acquisition strategies can approximate (Gigerenzer & Gaismaier, 2011; Klayman & Ha, 1987; Markant & Gureckis, 2012; Slowiaczek, Klayman, Sherman, & Skov, 1992) or even exactly implement (Nelson, 2005, 2009) particular statistical models. Previous studies have found varying rates of use of the split-half strategy. Eimas (1970) found use of the split-half strategy varied widely depending on the number of target items, number of available constraint questions, and saliency of stimuli. Among 2nd graders, the proportion ranged from 0% to 19%; among college students, from 13% to 75%.

1.2. Transfer effects and generalizable insight

Another important issue is whether intrinsically motivating games can instill generalizable intuitions about information-search strategies (“learning by playing”; Hirsh-Pasek & Golinkoff, 2003). Siegler (1977) began to explore this, by randomizing the order of structurally homologous letter and number guessing games, with 13–14-year-old children, in an experiment in which the use of informative question strategies was specifically encouraged. Siegler found that playing a number game beforehand led to improved performance on a letter game. He hypothesized that ordinal relationships among the numbers are more apparent than ordinal relationships among letters. We address whether positive transfer effects can occur between non-structurally-homologous games, in 8–10-year-old children, when instructions do not specifically encourage use of informative strategies.

2. Experiment

Theoretically speaking, the immediate statistics of the set of cards available should determine the questions that are asked. However, it may not be easy to immediately internalize the full joint distribution of persons and features. This suggests that it would make sense for people’s ideas of questions’ relative usefulness to be influenced in part by their own prior experience with the statistics of faces in the world. To address this, we examined search behavior while manipulating the statistical structure of the faces in the Person Game (the environment), and therefore the structure of the optimal question trees. We used

two statistical environments, a *Representative Environment* (Fig. 1, top left), with the gender distribution approximately equal (10 men, 8 women) and a *Nonrepresentative Environment* (Fig. 1, top right), in which the gender distribution was highly skewed (16 men, 2 women).

We derived the globally-optimal question trees for each environment through exhaustive search (Fig. 1, bottom). In the Representative Environment, Gender is the most informative first question. In the Nonrepresentative Environment, Beard (facial hair),¹ which is not a very useful question in the Representative Environment, is the best first question. The Nonrepresentative Environment is nonrepresentative in the sense that both the Beard and Gender feature proportions greatly differ from the real-world experiences of the children in our experiment, who have experienced roughly equal proportions of men and women, and only a minority of men with beards (cf. Nelson, 2005, Table 13). In both environments stepwise information gain and the split-half heuristic identified the optimal question tree.

2.1. Method

2.1.1. Participants and design

Participants were 60 fourth-grade children between age 8 and 10 (67% girls) from Ludwigsburg, Germany, who were not familiar with the Guess Who (“Wer ist es?”, by Hasbro) game from which the stimuli were taken. Factors “Person Game Environment” (Representative vs. Nonrepresentative) and “Order of Games” (Person Game First vs. Number Game First) were manipulated between subjects.

2.1.2. Materials and procedure

Stimuli were printed on cards. For the Person Game the cards showed 18 cartoon-like faces, placed in random arrangement in front of the child (Fig. 1). The experimenter explained that she would draw a random person from an identical set of people (face cards) and that the child’s task was to identify this target by asking as few yes–no questions as possible about the person’s features. To make clear that each card was equiprobable, in each round of the game the cards were shuffled face down and a random target face card was chosen. The experimenter also explained that if the child needed help identifying a question, they could refer to twenty available questions (physical features), which were printed on a different set of cards and placed near to the child. The face cards eliminated through a question were turned over by the child, if needed with help from the experimenter. Questioning continued until the target was identified. Each child played five rounds of the Person Game.

To explore the feasibility of using games to instill generalizable insight, we also included a non-structurally-homologous *Number Game*. The Number Game task was similar: to identify a randomly selected integer between 1 and 18, by asking yes–no questions. 18 number cards were ordered in front of the child. However, in the Number Game, arbitrary questions (pertaining to any subset of the numbers, e.g., “Is the number 7, 8, or 14?”) were allowed,² and no cards with possible

¹ The German word Bart was used in the experiment; it refers to various kinds of facial hair including a full beard (Vollbart), mustache (Schnurrbart), chin-only beard (Kinnbart), or goatee (Ziegenbart).

² In this case, the Huffman (1952) code identifies the optimal tree, and exhaustive search is not required.

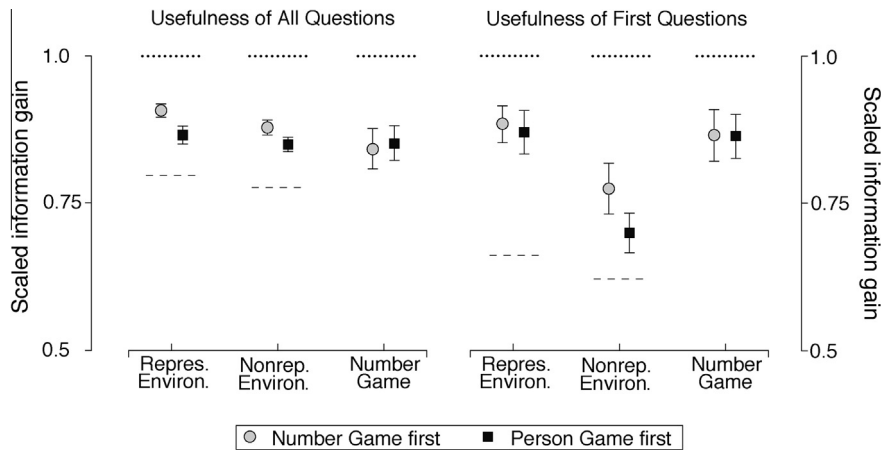


Fig. 2. Average informational value (\pm SEM) of questions asked, on the Person Game and Number Game. At left, usefulness of all questions. At right, usefulness of the first questions. Children who had first completed the Number Game asked higher-information-value questions on the Person Game (left side). On both the left and right halves, the first two sets of data points (Repres. and Nonrep. Environ.) are from the Person Game; the third set of data points is from the Number Game. In the Person Game, the first questions were more informative in the Representative Environment than in the Nonrepresentative Environment. Optimal performance, denoted with the dotted lines, would correspond to scaled information gain of 1. Random performance is denoted with dashed lines. Data were obtained by first averaging each game for each child, then averaging all games for each child, and finally by averaging across children. The random strategy information value was obtained through simulation of a strategy that picks at random from the list of 20 available constraint questions. Including the hypothesis-testing (“name”) questions would reduce the information value of the random strategy. In the Number Game, arbitrary questions were allowed. In general, where there are n items and all questions are allowed, there are $(2n - 2)/2$ informative and non-redundant possible questions. With 18 numbers this entails 131,071 potentially informative and nonredundant questions. It is not clear what random strategy would be equivalent to the Person Game random strategy, so no random performance is calculated for the Number Game.

questions were provided. The Number Game was played several times with random target numbers, for about 20 min.

3. Results

Because exhaustive search showed that stepwise information gain identifies the optimal question strategies in these environments, we use information gain to quantify questions’ usefulness. So that the best-available question always has a value of one, we report the *scaled* expected information gain (Hattori, 2002; Oaksford & Chater, 2003), which is obtained by dividing each available question’s information gain by the information gain of the most informative available question. Perfect use of the split-half heuristic leads to scaled information gain of 1 on every question.

Children asked questions that were more useful than a chance strategy, but less useful than the optimal strategy, in both the Person Game and in the Number Game (Fig. 2). Aggregate performance in the Person Game ($M = .87$, $Md = .88$) and Number Game ($M = .85$, $Md = .92$) was similar. However, the Number Game performance spanned a much wider range ($SD = .17$, range .38 to .996) than the Person Game performance ($SD = .05$, range .74 to .98). An F test revealed that the difference in variance is statistically reliable ($F(59, 59) = 11.05$, $p < .0001$); bootstrap sampling (which is robust to nonnormality) corroborated this result.

On the Person Game, children who had first played the Number Game asked higher-usefulness questions than children who played the Person Game first ($t(58) = 2.67$, two-tailed $p = .01$, Cohen’s $d = 0.69$). There was no transfer from the Person Game to the Number Game ($p = .83$).

Did children’s search behavior adapt to the statistical structure of each Person Game environment? While performance was high in both environments (Fig. 2), there

was a trend to ask higher-usefulness questions in the Representative Environment ($t(58) = 1.66$, two-tailed $p = .1$, $d = 0.43$). To explore this trend, we analyzed children’s search separately with respect to the first question asked, and for the other questions. From a normative perspective, the first question is the most important.³

When the first questions were excluded, questions’ mean scaled information gain did not differ between environments ($M_{\text{Representative}} = .89$ vs. $M_{\text{Nonrepresentative}} = .90$; $p = .28$). It thus appears that aggregate differences between environments were driven by the first question. Children asked higher-usefulness first questions in the Representative Environment than in the Nonrepresentative Environment ($t(58) = 3.82$, $p = .0003$, $d = 0.99$; Fig. 2, right). The Spearman rank correlation between first question frequency and objective usefulness was .75 in the Representative Environment and .53 in the Nonrepresentative Environment. In each task environment the most informative first question was the most frequent first question (Fig. 3). Adaptation to the statistical structure of the task environment was seen from the first round of the game (Fig. 4). Learning from experience over repeated games was not required for that adaptation, although there may be a learning trend across the five rounds of the game.

Fig. 3 shows the distribution of first questions for each environment, relative to the questions’ objective usefulness values. Gender was strongly preferred in the Representative Environment (55% of first questions), in which it is objectively most useful, but was also popular in the

³ When there are just two or three cards remaining, all informative questions have the same usefulness. In both environments, simulations show that the raw and scaled information gain of the random strategy increases gradually as the number of remaining face cards decreases.

Nonrepresentative Environment (24% of first questions), where it has low information value. The Beard question was seldom asked in the Representative Environment (4% of first questions), even though it tied for second-most-useful, but was the most frequent first question in the Nonrepresentative Environment (25% of first questions), where it was the most useful question.

4. Discussion

We observed a positive transfer effect from the Number Game to the Person Game. This shows that the games do not have to be structurally homologous for a facilitative transfer effect to occur, even among 4th grade children. Future research should explore a broad set of interventions

Representative Environment

Question	Split	IG	Scaled IG	Total	Percentage
Gender (female/male)	8:10	0.99	1.00	83	55.3%
Beard (present/absent)	5:13	0.85	0.86	6	4.0%
Nose (large/small)	5:13	0.85	0.86	2	1.3%
Glasses (present/absent)	5:13	0.85	0.86	6	4.0%
Mouth (large/small)	4:14	0.76	0.77	4	2.7%
Eyes (blue/brown)	4:14	0.76	0.77	2	1.3%
Earrings (present/absent)	4:14	0.76	0.77	7	4.7%
Hat (present/absent)	4:14	0.76	0.77	13	8.7%
Blond hair (present/absent)	4:14	0.76	0.77	3	2.0%
Brown hair (present/absent)	4:14	0.76	0.77	2	1.3%
Skin (dark/light)	3:15	0.65	0.66	1	0.7%
White hair (present/absent)	3:15	0.65	0.66	3	2.0%
Hair (present/absent)	3:15	0.65	0.66	9	6.0%
Red beard (present/absent)	2:16	0.50	0.51	1	0.7%
Black hair (present/absent)	2:16	0.50	0.51	1	0.7%
Red hair (present/absent)	2:16	0.50	0.51	4	2.7%
Blonde beard (present/absent)	1:17	0.31	0.31	0	0.0%
Black beard (present/absent)	1:17	0.31	0.31	0	0.0%
White beard (present/absent)	1:17	0.31	0.31	0	0.0%
Brown beard (present/absent)	0:18	0.00	0.00	0	0.0%
Theo	1:17	0.31	0.31	1	0.7%
Herman	1:17	0.31	0.31	2	1.3%

Nonrepresentative Environment

Question	Split	IG	Scaled IG	Total	Percentage
Beard (present/absent)	9:9	1.00	1.00	38	25.3%
Mouth (large/small)	7:11	0.96	0.96	4	2.7%
Nose (large/small)	6:12	0.92	0.92	5	3.3%
Hair (absent/present)	5:13	0.85	0.85	18	12.0%
Blond hair (present/absent)	4:14	0.76	0.76	3	2.0%
Skin (dark/light)	4:14	0.76	0.76	3	2.0%
Brown hair (present/absent)	3:15	0.65	0.65	3	2.0%
Black hair (present/absent)	3:15	0.65	0.65	7	4.7%
Hat (present/absent)	3:15	0.65	0.65	13	8.7%
Glasses (present/absent)	3:15	0.65	0.65	3	2.0%
Eyes (blue/brown)	3:15	0.65	0.65	10	6.7%
Gender (female/male)	2:16	0.50	0.50	36	24.0%
White hair (present/absent)	2:16	0.50	0.50	1	0.7%
Black beard (present/absent)	2:16	0.50	0.50	1	0.7%
Blonde beard (present/absent)	1:17	0.31	0.31	0	0.0%
Red hair (present/absent)	1:17	0.31	0.31	1	0.7%
Earrings (present/absent)	1:17	0.31	0.31	4	2.7%
Red beard (present/absent)	1:17	0.31	0.31	0	0.0%
White beard (present/absent)	1:17	0.31	0.31	0	0.0%
Brown beard (present/absent)	1:17	0.31	0.31	0	0.0%

Fig. 3. First questions asked in the Representative Environment (top) and the Nonrepresentative Environment (bottom). Split = initial feature distribution; IG = information gain. In each condition children received the same list of 20 constraint questions. In the Representative Environment, hypothesis-testing name questions (Theo, Herman) were asked a total of three times, although the name questions were not included in the set of suggested questions.




Fig. 4. Proportion of Gender and Beard questions as first question in the two environments of the Person Game, in each round of the game. The data indicate strong effects of adaptation to the environment; these effects are apparent from the very first round of the game. The Gender question was much more frequent in the Representative Environment than in the Nonrepresentative Environment, and the Beard question was much more frequent in the Nonrepresentative Environment than in the Representative Environment.

(perhaps small-group discussion, experimenting with different physical arrangements of cards, etc.) to enhance attainment of generalizable insights.

We found that information search was more efficient in a task environment that was representative of children's real-world experiences. Our results suggest key issues for further theory and model development.

Questions of varying information gain were asked. Could explore-exploit strategies, such as epsilon greedy or softmaxing (Sutton & Barto, 1998), together with the information value of the questions, explain this? These strategies entail occasional or proportional selection of low-information-value questions, and may be an important component of a full theory. However, they cannot explain why certain questions (e.g., Hat in both environments, Gender in the Nonrepresentative Environment) were much more prominent than other similarly low-information-value questions.

Does salience explain the results? Unfortunately, salience is an umbrella idea that encompasses many findings. In eye movement experiments, features may become salient because of abstract physical properties (Itti & Baldi, 2006), or because they have been useful previously (Nelson & Chenkov, 2010). In our experiment, questions could also be popular because of additional goals—such as differentiating between male and female, which correspond to stable conceptual categories—that are beyond the current modeling framework.

Perhaps the simplest explanation in the case of the Gender question is that the Nonrepresentative Environment statistics only partially overcame children's real-world experiences. Suppose that a child assumed that 31% of the faces were female, halfway between the true 11% base rate and the psychologically plausible 50%. In this case the Spearman rank correlation between first question frequency and the scaled information gain would increase

to .73 in the Nonrepresentative Environment, similar to the correlation in the Representative Environment.⁴ It is thus possible that the perceived proportion of female faces became closer to the true task statistics over repeated games, but that this shift was not dramatic enough to be apparent in the data.

What are the implications? Most, but not all, of the above accounts imply that experiments with novel, artificial stimuli will understate the efficiency of information search in the wild. It is therefore important to learn the extent to which each of these explanations is correct. Naturalistic stimuli, the relative representativeness of which can be manipulated, were required for the manipulation in the present study. To differentiate among the alternate explanations, however, future experiments should orthogonally manipulate physical feature salience, individual subjects' learning history, and the statistics of the immediate task environment.

4.1. Final thoughts

Both theoretical issues in the study of information acquisition, and the design of future experiments, stand to gain from bringing sequential search experimental paradigms from developmental experiments and statistical insights together. We used exhaustive search to find that in our Person Game tasks the split-half heuristic does in fact identify the most efficient question strategies. However, this is not the case in general (Hyafil & Rivest, 1976).

What search goals do people have if there is an unavoidable tradeoff between long-run efficiency and near-term information value? Meier and Blair (2013)

⁴ We thank Reviewer 2 for suggesting this analysis. Note that in the case of nonequal priors, Eq. (1) cannot be used, but the general definition of information gain (Cover & Thomas, 1991; Nelson, 2005) still applies.

found that people preferred a globally-more-efficient strategy, even if it entailed getting less information in the first query, in a situation in which a maximum of three queries were needed. In future research, one key theme to explore is whether, when, and how people identify efficient strategies in more complex sequential search tasks in which stepwise methods are suboptimal.

Acknowledgments

This research was supported by Grants NE 1713/1 to JDN, MA 1544/12 to LFM, and ME 3717/2 to BM, from the Deutsche Forschungsgemeinschaft (DFG) as part of the priority program “New Frameworks of Rationality” (SPP 1516). We are grateful to Ben Klimmek for introducing us to the Guess Who game. We thank Hasbro Germany for permission to use and reproduce the stimuli from their “Wer ist es?” (Guess Who) game. We thank Bryan Bergert, Henry Brighton, Vincenzo Crupi, Flavia Filimon, Matt Jones, Henrik Olsson, Amy Perfors, and Christine Szalay for extremely helpful comments on this research and manuscript.

References

- Amit, M., & Jan, I. (2006). Autodidactic learning of probabilistic concepts through games. In *Proceedings of the 30th conference of the international group for the psychology of mathematics education*, (Vol. 2, pp. 49–56). Charles University, Prague: PME.
- Austerweil, J. L., & Griffiths, T. L. (2011). Seeking confirmation is rational for deterministic hypotheses. *Cognitive Science*, 35, 499–526.
- Bicknell, K. (2011). *Eye movements in reading as rational behavior*. PhD dissertation, Departments of Linguistics and Cognitive Science, University of California at San Diego.
- Cover, T. M., & Thomas, J. A. (1991). *Elements of information theory*. New York, NY: Wiley.
- 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.
- Eimas, P. D. (1970). Information processing in problem solving as a function of developmental level and stimulus saliency. *Developmental Psychology*, 2, 224–229.
- Fischbein, E., Pampu, I., & Minzat, I. (1970). Comparison of ratios and the chance concept in children. *Child Development*, 41, 377–389.
- Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual review of psychology*, 62, 451–482.
- Hattori, M. (2002). A quantitative model of optimal data selection in Wason's selection task. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology*, 55, 1241–1272.
- Hirsh-Pasek, K., & Golinkoff, R. M. (2003). *Einstein never used flash cards: How our children really learn—and why they need to play more and memorize less*. Emmaus, PA: Rodale.
- Huffman, D. A. (1952). A method for the construction of minimum-redundancy codes. In *Proceedings of the I.R.E.* (pp. 1098–1102).
- Hyafil, L., & Rivest, R. L. (1976). Constructing optimal binary decision trees is NP-complete. *Information Processing Letters*, 5, 15–17.
- Itti, L., & Baldi, P. (2006). Bayesian surprise attracts human attention. In Y. Weiss, B. Schölkopf, & J. Platt (Eds.), *Advances in neural information processing systems* 18 (pp. 547–554). Cambridge, MA: MIT Press.
- Klayman, J., & Ha, Y.-W. (1987). Confirmation, disconfirmation, and information. *Psychological Review*, 94, 211–228.
- Legge, G. E., Klitz, T. S., & Tjan, B. S. (1997). Mr. Chips: An Ideal-Observer model of reading. *Psychological Review*, 104, 524–553.
- Lindley, D. V. (1956). On a measure of the information provided by an experiment. *Annals of Mathematical Statistics*, 27, 986–1005.
- Markant, D., & Gureckis, T. (2012). One piece at a time: Learning complex rules through self-directed sampling. In N. Miyake, D. Peebles, & R. Cooper (Eds.), *Proceedings of the 34th annual conference of the cognitive science society* (pp. 725–730). Austin, TX: Cognitive Science Society.
- Martignon, L., & Krauss, S. (2009). Hands-on activities for fourth graders: A tool box for decision-making and reckoning with risk. *International Electronic Journal of Mathematics Education*, 4, 227–258.
- 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. B. (2013). Waiting and weighting: Information sampling is a balance between efficiency and error-reduction. *Cognition*, 126, 319–325.
- Mosher, F. A., & Hornsby, J. R. (1966). On asking questions. In J. S. Bruner, R. R. Oliver, & P. M. Greenfield, et al. (Eds.), *Studies in cognitive growth* (pp. 86–102). New York, NY: Wiley.
- Nakamura, K. (2006). Neural representation of information measure in the primate premotor cortex. *Journal of Neurophysiology*, 96, 478–485.
- Navarro, D. J., & Perfors, A. F. (2011). Hypothesis generation, sparse categories, and the positive test strategy. *Psychological Review*, 118, 120–134.
- Najemnik, J., & Geisler, W. S. (2005). Optimal eye movement strategies in visual search. *Nature*, 434(7031), 387–391.
- Nelson, J. D. (2005). Finding useful questions: On Bayesian diagnosticity, probability, impact and information gain. *Psychological Review*, 112, 979–999.
- Nelson, J. D. (2009). Naïve optimality: Subjects' heuristics can be better-motivated than experimenters' optimal models. *Behavioral and Brain Sciences*, 32, 94–95.
- Nelson, J. D., & Chenkov, N. (2010, November). *Love at first feature learning: on the importance of salience in subjective informational value*. Program No. 532.6. 2010 Neuroscience Meeting Planner. San Diego, CA: Society for Neuroscience, 2010. Online.
- Nelson, J. D., & Cottrell, G. W. (2007). A probabilistic model of eye movements in concept formation. *Neurocomputing*, 70, 2256–2272.
- Nelson, J. D., McKenzie, C. R. M., Cottrell, G. W., & Sejnowski, T. J. (2010). Experience matters: information acquisition optimizes probability gain. *Psychological Science*, 21, 960–969.
- Nelson, J. D. (2008). Towards a rational theory of human information acquisition. In M. Oaksford & N. Chater (Eds.), *The probabilistic mind: Prospects for rational models of cognition* (pp. 143–163). Oxford: Oxford University Press.
- Nelson, J. D., Tenenbaum, J. B., & Movellan, J. R. (2001). Active inference in concept learning. In J. D. Moore & K. Stenning (Eds.), *Proceedings of the 23rd conference of the cognitive science society* (pp. 692–697). Mahwah, NJ: Erlbaum.
- 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.
- Oaksford, M., & Chater, N. (2003). Optimal data selection: Revision, review, and reevaluation. *Psychonomic Bulletin & Review*, 10, 289–318.
- Ruggeri, A., & Feufel, M. A. (2013). *How the level of inclusiveness affects object categorization*. Manuscript submitted for publication.
- Shannon, C. E. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, 27(379–423), 623–656.
- Siegler, R. S. (1977). The twenty questions game as a form of problem solving. *Child Development*, 48, 395–403.
- Siegler, R. S., & Stern, E. (1998). Conscious and unconscious strategy discoveries: A microgenetic analysis. *Journal of Experimental Psychology: General*, 127, 377–397.
- Slowiczek, 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.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.
- Thornton, S. (1982). Challenging “early competence”: A process-oriented analysis of children's classifying. *Cognitive Science*, 6, 77–100.
- Tourniaire, F., & Pulos, S. (1985). Proportional reasoning: A review of the literature. *Educational Studies in Mathematics*, 16, 181–204.
- Walker Renninger, L., Coughlan, J., Verghese, P., & Malik, J. (2005). An information maximization model of eye movements. In L. K. Saul, Y. Weiss, & L. Bottou (Eds.), *Advances in neural information processing systems* 17 (pp. 1121–1128). Cambridge, MA: MIT Press.