

THE DRIVE FOR KNOWLEDGE

The Science of Human Information Seeking

EDITED BY

IRENE COGLIATI DEZZA

University College London

ERIC SCHULZ

Max Planck Institute for Biological Cybernetics, Tübingen

CHARLEY M. WU

University of Tübingen



CAMBRIDGE
UNIVERSITY PRESS

What Makes a Good Query?
Prospects for a Comprehensive Theory of Human
Information Acquisition

Björn Meder, Vincenzo Crupi, and Jonathan D. Nelson

5.1 The Psychology of Human Information Acquisition

Searching for relevant information to support learning and reasoning is central to intelligent and goal-directed behavior. Cognitive development is guided by children's ability to actively acquire information about their physical and social environment. Doctors routinely perform tests to diagnose their patients. And, of course, scientists conduct experiments to test their theories and hypotheses.

Psychological research on how humans acquire information in a self-directed manner began in the wake of the cognitive revolution (Bruner et al., 1956; Mosher & Hornsby, 1966; Wason, 1960). Many of these earlier studies were inspired by Popper's (1959) philosophy of science and his method of falsification: that from a logical point of view scientific theories cannot be proven to be true, but they can be shown to be wrong if their predictions are inconsistent with the outcome of an experiment. Accordingly, queries (e.g., questions, tests, experiments) are only useful if they can yield data that could potentially falsify a hypothesis.

Adopting a logical framework and falsificationism as the normatively correct approach to information acquisition, psychologists devised empirical studies to find out whether people would intuitively seek out potentially disconfirming evidence. A prominent example is the *selection task*, in which participants can acquire information to test whether a conditional rule holds (Wason, 1968). Presented with four cards and a rule such as "If there is a vowel on one side of any card, then there is an even number on its other side," searchers could turn over one or multiple cards to attempt to falsify the rule. Contrary to the prescriptions of a logico-deductive method, and paralleling findings from related tasks (Wason, 1960), few participants selected queries in accordance with a falsificationist strategy. In line with

the emerging heuristics-and-biases program (Kahneman & Tversky, 1974), Wason and others suggested that people's hypothesis-testing strategies are prone to a "confirmation bias" and other suboptimalities (Klayman & Ha, 1987; Nelson & McKenzie, 2009; Nickerson, 1996).

Around the 1980s, psychologists started using a different theoretical framework for conceptualizing and evaluating human search, inspired by approaches developed in information theory (Shannon, 1948), philosophy of science (Good, 1950), statistics (Lindley, 1956), and decision theory (Savage, 1954). A key idea in this framework, foreshadowed by Chamberlin's (1890) method of multiple working hypotheses, is that the goal of scientific information acquisition is to discriminate among multiple possible hypotheses, rather than to falsify a single hypothesis. Commonly referred to as the *Optimal Experimental Design* (OED) framework, query selection is conceptualized as probabilistic inductive inference, where the outcome of a query is used to revise beliefs about the considered hypotheses. Updating beliefs in light of new data provides a certain amount of information, with different ways of measuring this quantity in relation to a searcher's beliefs and goals. These models have informed theory and empirical research on central issues in perception, developmental and cognitive psychology, and neuroscience.

Here, we discuss four key questions. First, how can we formalize the value of information and how do different models differ conceptually and mathematically? Second, how can these models support the theoretical and empirical analysis of human information acquisition? Third, what normative and computational principles govern sequential search, and how can the tension between short- and long-run optimality inform research on people's ability to conduct an efficient series of queries? Fourth, what are the relations between probabilistic (Bayesian) models of the value of information and heuristic approaches to information acquisition, and what insights can be gained from bridging different levels of analysis? We conclude by discussing current challenges and prospects for a comprehensive theory of human information search.

5.2 What Makes a Good Query?

How can we quantify the informational value of a query – a verbal question, a medical test, an eye movement, or an experiment? Several models to address this issue have been proposed in different fields, including information theory, philosophy of science, statistics, and psychology. Placed within the broad OED framework, these models

evaluate queries according to different conceptual ideas and corresponding measures, such as the expected reduction in uncertainty (Lindley, 1956) or the expected improvement in classification accuracy (Baron, 1985). The quantities underpinning OED models (e.g., entropy measures) are widely used, including applications in ecology (Crupi, 2019; Simpson, 1949), economics (Gini, 1921), machine learning (Settles, 2010), and physics (Beck, 2009). They also provide the foundation of active learning strategies used to adaptively sample high-dimensional search spaces, for instance in material sciences (Lookman et al., 2019), drug discovery (Murphy, 2011), or learning of protein interactions (Mohamed et al., 2010). How humans search large decision spaces has also been addressed in recent psychological research (Meder et al., 2021; Schulz et al., 2019; Wu et al., 2018).

In cognitive science, OED models serve as descriptive or normative models for assessing human information acquisition in different domains (for reviews, see Coenen et al., 2019; Crupi et al., 2018; Nelson, 2005; Gureckis & Markant, 2012). Key applications include perceptual tasks (Najemnik & Geisler, 2005; Nelson & Cottrell, 2007), categorization (Markant & Gureckis, 2014; Nelson et al., 2010), associative learning (Kruschke, 2008), causal induction (Bramley et al., 2015; Steyvers et al., 2003), and hypothesis testing (Austerweil & Griffiths, 2011; Crupi et al., 2009; Oaksford & Chater, 1994, 1996; Skov & Sherman, 1986). OED principles are also frequently used in developmental research (Kachergis et al., 2016; Nelson et al., 2014; Ruggeri & Lombrozo, 2015; Ruggeri et al., 2015) and neuroscience (Filimon et al., 2020; Nakamura, 2006). Applied issues include eye witness identification (Wells & Lindsay, 1980), medical diagnosis (Benish, 1999), reading (Legge et al., 1997), and the design of psychological experiments (Myung & Pitt, 2009).

5.3 Optimal Experimental Design (OED): Probabilistic Models of the Value of Information

OED models provide a formal account of the epistemic value of information and the usefulness of possible queries. Typically, OED models are defined solely in terms of the relevant probability model, and are therefore most adequate for situations where the goal of the searcher is purely epistemic. Accordingly, OED models provide *informational* (or *epistemic*) *utility functions*.

The general method for evaluating a query Q based on OED principles requires specifying the prior probabilities of the hypotheses, the likelihood

of each query outcome given the hypotheses, and a measure of the epistemic utility of each outcome. For instance, the information value of a medical test (Benish, 1999) that can come out positive or negative is fully specified by the prior probabilities of the condition (e.g., prior probability of the disease, for instance derived from epidemiological data), and the likelihoods of a positive test given that the medical condition is present or absent, respectively. Other examples are classification tasks, where different features can be queried before making a prediction (Nelson et al., 2010), or visual search tasks, where eye movements must be directed toward informative parts of a scene to find a target item (Najemnik & Geisler, 2005).

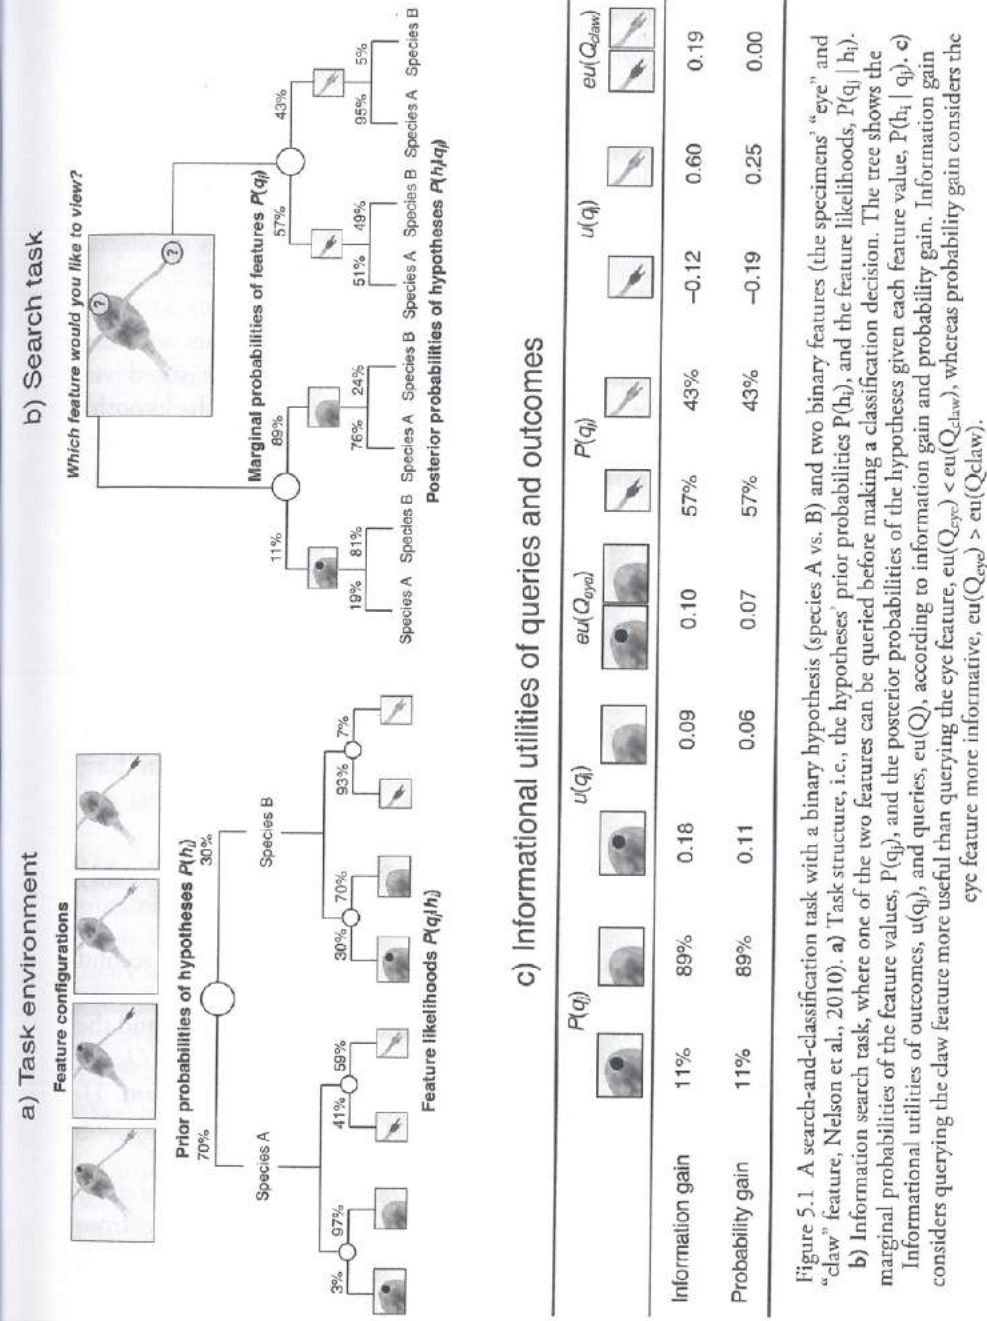
Based on the probability model, a *preposterior analysis* (Raiffa & Schlaifer, 1961) can be conducted, where the amount of obtained information is quantified by comparing the prior probability distribution over the hypotheses *before* asking query Q with the posterior distribution *after* asking Q and obtaining outcome q_j . (Capital letters refer to random variables, such as a query Q , whose answer is not yet known. Lowercase variables refer to known values: for instance, q_j is a specific obtained answer.) If a query Q has m possible answers, then its expected informational utility, $eu(Q)$, is the weighted average usefulness of its possible answers:

$$eu(Q) = \sum_{j=1}^m P(q_j)u(q_j) \quad (1)$$

where $P(q_j)$ denotes the probability of observing outcome q_j and $u(q_j)$ denotes its information value.

Figure 5.1 illustrates the basic rationale of determining queries' information value using a search-and-classification task with a binary hypothesis space (species A vs. species B) and two binary features, the specimen's "eye" and "claw" (Nelson et al., 2010). The probabilistic structure of the environment is shown in Figure 5.1a – that is, the hypotheses' prior probabilities, $P(h_i)$, and the individual feature likelihoods $P(q_j | h_i)$ (numbers from Wu et al., 2017, experiment 1). The goal is to classify the specimen as species A or B, after querying one of the two features (Figure 5.1b). Each query could yield two outcomes (feature values) with varying implications for the hypotheses' posterior probabilities – that is, for each $P(h_i | q_j)$. But how exactly can we use these quantities to determine the information value of the two queries?

Many different value-of-information models within the broad OED framework have been proposed for quantifying the value of queries. These



models can make strongly diverging claims about the informational utility of a query (Nelson, 2005). Here, we briefly introduce some prominent entropy-based models of the value of information, which constitute the most widely used class of measures across different disciplines. Importantly, there are deep mathematical relationships among seemingly different models that allow them to be treated in a unified mathematical framework (Crupi et al., 2018; Sharma & Mittal, 1975).

5.4 Information Gain: Reducing Uncertainty

Information gain (IG; Lindley, 1956) values queries in accordance with the expected reduction in uncertainty, where uncertainty is measured via Shannon (1948) entropy. The model contrasts the entropy in the hypothesis space before conducting a query with the expected entropy after the query, with the information value of a query being the expected amount of entropy reduction.

The Shannon entropy of a random variable H , $ent(H)$, corresponding to the true hypothesis (e.g., the species in Figure 5.1), is defined as

$$ent_{Shannon}(H) = \sum_{i=1}^n P(h_i) \log_2 \frac{1}{P(h_i)} \quad (2)$$

where there are n possible hypotheses h_1, h_2, \dots, h_n . (The choice of base of logarithm is arbitrary; we use \log_2 such that IG is measured in bits; Figure 5.2a.) The corresponding informational utility of query outcome q_j is

$$u_{IG}(q_j) = \sum_{i=1}^n P(h_i) \log_2 \frac{1}{P(h_i)} - \sum_{i=1}^n P(h_i|q_j) \log_2 \frac{1}{P(h_i|q_j)} \quad (3)$$

where the first term corresponds to the prior entropy, $ent(H)$, and the second term denotes the expected posterior entropy of H , given answer q_j .

Figure 5.1c shows the IG of the individual query outcomes, $u(q_j)$, and the expected informational utility of the two queries, $eu(Q_{eye})$ and $eu(Q_{claw})$, given by the weighted average utility of the outcomes (Equation 1). According to IG, querying the claw feature is more useful than querying the eye feature, $eu(Q_{eye})=0.1$ bits vs. $eu(Q_{claw})=0.19$ bits.

IG is probably the most widely used measure for quantifying uncertainty reduction in information acquisition. Applications in psychology range from hypothesis testing (Austerweil & Griffiths, 2011; Nelson et al., 2001; Oaksford & Chater, 1994, 1996) and causal induction (Bramley et al.,

2015; Steyvers et al., 2003) to the analysis of eye movements (Legge et al., 1997; Najemnik & Geisler, 2005; Nelson & Cottrell, 2007) and cognitive development (Nelson et al., 2014; Ruggeri et al., 2015).

5.5 Probability Gain: Reducing Prediction Error

Probability gain (PG; Baron, 1985) measures improvement in classification accuracy. It can alternately be viewed as measuring a kind of information gain (see Crupi et al., 2018) in which error entropy (classification error), rather than Shannon entropy, is used to measure uncertainty. PG values answers in accordance with how much they improve the probability of making a correct prediction, assuming that if a searcher has to guess which hypothesis is correct (e.g., which disease is present or the species of a biological specimen) they will select the most likely hypothesis given all information obtained so far. The prediction error, $1 - \max P(h_i)$, provides the entropy function (Figure 5.2a):

$$ent_{Error}(H) = 1 - \max_{1 \leq i \leq n} P(h_i), \quad (4)$$

with the corresponding epistemic utility of a query outcome defined as

$$\begin{aligned} u_{Error}(q_j) &= \max_{1 \leq i \leq n} P(h_i|q_j) - \max_{1 \leq i \leq n} P(h_i) \\ &= \left(1 - \max_{1 \leq i \leq n} P(h_i)\right) - \left(1 - \max_{1 \leq i \leq n} P(h_i|q_j)\right). \end{aligned} \quad (5)$$

In Figure 5.1, the probability of making a correct prediction before obtaining additional information is 70%, assuming one picks the most likely hypothesis (species A). According to PG, querying the eye feature is more useful than querying the claw feature, as the eye feature leads to 77% classification accuracy, $eu(Q_{eye})=0.07$ (Figure 5.1c). In contrast, the claw feature does not improve on the initial 70% accuracy based on the prior probabilities alone, $eu(Q_{eye})=0$. Thus, a reduction in uncertainty measured via Shannon entropy does not necessarily reduce prediction error.

PG has been widely used in psychology, including probabilistic multiple-cue categorization tasks (Nelson et al., 2010; Meder & Nelson, 2012; Wu et al., 2017), medical diagnostic test selection (Baron et al., 1988), and reasoning in the pseudodiagnosticity paradigm (Crupi et al., 2009). This research has also revealed important constraints on the ability of OED models to predict human behavior,

namely the ways in which probability information is conveyed. In experience-based category learning with hundreds of learning trials, PG provides a robust predictive model of human query selection, compared to several other models (Nelson et al., 2010). With other presentation formats, for instance when probabilities are communicated with words-and-numbers or using visualizations of frequency information, the proportion of people searching in accordance with probability gain varies strongly, being as low as about 20% and as high as about 90% (Wu et al., 2017; where chance would be 50%). Thus, the predictive power of OED models to account for behavior cannot be judged in isolation, but requires a precise characterization of the task circumstances, especially in relation to how environmental probabilities are conveyed.

5.6 A Unifying Framework for Entropy and the Value of Information

5.6.1 The Sharma-Mittal-Space of Entropy Measures

Information gain and probability gain are two prominent measures of the epistemic utility of queries, with different entropies underpinning them. However, a variety of further entropy measures exist and are being used in different fields, including the measurement of species diversity in ecology (Keylock, 2005), genetic variability in evolutionary biology (Lewontin, 1972), and applications in physics (Beck, 2009). Different entropies offer alternative ways to formalize the notion of uncertainty contained in a probability distribution, with distinct mathematical properties of interest for theory and application (Crupi & Tentori, 2014; Nelson 2005). Analogously to Shannon and error entropy, these measures give rise to different informational utility functions, with the value of queries valued in terms of their ability to reduce alternative forms of entropy (Crupi et al., 2018).

One prominent model is Hartley (1928) entropy, which was an important precursor in the development of Shannon's seminal work:

$$ent_{Hartley}(H) = \log \sum_{i=1}^n P(h_i)^0. \quad (6)$$

Hartley entropy computes the logarithm of the number of non-zero-probability hypotheses (assuming that $0^0=1$ and $P(h_i)^0=1$ for all

$P(h_i) > 0$; Figure 5.2a). A value-of-information model based on Hartley entropy resembles Popper's (1959) method of falsification, as queries are valued by their ability to potentially rule out (falsify) at least one of the hypotheses considered. (Note though that ruling out any hypothesis – and not only the target hypothesis – would be equally valued in this information gain model.) If no outcome can potentially rule out a hypothesis, the query has no epistemic utility. This would be the case for the scenario in Figure 5.1, where no outcome yields certainty about the hypotheses.

Another important measure is *quadratic entropy* (Vajda & Zvárová, 2007; also known as *Gini-* or *Simpson-Gini index*), formally defined as

$$ent_{quadratic}(H) = 1 - \sum_{i=1}^n P(h_i)^2. \quad (7)$$

Quadratic entropy (Figure 5.2a) has been widely applied in many fields, including in economics to quantify income inequality (Gini, 1921) and in biology to measure species diversity (Patil & Taillie, 1982; Simpson, 1949). It is also frequently used in machine learning (Settles, 2010).

Different entropy measures are based on diverse conceptual and mathematical ideas about uncertainty and information. Notwithstanding these differences, several entropy measures arise as special cases in a unifying mathematical framework: the *Sharma-Mittal space of entropies* (SM; Sharma & Mittal, 1975; Crupi et al., 2018). The generalized SM-entropy measure has two parameters, the order parameter r and the degree parameter t :

$$ent_{SM}(H) = \frac{1}{t-1} \left[1 - \left(\sum_{i=1}^n P(h_i)^r \right)^{\frac{t-1}{r}} \right]. \quad (8)$$

The family of Sharma-Mittal entropies contains several other entropy measures as special cases (many of which exist only for well-defined limits): Shannon entropy is recovered for $r = t = 1$, Hartley entropy for $r = 0$ and $t = 1$, Quadratic entropy for $r = t = 1$, and Error entropy for $t = 2$ and $r \rightarrow \infty$ (Figure 5.2b). The family of Rényi (1961) entropies is recovered for $t = 1$, and the family of Tsallis (1988) entropies for $r = t$. Each of these entropy measures can be used to implement a corresponding model of the value of information that values query in accordance with the expected reduction of the chosen entropy.

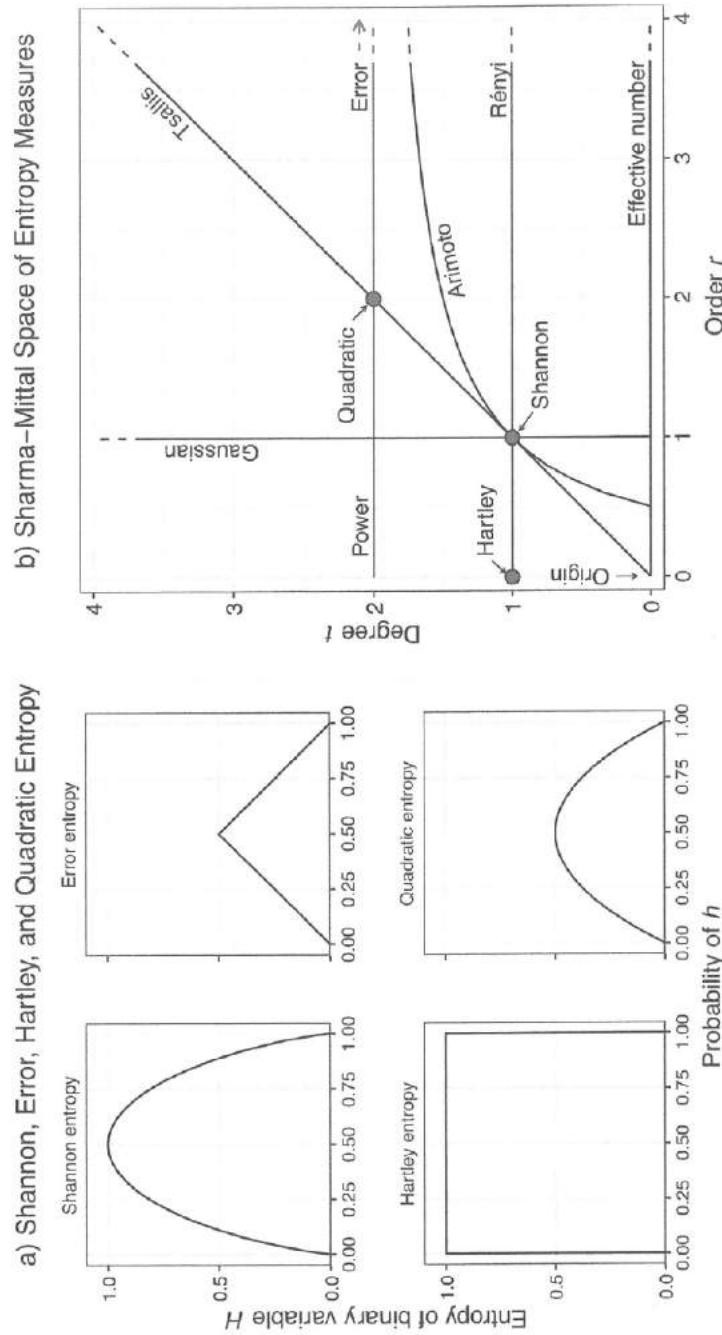


Figure 5.2 Example of entropy functions and the biparametric Sharma-Mittal family of entropy (Crupi et al., 2018). a) Examples of four prominent entropy functions for a binary hypothesis space. b) Sharma-Mittal-Space with order parameter r and degree parameter t . Each point represents a specific entropy measure (e.g., Shannon entropy); each line corresponds to a distinct generalized entropy function, such as the families of Rényi (1961) and Tsallis (1988) entropies.

In addition to providing a unified framework for quantifying entropy and uncertainty, the Sharma-Mittal framework offers prospects for cognitive modeling and empirical research. For instance, the logico-deductive interpretation of what constitutes rational information acquisition in Wason's (1968) selection task can be recovered in the Sharma-Mittal framework by using the reduction of Hartley entropy as a measure of epistemic utility (Crupi et al., 2018). Such analyses help to bridge seemingly incommensurate perspectives and models in a unifying mathematical framework, thereby fostering theory integration and providing guidance for empirical studies.

5.7 Short- vs. Long-Run Optimality of OED Models

OED models assess the epistemic utility of queries by considering each question's potential outcomes and their immediate implications for the hypotheses. These models were originally envisaged for situations in which only a single query can be conducted and thus they are *stepwise-optimal* in the sense that they optimize a particular informational utility function in the next time step. Since they disregard any future queries that could be conducted, such methods are also known as *myopic* or *greedy* models.

However, in many situations multiple queries can be conducted, for instance in medical diagnosis. In such cases, it is critical to search *efficiently* for information – to obtain as much information as possible with as few queries as possible (e.g., because medical tests can have intrinsic harms and are costly). Are stepwise OED models optimal in the long run when multiple queries can be conducted? The answer to this question is in the negative: Stepwise-optimal methods can provide a reasonable and tractable account for specific scenarios (Meder et al., 2019; Nelson et al., 2018), but they do not generally identify the most efficient sequence of queries – which in fact is a computationally intractable problem (Hyafil & Rivest, 1976).

The fact that what is in some sense optimal for a one-shot search decision can be distinctly suboptimal when multiple queries can be conducted has important implications for the theoretical and empirical analysis of human sequential search. For instance, stepwise models correspond to a valuation of individual queries, whereas planning a sequence of queries requires evaluating alternative search trees comprising multiple queries. However, in stark contrast to other domains such as reward-based learning (Bellman, 1957; Sutton & Barto, 1998), there is little research on the conflict between short- and long-term optimality and people's ability to

search efficiently in situations where stepwise methods are distinctly suboptimal.

An exception is Meier and Blair (2013), who used a probabilistic classification task with four hypotheses and three (binary) queries, where the most informative first query according to several stepwise methods was not the most efficient first query in the long run (Figure 5.3a). Searchers were able to identify the most efficient query, but often required dozens or hundreds of learning trials to consistently select the most efficient feature first. These results show that people can learn from experience to search efficiently and overcome the limitations of stepwise approaches. However, further research is needed to achieve a better understanding of the interplay between learning processes, normative constraints imposed by sequential search scenarios, and the cognitive processes underlying query selection when stepwise methods are suboptimal.

The limitations of stepwise methods also apply in simpler situations with uniform priors over the hypotheses and deterministic likelihoods. Figure 5.3b shows a sequential search scenario based on the twenty-questions game (Meder et al., 2019), an experimental paradigm widely used in developmental and cognitive psychology to investigate sequential search in children and adults. The goal is to identify a randomly chosen hypothesis (here: monster) by asking as few yes–no questions about its features as possible (e.g., “Is the monster blue?”). Through a series of questions searchers move from a state of maximal uncertainty with equiprobable hypotheses to knowing the true hypothesis with certainty.

A common benchmark on this task is stepwise information gain (Kachergis et al., 2017; Nelson et al., 2014; Ruggeri et al., 2016, 2017). However, selecting questions according to their maximal stepwise information gain does not necessarily entail conducting as few queries as possible. For instance, in the task environment shown in Figure 5.3b, the color and shape feature tie for the maximum information gain query, but starting with the color feature is more efficient in the long run, because it allows for more informative queries on subsequent steps (Meder et al., 2019). Thus, similar to Meier and Blair (2013), selecting queries in accordance with the highest stepwise expected reduction of Shannon entropy (as well as various other entropy models) fails to identify the most efficient first query. Interestingly, in this task both children and adults show very limited sensitivity to long-run considerations, with behavior better accounted for by stepwise models (Meder et al., 2019). One explanation is that learning from experience is required to search efficiently across multiple queries.




Figure 5.3 Suboptimality of stepwise methods and relationships between information gain and the split-half heuristic. a) Search trees in the classification task of Meier and Blair (2013). Q_1 to Q_3 denote possible queries (binary physical features of stimuli), the h_i denote four possible hypotheses (categories); + and – indicate the observed feature values. According to stepwise information gain (IG) and probability gain (PG), Q_2 is the most informative first query. In contrast, the globally optimal (=most efficient) search tree starts with Q_1 . b) Stimuli from a twenty-questions game where stepwise-optimal methods fail to identify the most efficient first query (Meder et al., 2019). c) Relation between the split induced by a query and its stepwise information gain.

Another possibility is the use of heuristic methods that may somehow correspond to stepwise OED models.

The fact that selecting queries in accordance with stepwise methods does not ensure long-run efficiency is a critical issue for both the theoretical and the empirical analysis of human information acquisition, because it entails that stepwise OED models do not generally provide adequate benchmarks for evaluating people's search efficiency. While globally optimal solutions are generally unattainable due to the quickly increasing computational complexity, for smaller numbers of hypotheses and queries the most efficient search tree(s) can be determined through dynamic programming, such that empirical paradigms can be constructed that dissociate between short- and long-run efficiency (Meder et al., 2019; Nelson et al., 2018). Such paradigms also enable contrasting stepwise and multistep models of information acquisition, which can be more efficient because they plan multiple queries ahead (Nelson et al., 2018). To what extent multistep approaches can predict human sequential search is an important open question for developing a comprehensive theory of human sequential search.

5.8 From Heuristics to OED Models and Back

In cognitive science, OED models often serve as normative benchmarks that specify the problem faced by the searcher and the quantities required to solve the task from a computational point of view. Descriptively, OED models are able to predict human behavior well on a wide range of tasks, raising questions about the nature of the cognitive processes underlying query selection. It could be that the brain indeed carries out the relevant computations without searchers being consciously aware of these processes (e.g., in visual search tasks that recruit a dedicated neural machinery), but typically OED methods are considered computational-level models of behavior which provide a "rational description" rather than a "rational calculation" (Chater et al., 2003).

One promising approach to bridge different levels of analysis is to consider heuristic strategies that choose queries without resorting to the explicit calculation of the quantities underpinning OED models. In psychology, the use of heuristics is often assumed to lead to systematic errors in thinking and reasoning (Tversky & Kahneman, 1974). In a similar vein, it has frequently been argued that human search behavior does not conform to (supposedly) normative principles and suffers from a "confirmation bias" or other peculiarities. However, research has also

revealed that the performance of heuristics can be on par with more computationally heavy models and sometimes even outperform them (Gigerenzer & Gaissmaier, 2011). A similar picture has emerged in the psychology of human inquiry, where intriguing connections between OED models and simple heuristics for query selection have been identified.

Consider the *split-half heuristic* (Navarro & Perfors, 2011; Nelson et al., 2014), which is applicable in the twenty-questions game (Figure 3b). This strategy does not rely on explicitly computing entropy, but rather evaluates queries in accordance with how they split the current set of hypotheses (items). According to the split-half heuristic, the closer a query comes to inducing a 50:50 split, the more informative it is. In Figure 3b, color and shape would be considered the most informative queries, as they induce a more even split (3:9) than the pattern feature (2:10) or querying individual monsters (1:11). Several papers have reported that human searchers frequently rely on this strategy, with the propensity to select features in accordance with the evenness of the induced split increasing with age (e.g., Denney & Denney, 1973; Mosher & Hornsby, 1966; Siegler, 1977).

How does the split-half heuristic relate to OED theories? It turns out that in the twenty-questions game the stepwise information gain of a query is a direct function of the induced split: The closer the split comes to 50:50, the higher the information gain of the query (Figure 5.3c). Thus, a searcher using the split-half heuristic will invariably select questions in accordance with their information gain. The analytic relation between the informativeness of a query and induced split also holds for several other entropy measures (Crupi et al., 2018) and particular nonentropic OED models (Nelson et al., 2018).

Another heuristic is the *likelihood difference heuristic* (Nelson, 2005; Nelson et al., 2020), which is sometimes used in binary classification tasks when probability information is conveyed through words and numbers (Slowiczek et al., 1992). The heuristic is applicable in probabilistic search tasks with binary hypotheses and features, and values queries in accordance with the absolute difference in outcome likelihoods. Ranking queries in accordance with their likelihood difference corresponds to ranking queries in accordance with the OED model impact (Wells & Lindsay, 1980), which for the special case of equal priors also corresponds to ranking queries according to probability gain (Nelson, 2005). In particular circumstances the likelihood difference heuristic is provably optimal for selecting queries, whereas several OED models are not (Nelson et al., 2020).

5.9 Concluding Remarks: Open Questions and Challenges

Recent years have seen strong progress in building a comprehensive computational account of human inquiry based on applying OED principles to the theoretical and empirical analysis of cognition and behavior. What challenges need to be met to further advance theory, and what are important venues for empirical research?

One important issue concerns the relation between informational utilities and other kinds of reward structures. For instance, in medical diagnosis the information provided by a test is of course essential, but further costs and rewards are also highly relevant (Baron & Hershey, 1988; Nelson et al., 2020). This is because tests themselves can be intrinsically harmful (e.g., radiation) or because their outcomes have strong implications for subsequent actions (e.g., erroneously failing to quarantine an infectious person due to a false test result). While OED models evaluate queries in terms of their epistemic value, the general approach can be extended to incorporate situation-specific utilities (e.g., rewards and costs associated with decisions based upon obtained information). In this case, queries can be valued in accordance with their expected *utility gain*, rather than pure information value (Markant & Gureckis, 2012; Meder & Nelson, 2012). Incorporating other rewards and costs in the analysis of query selection might require invoking additional methods and considerations (Sharot & Sunstein, 2020). For instance, a doctor may need to choose between a highly informative test whose results may take several days to obtain, and a less informative test whose results are immediately available. Investigating how searchers balance information value and time (i.e., temporal discounting) and how they integrate other factors (e.g., hedonic value of query outcomes and other emotions related to the anticipation of information; Li et al., 2020) are important challenges for a broader characterization of information acquisition when searchers' goals are not purely epistemic.

A promising venue for applying the Sharma–Mittal framework to such questions is *deliberate ignorance* – the observation that people sometimes prefer to remain in a state of uncertainty rather than seeking out information that would resolve it (Hertwig & Engel, 2016). People may not want to know the sex of their child before birth or prefer to not know whether they have a genetic predisposition for a disease. Such behaviors raise interesting questions for theories of human inquiry, because OED models based on entropy reduction are usually strictly nonnegative, meaning that in the expectation a query can never have negative epistemic utility. The Sharma–Mittal framework, however, has the expressive power to devise models where queries can

have negative information value, depending on which entropy measure is chosen to represent uncertainty and information (Crupi et al., 2018). Thus, the framework provides novel pathways to model situations where searchers consider information as contrary to their goals even in strictly informational terms and without introducing situation-specific costs.

Another key issue is the interplay between reinforcement learning and human inquiry. This is of particular importance in light of empirical findings showing that searchers can harness their learning experiences to overcome the limitations of stepwise methods (Meier & Blair, 2013), whereas absent such experiences behavior is better accounted for by stepwise models (Meder et al., 2019). Similarly, OED models provide a much better account of human behavior when learning environmental probabilities through experience, compared to other means of conveying probability information (Meder & Nelson, 2012; Nelson et al., 2010, 2017). Within the reinforcement learning framework (Sutton & Barto, 1998), inquiry could be modeled by using informational utilities as reward functions, extending the approach from reward-based to information-based reinforcement learning (Li et al., 2019). Importantly, the reinforcement learning framework offers different proposals to address the tension between short- and long-run optimality by assuming that future outcomes are temporally discounted. Stepwise OED methods could be implemented by setting the temporal discounting parameter to zero, such that only the immediate implications of query outcomes are considered. Multistep models with a longer planning horizon could be implemented by increasing the temporal discounting factor, in which case later inquiry steps are also considered (Butko & Movellan, 2010). Integrating ideas and models from the OED and reinforcement learning literature would offer new perspectives on the learning processes and how they are constrained by OED principles.

Besides these pathways for future research, several open questions about the psychology of human inquiry and OED models exist (Coenen et al., 2019). For instance, whereas OED models require a well-defined hypothesis space and probability model, many forms of human inquiry are more open ended in that the goal is not to discriminate between existing hypotheses, but to *generate* hypotheses in the first place, often in the absence of explicit learning goals or situation-specific incentives (e.g., self-guided inquiry in childhood or science). While the broad goal in such cases may also be epistemic in the sense of reducing uncertainty about the world more generally, fully accounting for such forms of open-ended inquiry might require integrating OED methods with other frameworks, such as theories of curiosity (Dubey & Griffiths, 2020). Conversely, applying

OED methods can be challenging in practice – for instance, in scientific experimentation, when the goal is to discriminate among complex models with free parameters and the data-generating distribution is unknown (Kleinegesse & Gutmann, 2020).

In sum, OED theories provide a computational framework for evaluating the informational value of different kinds of queries, including verbal questions, eye movements, medical tests, and experiments. Recent decades have yielded substantial progress in the understanding of the computational and behavioral principles underlying human information acquisition in different domains, from neuroscience to developmental and cognitive psychology. Notwithstanding these advances and breadth of applicability, many open questions remain, and addressing these challenges will be key to making further progress toward a comprehensive theory of human information acquisition.

References

- Austerweil, J. L., & Griffiths, T. L. (2011). Seeking confirmation is rational for deterministic hypotheses. *Cognitive Science*, 35(3), 499–526. <https://doi.org/10.1111/j.1551-6709.2010.01161.x>.
- Baron, J. (1985). *Rationality and intelligence*. Cambridge University Press.
- Baron, J., & Hershey, J. C. (1988). Heuristics and biases in diagnostic reasoning: I. Priors, error costs, and test accuracy. *Organizational Behavior and Human Decision Processes*, 41(2), 259–279. [https://doi.org/10.1016/0749-5978\(88\)90030-1](https://doi.org/10.1016/0749-5978(88)90030-1).
- Beck, C. (2009). Generalised information and entropy measures in physics. *Contemporary Physics*, 50(4), 495–510. <https://doi.org/10.1080/00107510902823517>.
- Bellman, R. (1957). *Dynamic programming*. Princeton University Press.
- Benish, W. A. (1999). Relative entropy as a measure of diagnostic information. *Medical Decision Making*, 19(2), 202–206. <https://doi.org/10.1177/0272989X9901900211>.
- 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(3), 708–731. <https://doi.org/10.1037/xlm0000061>.
- Bruner, J. S., Goodnow, J. J., & Austin, G. A. (1956). *A study of thinking*. John Wiley and Sons.
- Butko, N. J., & Movellan, J. R. (2010). Infomax control of eye movements. *IEEE Transactions on Autonomous Mental Development*, 2(2), 91–107. <https://doi.org/10.1109/TAMD.2010.2051029>.
- Chamberlin, T. C. (1890). The method of multiple working hypotheses. *Science*, 15, 92–96. <https://doi.org/10.1126/science.148.3671.754>.
- Chater, N., Oaksford, M., Nakisa, R., & Redington, M. (2003). Fast, frugal, and rational: How rational norms explain behavior. *Organizational Behavior and Human Decision Processes*, 90(1), 63–86. [https://doi.org/10.1016/S0749-5978\(02\)00508-3](https://doi.org/10.1016/S0749-5978(02)00508-3).
- Coenen, A., Nelson, J. D., & Gureckis, T. M. (2019). Asking the right questions about the psychology of human inquiry: Nine open challenges. *Psychonomic Bulletin & Review*, 26, 1548–1587. <https://doi.org/10.3758/s13423-018-1470-5>.
- Crupi V. (2019). Measures of biological diversity: Overview and unified framework. In E. Casetta, J. Marques da Silva, and D. Vecchi (Eds.), *From assessing to conserving biodiversity* (pp. 123–136). Springer.
- 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(5), 1410–1456. <https://doi.org/10.1111/cogs.12613>.
- Crupi, V., & Tentori, K. (2014). State of the field: Measuring information and confirmation. *Studies in History and Philosophy of Science Part A*, 47, 81–90. <https://doi.org/10.1016/j.shpsa.2014.05.002>.
- Crupi, V., Tentori, K., & Lombardi, L. (2009). Pseudodiagnosticity revisited. *Psychological Review*, 116(4), 971–985. <https://doi.org/10.1037/a0017050>.
- Denney, D. R., & Denney, N. W. (1973). The use of classification for problem solving: A comparison of middle and old age. *Developmental Psychology*, 9(2), 275–278. <https://doi.org/10.1037/h0035092>.
- Dubey, R., & Griffiths, T. L. (2020). Reconciling novelty and complexity through a rational analysis of curiosity. *Psychological Review*, 127(3), 455–476. <https://doi.org/10.1037/rev0000175>.
- Filimon, F., Nelson, J. D., Sejnowski, T. J., Sereno, M. I., & Cottrell, G. W. (2020). The ventral striatum dissociates information expectation, reward anticipation, and reward receipt. *Proceedings of the National Academy of Sciences*, 117(26), 15200–15208. <https://doi.org/10.1073/pnas.1911778117>.
- Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual Review of Psychology*, 62, 451–482. <https://doi.org/10.1146/annurev-psych-120709-145346>.
- Gini, C. (1921). Measurement of inequality of incomes. *The Economic Journal*, 31(121), 124–126. <https://doi.org/10.2307/2223319>.
- Good, I. J. (1950). *Probability and the weight of evidence*. Charles Griffin & Co.
- Gureckis, T. M., & Markant, D. B. (2012). Self-directed learning: A cognitive and computational perspective. *Perspectives on Psychological Science*, 7(5), 464–481. <https://doi.org/10.1177/1745691612454304>.
- Hartley, R. V. (1928). Transmission of information. *Bell System Technical Journal*, 7(3), 535–563.
- Hertwig, R., & Engel, C. (2016). Homo Ignorans: Deliberately choosing not to know. *Perspectives on Psychological Science*, 11(3), 359–372. <https://doi.org/10.1177/1745691616635594>.
- Hyafil, L., & Rivest, R. L. (1976). Constructing optimal binary decision trees is NP-complete. *Information Processing Letters*, 5(1), 15–17. [https://doi.org/10.1016/0020-0190\(76\)90095-8](https://doi.org/10.1016/0020-0190(76)90095-8).

- Kachergis, G., Berends, F., Kleijn, R. D., & Hommel, B. (2016). Human reinforcement learning of sequential action. In A. Papafragou, D. Grodner, & D. Mirman (Eds.), *Proceedings of the 38th Annual Meeting of the Cognitive Science Society (CogSci 2016)*, pp. 193–198.
- Kachergis, G., Rhodes, M., & Gureckis, T. (2017). Desirable difficulties during the development of active inquiry skills. *Cognition*, 166, 407–417. <https://doi.org/10.1016/j.cognition.2017.05.021>.
- Keylock, C. J. (2005). Simpson diversity and the Shannon–Wiener index as special cases of a generalized entropy. *Oikos*, 109(1), 203–207. <https://doi.org/10.1111/j.0030-1299.2005.13735.x>.
- Klayman, J., & Ha, Y.-W. (1987). Confirmation, disconfirmation, and information in hypothesis testing. *Psychological Review*, 94(2), 211–228. <https://doi.org/10.1037/0033-295X.94.2.211>.
- Kleinegesse, S., & Gutmann, M. U. (2020). Bayesian experimental design for implicit models by mutual information neural estimation. *Proceedings of Machine Learning Research*, 119, 5316–5326.
- Kruschke, J.K. (2008). Bayesian approaches to associative learning: From passive to active learning. *Learning & Behavior*, 36, 210–226. <https://doi.org/10.3758/LB.36.3.210>.
- Legge, G. E., Klitz, T. S., & Tjan, B. S. (1997). Mr. Chips: An ideal-observer model of reading. *Psychological Review*, 104(3), 524–553. <https://doi.org/10.1037/0033-295X.104.3.524>.
- Lewontin R.C. (1972) The apportionment of human diversity. In T. Dobzhansky, M. K. Hecht, and W. C. Steere (Eds.), *Evolutionary biology* (pp. 381–398). New York, NY: Springer. https://doi.org/10.1007/978-1-4684-9063-3_14.
- Li, S., Sun, Y., Liu, S., Sun, Y., Gureckis, T. M., & Bramley, N. R. (2019). Active physical inference via reinforcement learning. *Proceedings of the Cognitive Science Society* (pp. 2126–2132). Austin, TX: Cognitive Science Society.
- Li, Z., Bramley, N. R., & Gureckis, T. M. (2020). Expectations about future learning influence moment-to-moment feelings of suspense. <https://doi.org/10.31234/osf.io/532tw>.
- Lindley, D. V. (1956). On a measure of the information provided by an experiment. *The Annals of Mathematical Statistics*, 27(4), 986–1005.
- Lookman, T., Balachandran, P. V., Xue, D., & Yuan, R. (2019). Active learning in materials science with emphasis on adaptive sampling using uncertainties for targeted design. *NPJ Computational Materials*, 5(1), 1–17. <https://doi.org/10.1038/s41524-019-0153-8>.
- Markant, D., & Gureckis, T. M. (2012). Does the utility of information influence sampling behavior? *Proceedings of the 34th Annual Conference of the Cognitive Science Society* (pp. 719–724). Austin, TX: Cognitive Science Society.
- 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(1), 94–122. <https://doi.org/10.1037/a0032108>.

- Meder, B., & Nelson, J. D. (2012). Information search with situation-specific reward functions. *Judgment and Decision Making*, 7(2), 119–148.
- Meder, B., Nelson, J. D., Jones, M., & Ruggeri, A. (2019). Stepwise versus globally optimal search in children and adults. *Cognition*, 191, Article 103965. <https://doi.org/10.1016/j.cognition.2019.05.002>.
- Meder, B., Wu, C. M., Schulz, E., & Ruggeri, A. (2021). Development of directed and random exploration in children. *Developmental Science*. e13095. <https://doi.org/10.1111/desc.13095>.
- Meier, K. M., & Blair, M. R. (2013). Waiting and weighting: Information sampling is a balance between efficiency and error-reduction. *Cognition*, 126 (2), 319–325. <https://doi.org/10.1016/j.cognition.2012.09.014>.
- Mohamed, T. P., Carbonell, J. G., & Ganapathiraju, M. K. (2010). Active learning for human protein-protein interaction prediction. *BMC Bioinformatics*, 11, S57. <https://doi.org/10.1186/1471-2105-11-S1-S57>.
- 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). Wiley.
- Murphy, R. F. (2011). An active role for machine learning in drug development. *Nature Chemical Biology*, 7(6), 327–330. <https://doi.org/10.1038/nchembio.576>.
- Myung, J. I., & Pitt, M. A. (2009). Optimal experimental design for model discrimination. *Psychological Review*, 116(3), 499–518. <https://doi.org/10.1037/a0016104>.
- Najemnik, J., & Geisler, W. (2005). Optimal eye movement strategies in visual search. *Nature*, 434, 387–391. <https://doi.org/10.1038/nature03390>.
- Nakamura, K. (2006). Neural representation of information measure in the primate premotor cortex. *Journal of Neurophysiology*, 96(1), 478–485. <https://doi.org/10.1152/jn.01326.2005>.
- Navarro, D. J., & Perfors, A. F. (2011). Hypothesis generation, sparse categories, and the positive test strategy. *Psychological Review*, 118(1), 120–134. <https://doi.org/10.1037/a0021110>.
- Nelson, J. D. (2005). Finding useful questions: On Bayesian diagnosticity, probability, impact, and information gain. *Psychological Review*, 112(4), 979–999. <https://doi.org/10.1037/0033-295X.112.4.979>.
- Nelson, J. D., & Cottrell, G. W. (2007). A probabilistic model of eye movements in concept formation. *Neurocomputing*, 70, 2256–2272. <https://doi.org/10.1016/j.neucom.2006.02.026>.
- 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(1), 74–80.
- Nelson, J. D., & McKenzie, C. R. M. (2009). Confirmation bias. In M. Kantan (Ed.), *The Encyclopedia of Medical Decision Making* (pp. 167–171). Sage.
- Nelson, J. D., McKenzie, C. R. M., Cottrell, G. W., & Sejnowski, T. J. (2010). Experience matters: Information acquisition optimizes probability gain. *Psychological Science*, 21(7), 960–969. <https://doi.org/10.1177/0956797610372637>.

- Nelson, J. D., Meder, B., & Jones, M. (2018). Towards a theory of heuristic and optimal planning for sequential information search. *PsyArXiv*.
- Nelson, J. D., Rosenauer, C., Crupi, V., Tentori, K., & Meder, B. (2020). *On the likelihood difference heuristic and the objective utility of possible medical tests*. Manuscript submitted for publication.
- 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*, 692–697.
- Nickerson, R. S. (1996). Hempel's paradox and Wason's selection task: Logical and psychological puzzles of confirmation. *Thinking & Reasoning*, 2, 1–31. <https://doi.org/10.1080/135467896394546>.
- Oaksford, M., & Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review*, 101(4), 608–631. <https://doi.org/10.1037/0033-295X.101.4.608>.
- Oaksford, M., & Chater, N. (1996). Rational explanation of the selection task. *Psychological Review*, 103(2), 381–391. <https://doi.org/10.1037/0033-295X.103.2.381>.
- Patil, G. P., & Taillie, C. (1982). Diversity as a concept and its measurement. *Journal of the American Statistical Association*, 77(379), 548–561. <https://doi.org/10.1080/01621459.1982.10477845>.
- Popper, K. R. (1959). *The logic of scientific discovery*. Hutchinson.
- Raiffa, H., & Schlaifer, R. O. (1961). *Applied statistical decision theory*. Cambridge, MA: Division of Research, Graduate School of Business Administration, Harvard University.
- Rényi, A. (1961). On measures of entropy and information. In J. Neyman (Ed.), *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability I* (pp. 547–556). University of California Press.
- Ruggeri, A., & Lombrozo, T. (2015). Children adapt their questions to achieve efficient search. *Cognition*, 143, 203–216. <https://doi.org/10.1016/j.cognition.2015.07.004>.
- Ruggeri, A., Lombrozo, T., Griffiths, T. L., & Xu, F. (2016). Sources of developmental change in the efficiency of information search. *Developmental Psychology*, 52(12), 2159–2173. <https://doi.org/10.1037/dev0000240>.
- Ruggeri, A., Sim, Z. L., & Xu, F. (2017). "Why is Toma late to school again?" Preschoolers identify the most informative questions. *Developmental Psychology*, 53(9), 1620.
- Savage, L. J. (1954). *The foundations of statistics*. Wiley.
- Schulz, E., Wu, C. M., Ruggeri, A., & Meder, B. (2019). Searching for rewards like a child means less generalization and more directed exploration. *Psychological Science*, 30(11), 1561–1572. <https://doi.org/10.1177/0956797619863663>.
- Settles, B. (2010). *Active learning literature survey*. Technical Report, University of Wisconsin-Madison.
- Shannon, C. E. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, 27, 379–423. <https://doi.org/10.1002/j.1538-7305.1948.tb01338.x>.

- Sharma, B., & Mittal, D. (1975). New non-additive measures of entropy for discrete probability distributions. *Journal of Mathematical Sciences*, 10, 28–40.
- Sharot, T., & Sunstein, C. R. (2020). How people decide what they want to know. *Nature Human Behaviour*, 4, 14–19. <https://doi.org/10.1038/s41562-019-0793-1>.
- Siegler, R. S. (1977). The twenty questions game as a form of problem solving. *Child Development*, 395–403.
- Simpson, E. H. (1949). Measurement of diversity. *Nature*, 163, 688. <https://doi.org/10.1038/163688a0>.
- Skov, R. B., & Sherman, S. J. (1986). Information-gathering processes: Diagnosticity, hypothesis-confirmatory strategies, and perceived hypothesis confirmation. *Journal of Experimental Social Psychology*, 103, 278–282. <https://doi.org/10.1016/j.beproc.2014.01.014>.
- 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(4), 392–405. <https://doi.org/10.3758/BF03210923>.
- Steyvers, M., Tenenbaum, J. B., Wagenmakers, E. J., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, 27, 453–489. [https://doi.org/10.1016/S0364-0213\(03\)00010-7](https://doi.org/10.1016/S0364-0213(03)00010-7).
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.
- Tsallis, C. (1988). Possible generalization of Boltzmann-Gibbs statistics. *Journal of Statistical Physics*, 52(1–2), 479–487. <https://doi.org/10.1007/BF01016429>.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1124–1131. <https://doi.org/10.1126/science.185.4157.1124>.
- Vajda I. & Zvárová J. (2007). On generalized entropies, Bayesian decisions, and statistical diversity. *Kybernetika*, 43(5), 675–696.
- Wason, P. C. (1960). On the failure to eliminate hypotheses in a conceptual task. *Quarterly Journal of Experimental Psychology*, 12, 129–140. <https://doi.org/10.1080/17470216008416717>.
- Wason, P. C. (1968). Reasoning about a rule. *Quarterly Journal of Experimental Psychology*, 20, 273–281. <https://doi.org/10.1080/14640746808400161>.
- Wells, G. L., & Lindsay, R. C. (1980). On estimating the diagnosticity of eyewitness nonidentifications. *Psychological Bulletin*, 88(3), 776–784. <https://doi.org/10.1037/0033-2909.88.3.776>.
- 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(8), 1274–1297. <https://doi.org/10.1037/xlm0000374>.
- Wu, C., Schulz, E., Speekenbrink, M., Nelson, J. D., & Meder, B. (2018). Generalization guides human exploration in vast decision spaces. *Nature Human Behavior*, 2, 915–924. <https://doi.org/10.1038/s41562-018-0467-4>.