

Self-Directed Learning: A Cognitive and Computational Perspective

Todd M. Gureckis and Douglas B. Markant

New York University

Perspectives on Psychological Science
7(5) 464–481

© The Author(s) 2012

Reprints and permission:

sagepub.com/journalsPermissions.nav

DOI: 10.1177/1745691612454304

<http://pps.sagepub.com>



Abstract

A widely advocated idea in education is that people learn better when the flow of experience is under their control (i.e., learning is self-directed). However, the reasons why volitional control might result in superior acquisition and the limits to such advantages remain poorly understood. In this article, we review the issue from both a cognitive and computational perspective. On the cognitive side, self-directed learning allows individuals to focus effort on useful information they do not yet possess, can expose information that is inaccessible via passive observation, and may enhance the encoding and retention of materials. On the computational side, the development of efficient “active learning” algorithms that can select their own training data is an emerging research topic in machine learning. This review argues that recent advances in these related fields may offer a fresh theoretical perspective on how people gather information to support their own learning.

Keywords

self-directed learning, active learning, machine learning, self-regulated study, intervention-based causal learning

Some information is provided to us by the environment, and the timing and sequence of presentation is not under our immediate control (e.g., watching TV without a remote control or attending a lecture). Other information becomes accessible as a direct result of our own actions and choices, such as when we search for information online, interact with an unfamiliar device or mechanism, or ask questions of those around us. The goal of the present article is to consider the implications of these two modes of learning from both a cognitive and computational perspective.

The distinction between “active” and “passive” information acquisition is a perennial and hugely influential topic in the learning sciences (Bruner, Jolly, & Sylva, 1976; Montessori, 1912/1964; National Research Council, 1999; Piaget, 1930). For example, instruction methods such as discovery learning (Bruner, 1961), experiential learning (Kolb, 1984), and inquiry learning (Papert, 1980) all advocate situations in which learners engage in active hypothesis testing, interaction with learning materials, and self-directed exploration. In many of these accounts, self-directed acquisition of information is seen not only as a pedagogical tool but also as a “motivating force” on the desire to learn (Berlyne, 1960; Hebb, 1955). Similar ideas are prominent in theories of cognitive development (Adolph & Eppler, 1998; Gibson, 1988; Kuhn, Black, Keselman, & Kaplan, 2000; Montessori, 1912/1964; Piaget, 1930; Schulz & Bonawitz, 2007).

However, relative to the widespread enthusiasm for an “active” view of human learning, there has often been less

attention given to self-directed information acquisition in cognitive psychology and cognitive neuroscience. In fact, empirical studies of human learning and memory are most typically passive in that the experimenter tightly controls (and manipulates) what information is presented to the learner on every trial. As a result, basic research in learning and memory has sometimes failed to make contact with core issues in educational research. Meanwhile, the many uses of the term “active learning” or “discovery learning” throughout the learning sciences have led to increasingly divergent conceptions of the basic issues (Chi, 2009).

In this article, we provide a synthesis of research in cognitive science and machine learning considering the interplay of learning, decision making, and information gathering. In particular, we focus on situations in which learners are in control of the information they experience by way of their ongoing decisions (i.e., learning is self-directed). Our review demonstrates why self-directed learning is not simply a special case of passive learning but has important and varied implications for both what is learned from any experience and for what is learnable. In addition, we explore how self-directed learning in humans can be understood in terms of key computational principles borrowed from “active learning” research in the

Corresponding Author:

Todd Gureckis, New York University, 6 Washington Place, Room 280, New York, NY 10012

E-mail: todd.gureckis@nyu.edu

machine learning literature. This subfield of computer science seeks optimized learning algorithms that can construct or control their own training experience. Finally, we address the fundamental dilemma regarding self-directed learning that lies at the heart of recent debates in the educational literature (e.g., Klahr & Nigam, 2004; Mayer, 2004): When does self-directed learning improve learning, retention, or transfer, and when do learners fall prey to biases that limit their ability to effectively gather information? Although much of the research surveyed in this review focuses on basic cognitive processes involved in learning and memory, we highlight how this emerging cluster of ideas may apply to educationally relevant scenarios. We conclude that self-directed learning remains a relatively understudied issue in cognitive science (at least in comparison to the education literature) but one that holds fundamental implications for theories of learning and memory.

What Is Meant by “Self-Directed” Learning?

Although the idea that learning should be “active” or “self-directed” is a long-standing and influential idea, there is often a lack of agreement about exactly what this means (Chi, 2009). For example, self-directed learning is alternately associated with physical activity during a task (e.g., Harman, Humphrey, & Goodale, 1999), the generation effect (i.e., enhanced long-term memory for material that is actively retrieved; Crutcher & Healy, 1989; Jacoby, 1978), or with elaborative cognitive processes such as providing self-generated explanations (Lombrozo, 2006; Roscoe & Chi, 2007, 2008). The present article focuses on a single dimension of self-directed learning, namely, the consequence of allowing learners to make decisions about the information they want to experience (see Fig. 1). Our assertion is that interactions between information sampling behavior (i.e., the decision to access or gather some piece of information) and learning is one domain in which education, cognitive science, cognitive neuroscience, and

machine learning research have the greatest immediate potential for cross-fertilization.

However, distinguishing between these various senses of self-directed learning is difficult in most realistic learning situations. For example, in a passive learning environment wherein choices about information selection are limited, learners can still choose to selectively attend to different cues or features of the presented stimulus (e.g., Rehder & Hoffman, 2005). Even a teacher-led, “passive” student might be cognitively active in the sense of mentally evaluating hypotheses or explanations, just as a self-directed learner may engage in self-explanation in order to decide what information to access. Likewise, the degree of engagement of individual learners in a task (i.e., their level of “cognitive activity”) may be influenced by whether they are physically active during learning. Nevertheless, as our review will summarize, there are important psychological implications simply from allowing learners to make decisions about what information they want to access.

Experimental Approaches to Self-Directed Learning

As defined above, self-directed learning situations are relevant to a broad range of cognitive tasks. To highlight the basic distinction, the following section provides examples of popular learning tasks from cognitive psychology and compares them with nearly identical “self-directed” alternatives (see Table 1 for a summary). As we highlight later in the review, even relatively small changes to a learning task can have dramatic consequences for what is learned and retained.

Memory encoding

One key element of classroom learning is memorizing new facts. In many contemporary laboratory tasks used to study memory, the experimenter determines the sequence and timing




Fig. 1. An example of self-directed learning in everyday life. In the scene, a young child is flipping through the pages of a storybook. At some point, the child comes to a picture she finds interesting and requests the name of the object from the caregiver. A key feature of this example is that the learner herself, as opposed to the parent or teacher, controls the learning sequence through her choices and actions.

Table 1. Comparisons Between Traditional Experimental Tasks Used in Cognitive Psychology and Analogous Self-Directed Alternatives

Traditional cognitive task	Self-directed alternative	Additional measurements	Example papers
Memory encoding: Individual list items are presented one at a time by the experimenter.	“Flash card” study: The participant chooses the timing and ordering of each studied item.	Order, timing, and spacing of study decisions	Kornell and Metcalfe (2006); Metcalfe (2002); Voss, Gonsalves, Federmeier, Tranel, and Cohen (2011)
Category learning: Category exemplars are sampled randomly from an experimenter-defined distribution and presented one at a time.	Active category learning: The learner can point to particular items to query the category label or can “design” items to test.	Identity of queried examples, sequence of queries	Huttenlocher (1962); Markant and Gureckis (2010, 2012b)
Causal learning: Learners observe pairings between causal events and their consequences and use the observed contingencies to estimate causal strength.	Intervention-based causal learning: Learners actively design “experiments” or intervene on variables in a causal system then observe the consequences.	Pattern or sequence of interventions	Lagnado and Sloman (2004); Sobel and Kushnir (2006); Steyvers, Tenenbaum, Wagenmakers, and Blum (2003)
Decision making: Participant decides between a set of prospects that have been described by the experimenter.	Active gathering of information: Incomplete information is given about each prospect and the decision maker must gather additional information via active information search.	Sequence of information search decisions, how much information is gathered prior to making a decision	Edwards (1965); Tversky and Edwards (1966); Hau, Pleskac, Kiefer, and Hertwig (2008); Juni, Gureckis, and Maloney (2011)

of study items. In contrast, in a self-directed memory task, participants make decisions about how to study a set of items in preparation for a future test. For example, learners might control a “window” that reveals individual items hidden within an array so that they can devote varying amounts of study time to different items (Voss, Gonsalves, Federmeier, Tranel, & Cohen, 2011). Similarly, researchers have examined self-directed study time allocation in common educational scenarios such as studying with flashcards (Kornell & Bjork, 2007; Kornell & Metcalfe, 2006; Metcalfe, 2002, 2009; Metcalfe & Kornell, 2003; T. O. Nelson & Narens, 1994). In both of these situations, the learner’s choices (rather than the experimenter’s) determine what is learned, how much time is spent per item, and the sequence of information presented. Note that in comparison to more recent approaches, early research on memory actually gave participants more control (e.g., when memorizing a list of words printed on paper, participants could decide how to allocate effort or study time to different elements). However, on the whole, study strategies or exploration behaviors are less frequently a focus of investigation in memory research.

Category learning

Category learning research aims to understand how people discover the natural grouping of objects into classes (e.g., learning that a particular group of four-legged animals are “dogs”). Category learning differs from memory tasks in that

it examines not just the acquisition of new facts but the generalization of learned information to new situations (e.g., correctly classifying a completely novel dog). In the traditional “passive” version of a category learning task used in the laboratory, category members are sampled pseudo-randomly from an experimenter-defined statistical distribution and presented one at a time (Ashby & Gott, 1988; Shepard, Hovland, & Jenkins, 1961). Of particular interest here is the fact that the learner has no control over the order or nature of the stimuli. In a self-directed learning task, subjects might be able to design category members they would like to learn about (Markant & Gureckis, 2010, 2012b) or query individual category members by pointing at them in an array. Such procedures are analogous to a child asking a parent whether an unfamiliar object is a “dog” rather than waiting for the parent to name it (see a related literature on question-asking behavior: Berlyne & Frommer, 1966; Chouinard, 2007; Kemler Nelson, Egan, & Holt, 2004; Mills, Legare, Bills, & Mejias, 2010). Self-directed learning procedures were common in early work on category acquisition (Bruner, Goodnow, & Austin, 1956; Huttenlocher, 1962), but such paradigms have attracted less overall attention despite their relevance to classroom learning (see Kornell & Bjork, 2007, for a review). A related approach is found in eye-tracking studies that have explored how learners visually scan various stimulus properties during the course of learning (Blair, Watson, Calen Walshe, & Maj, 2009; Rehder & Hoffman, 2005), although eye movements include both voluntary and involuntary components.

Causal learning

Real world learning is not just about acquiring new facts or memories but often involves understanding the causal relations between events in the world. For example, a child might learn that medicine can get rid of a tummy ache, that heating up a liquid causes it to bubble, and that engaging in certain bad behaviors can cause an adult to get angry. In each of these cases, some event in the world (taking medicine, applying heat to a liquid, or behaving badly) causes certain other events. Philosophers, statisticians, and psychologists often distinguish between correlation and causation and have noted that it is impossible to empirically separate these ideas without the ability to actively intervene on the environment (e.g., Mill, 1843/1950). For example, although a child could observe that a sibling's behavior is correlated with the mood of their parent, children can better establish which behaviors cause their parent to get angry by themselves behaving in different ways in different situations and observing the resulting effects on their parent's disposition. This type of learning is often referred to as intervention-based causal learning because the learner is intervening or manipulating a variable in the environment (their behavior) just as a scientist might manipulate an independent variable in the design of an experiment.

A number of recent studies have explored the impact of active, intervention-based learning on causal learning in children and adults (Gopnik et al., 2004; Lagnado & Sloman, 2004; Rottman & Keil, 2012; Sobel & Kushnir, 2006; Steyvers, Tenenbaum, Wagenmakers, & Blum, 2003). In such tasks, participants learn by interacting with an ambiguous system, setting or manipulating variables and observing the effect these interventions have on other variables. For example, Sobel and Kushnir (2006) had adults interact with a virtual circuit board presented on a computer display. The circuit involved four differently colored lights that were linked in various ways such that turning on one light caused other lights to turn on (the causal relationships were probabilistic). Learners experimented with the system by setting various lights to the "on" or "off" state and observing the effect on other lights in the system. Later, they were tested for their knowledge of the causal relationships between the lights. Effective learning in this task requires the learner to design informative interventions or "experiments" that provide information about the causal links. Similar tasks have a long history in the developmental literature (e.g., Kuhn & Brannock, 1977; Kuhn & Ho, 1980; Piaget, 1930). At a fundamental level, these types of learning environments emphasize self-directed information gathering because the learner is in charge of which interventions to perform at each point in time. There is growing evidence that interventions are planned specifically with the objective of acquiring useful information about causal structure (e.g., Schulz & Bonawitz, 2007; Steyvers et al., 2003). For example, Schulz and Bonawitz (2007) found that young children play more with a toy after being shown confounded information about how it works, suggesting that one goal of exploratory play is to reduce uncertainty about causal structure.

Gathering information for making decisions

A final domain that frequently involves a component of self-directed information acquisition is decision making. In traditional decision-making tasks, individuals must choose between pairs of options that are described by the experimenter (e.g., "Would you prefer \$5 now or \$10 in two weeks?"). However, in order to make effective real-world decisions, people often have to first gather information about various options (e.g., Hertwig, Barron, Weber, & Erev, 2004). For example, when purchasing a car, a buyer might consult various resources concerning the reliability and pricing of different vehicles. In these cases, decision making is preceded by a period of self-directed information sampling (Edwards, 1965; Gureckis & Markant, 2009; Hau, Pleskac, Kiefer, & Hertwig, 2008; Hills & Hertwig, 2010; Juni, Gureckis, & Maloney, 2011; Markant & Gureckis, 2012a; Todd & Dieckmann, 2005; Todd, Hills, & Robbins, in press; Tversky & Edwards, 1966). The key questions in this research are what information people sample prior to making a decision, how much information is needed before making a choice, and how information sampling influences later choices. Experimental techniques such as the Mouselab protocol (Payne, Bettman, & Johnson, 1993; Payne, Braunstein, & Carroll, 1978) or eye tracking allow detailed measurement of self-directed information sampling prior to choice.

Isolating the Effects of Individual Choice and Control Through "Yoked" Designs

The previous examples clearly highlight how self-directed information sampling can alter the course of learning, but how can we tell whether it actually conveys advantages for acquiring new knowledge? Isolating the contribution of individual choice to learning and memory requires establishing appropriate experimental controls. One popular method is the use of "yoked" learning designs (see Huttenlocher, 1962, for an early example of this technique). In these experiments, one participant performs the task in a self-directed way, and the experimenter records the information that the learner requests or accesses. The same sequence of observations is then presented to another subject—the yoked control—who receives the same data passively (i.e., not under his or her control). Yoked learners experience a situation not unlike the main character in the movie *Being John Malkovich*, wherein they see the learning task through the "eyes" of another individual. By ensuring that there are no differences in the data experienced, differences in learning outcomes between these conditions help to highlight the effect of individual choice during learning (see Fig. 2).

However, a variety of factors suggest caution when interpreting yoked learning studies. First, it must be considered whether the yoking design systematically disadvantages yoked learners. For example, the lack of control that yoked participants feel may be distracting and could interfere with learning independently of their self-directed partner's behavior. Alternatively,




Fig. 2. A simple approach to yoked studies uses two computers that display identical information. As the self-directed learner interacts with the task on the host computer, a yoked participant seated in another room attempts to learn from the same display without the advantage of control.

whereas self-directed learners may be aware of the “next step” in the learning sequence as soon as they make a decision, yoked learners may lag behind when directing attention to new stimuli and preparing to learn. In addition, self-directed learners may be more engaged in a task relative to the yoked learner either because they are more physically active or because they are making more decisions. An important goal for yoked experiment designs is to control for engagement and attention in order to isolate the specific aspects of self-directed decision making that affect learning.

A more interesting factor to consider is whether, despite equating the sequence of data, the usefulness of the information experienced is tied to the beliefs of the self-directed learner. In particular, self-directed learners may gather data that specifically tests a hypothesis they have in mind, leading to a more individually informative training experience. Because different individuals bring different background knowledge to a task, this divergence could result in yoked learners gaining less from the same sequence of observations. As discussed below, a key challenge for yoked experiment designs is to disentangle “data-driven” processes from “decision-driven” processes that are involved in the act of gathering information.

A Cognitive Perspective on the Advantages of Self-Directed Learning

One reason cognitive scientists should be interested in self-directed learning is the fact that it is widely thought to improve learning, particularly in educational contexts. However, there is less understanding of why these benefits occur and considerable debate about the generality of such claims, particularly in the education literature (e.g., Klahr & Nigam, 2004; Mayer, 2004). We will begin our review by laying out a number of contributing cognitive factors. Our synthesis draws from a broad set of disparate psychological literatures that have explored the interactions between learning and choice behavior.

Data-driven or informational explanations

Rather than being limited by the flow of information from passive experience, self-directed learners are free to choose which information they want to learn. For example, by preferentially selecting information that reduces their current uncertainty, people may be able to optimize their experience (e.g., avoiding redundant data and focusing effort on parts of the environment that are not well understood). As a result, more can be learned with less training, because each experience is more useful or informative. For example, Markant and Gureckis (2010) found that when learning novel perceptual categories, participants who were free to query individual exemplars in a self-directed way outperformed participants who viewed examples that were randomly generated from a distribution (the typical setup in this kind of task). Analysis of self-directed participants’ information sampling decisions suggested that these individuals learned more quickly by avoiding redundant exemplars they were already able to confidently classify. Similar patterns of uncertainty-driven information gathering have been observed in the exploratory play of young children (Schulz & Bonawitz, 2007).

As noted above, some types of information are accessible only via interaction with the environment. For example, as in science, if two variables are correlated (e.g., lung cancer and smoking), actively manipulating one variable and observing the effect is necessary to establish the direction of causality, thus providing information that is inaccessible using observation alone (Mill, 1842/1950; Pearl, 2000). Consistent with this perspective, a number of recent human studies have found learning advantages for intervention-based learning in causal settings (Lagnado & Sloman, 2004; Sobel & Kushnir, 2006; Steyvers et al., 2003). In these cases, self-direction can be viewed as providing not just “better” data than passive learning but fundamentally distinct information that supports stronger inferences.

Self-directed information sampling can strongly shape our knowledge about the world. As a familiar example from social

psychology, if an individual interacts with others only following a positive social experience, this will result in more information about friends and biased information about others (because initial negative impressions are never corrected; Denrell, 2005; Fazio, Eiser, & Shook, 2004). Biased data gathering may underlie a variety of other cognitive biases, such as risk aversion (Denrell, 2007; Hertwig et al., 2004; March, 1996), overconfidence (Juslin, Winman, & Hansson, 2007), and illusory correlation perception (Fiedler, 2000). Indeed, psychologists are increasingly finding that interactions between choice behavior and learning have broad implications for many areas of knowledge representation, learning, and decision making (Fiedler & Juslin, 2006).

Effort or encoding optimization

Similar to the advantage of collecting data to test different hypotheses, self-directed learning allows people to decide how to study a set of material to maximize retention. Metcalfe and colleagues have examined how people allocate study effort across materials of varying difficulty in preparation for future memory tests (Metcalfe, 2002; Metcalfe & Kornell, 2003; Son & Kornell, 2008; Son & Metcalfe, 2000; see also Dunlosky & Hertzog, 1998; T. O. Nelson & Leonesio, 1988). One hypothesis is that materials or concepts just beyond the grasp of the learner are most amenable to learning (similar to Vygotsky's [1987] "zone of proximal development"). Consistent with this view, study time allocation research has found that learners often allocate the most effort to easier items and progress to harder items only later in a study session, particularly when given limited study time. For example, Metcalfe (2002) had people learn Spanish–English word pairs that were easy (e.g., "fantastico"–"fantastic"), medium, or hard (e.g., "chafarrinada"–"stain"). Rather than devote time to the most difficult pairs, participants focused on items of easy and medium difficulty that could be learned in the brief amount of study time available. Formal analyses suggest that such strategies are more effective for optimizing memory performance because starting with the most difficult items may mean fewer total items are successfully encoded (Atkinson, 1972a, 1972b; Son & Sethi, 2010).

Even when the set of material to be learned does not vary dramatically in difficulty, making decisions about the timing, spacing, and order of information experienced may enhance encoding. In a spatial memory task in which the goal was to memorize grids of commonplace objects, Voss et al. (2011) reported that individuals who controlled the time spent studying each item and the order visited (using a joystick) had superior memory at test compared with "yoked" individuals who viewed a replay of another subject's active scan path. In this study, self-directed encoding was also associated with increased coordination in cortico-hippocampal activity compared with the yoked control (assessed using fMRI), suggesting that self-directed study enhances neural processes related to successful encoding. Some of the observed benefits in this

task may also result from the link between the data experienced by learners and their current knowledge state (Atkinson, 1972b; Kornell & Metcalfe, 2006). For example, if a participant has prior experience with a particular object in the array, this may facilitate learning on the initial exposure and lead to avoiding that item in subsequent study opportunities. The same object may be less familiar to a yoked participant and thus would benefit more from repeated study.

Although these studies demonstrate potential benefits of self-directed study, other work suggests that people often fail to account for properties of their own memory when selecting study strategies. For example, memory is typically better for items that are spaced (i.e., repeated study events are distributed in time) rather than massed (i.e., study events occur repeatedly in close succession; Dempster, 1989; Glenberg, 1979). However, people often fail to take advantage of this effect when deciding how to study. For example, people believe that massed practice is a more effective learning mode (Kornell & Bjork, 2008; Simon & Bjork, 2001) and when given the opportunity prefer to mass practice difficult items and space easy items (Son, 2004, although see Son & Kornell, 2009, for a counterexample). Taken together, however, work in this area strongly suggests that even memory encoding—a mental function typically viewed as passive or automatic—might be better understood as involving some component of active, strategic choice.

Inductive inference and sampling assumptions

In addition to modifying the flow of experience, knowledge about the way information was gathered may influence learning and generalization. For example, Xu and Tenenbaum (2007) explored how children generalize words to novel objects (see also Gweon, Tenenbaum, & Schultz, 2010; Rhodes, Gelman, & Brickman, 2010). In one condition, a set of four novel objects was verbally labeled by a knowledgeable teacher (i.e., the experimenter), whereas in the other condition the same four objects were labeled after the children themselves pointed to them and requested their label (i.e., a self-directed condition; see Figure 3 for an illustrative example). In this study, children made more restrictive generalizations of word labels (i.e., only to other nearly identical items) when examples were selected and labeled by a teacher, compared with when they selected examples themselves, even though the final information conveyed in both conditions was identical.

One explanation of this finding is that children inform their generalizations with knowledge of how data was generated or gathered. Such sampling assumptions have been a topic of extensive discussion both in machine learning (Mitchell, 1997) and psychology (Griffiths & Tenenbaum, 2001; Shepard, 1987; Tenenbaum, 1999). Under "strong" sampling, training examples are selected from the space of items that fall within the target concept. This assumption may be appropriate when the process by which examples were collected is selective, such as being picked by a helpful and knowledgeable teacher (Xu &




Fig. 3. An illustrative example that conceptually matches the task used by Xu and Tenenbaum (2007). At left is a bucket of different colored balls. In one scenario, a knowledgeable teacher picks out four red balls and labels each a “Dax” (middle). In the second scenario (at right), the learner herself picks up four balls from the bucket that just happen to be red and learns via feedback that all four objects are called “Dax.” The learner may then be asked whether the label “Dax” should also apply to other balls in the bucket that are not red. Different inferences in these two scenarios might stem from the learner’s assumptions about the process that gathered the examples.

Tenenbaum, 2007). This is also consistent with ideas of pedagogical or intentional sampling (e.g., Gweon et al., 2010; Shafto & Goodman, 2008; Shafto, Goodman, & Frank, in press). The assumption of strong (or pedagogical) sampling allows more restrictive generalization from smaller sets of examples. However, the strong sampling assumption may be inappropriate when information is gathered independently of the target concept (such as when samples are gathered randomly or by a self-directed learner who is ignorant of the true pattern). In these cases, self-directed learners might at first default to a “weak” sampling assumption. Weak sampling implies less restrictive generalization because the sampling process itself conveys no information about the to-be-learned concept. These ideas have been formalized within a Bayesian learning framework that successfully explains human inference and generalization across a variety of situations (Griffiths & Tenenbaum, 2001; Tenenbaum, 1999). Critically, this line of work shows that even in cases where two people experience identical information, learning may depend on the role of each individual in gathering it.

Decision-driven explanations

Independent of different data or sampling assumptions, the very act of planning interventions or deciding which information to collect may necessitate a more thorough evaluation of the problem structure and of how observed experience relates to different hypotheses (Bruner, 1961). Sobel and Kushnir (2006) showed that learners who designed their own interventions on a causal system learned better than yoked participants who either passively observed the same sequence of actions or re-created the same choices made by others (see also Markant & Gureckis, 2010; Steyvers et al., 2003). Because the data experienced by both groups is identical, “yoked” learning studies such as these isolate the influence of individual choice from differences in information.

As mentioned earlier, enhanced memory under volitional memory encoding has been associated with greater coordination across a network of brain regions involved in executive control, attention, and memory encoding (Voss et al., 2011). However, further work is needed to determine how different components of self-directed decision making contribute to this effect. For example, deciding what to study often relies on a metacognitive judgment about what has already been learned, an introspective process that can enhance memory independently of any further study (Kimball & Metcalfe, 2003). Even in the absence of such metacognitive monitoring, however, learning may be enhanced by basic processes related to making decisions. Explanations for such decision-driven advantages include the psychological benefits of free choice (Leotti, Iyengar, & Oshsner, 2010), greater engagement with the learning task, and memory enhancements related to active exploration (Kaplan et al., 2012).

It is important to note that decision-driven effects may also interact with the data-driven differences described above. For example, as argued in Markant and Gureckis (2010), self-directed learners can select information that would test the set of hypotheses they are currently entertaining whereas yoked participants (who may be considering different hypotheses about a task) may gain less from the same sequence of data. Markant and Gureckis described a simple computational model that accounts for self-directed learning advantages in terms of the congruence between experienced data and the knowledge held by the individual.

In summary, self-directed sampling does not appear to simply be a special case of passive learning in which an action is first required by the learner but may have broad implications both for what is learned and what is learnable. In particular, self-directed information sampling appears to improve the rate of learning and the fidelity of memory, can reveal novel or useful information about the structure of the environment, can

influence patterns of inference and generalization, and can alter subsequent decision-making strategies.

Curious Machines: A Computational Perspective on Self-Directed Learning

Paralleling this psychological literature are recent developments in machine learning under the name “active learning.” Unlike traditional learning models that involve passively fed training data, this work has explored algorithms that gather their own training data (see Settles, 2009 or Sutton & Barto, 1998, for reviews). For example, consider the document classification problem confronting many Internet search engine companies in which data (e.g., Web pages, videos) are nearly infinite, but obtaining information about the content of each document may involve costly human operators. Such problems are analogous to the data-rich, instruction-sparse environments that confront a young child. In these cases, it would be ideal if the classifier system could make intelligent decisions about which information is expected to be nonredundant and request additional information for only those items. This motivation has been used to develop “curious” machine learners that can perform as well as passive approaches but with less training. Such techniques have been recently applied to a wide range of learning problems, including sequential decision making (i.e., reinforcement learning), causal learning, and categorization.

The formal methods developed in this line of work have the potential to make important contributions to research on human learning. First, research in this area has attempted to formally delineate the tasks and environments in which self-directed sampling may result in improvements in learning efficiency (Angluin, 1988; Cohn, Atlas, & Ladner, 1994; Dasgupta, Kalai, & Monteleoni, 2005). Typically these analyses take the form of mathematical proofs establishing the best- and worst-case advantages for various information-gathering strategies and learning models. To the degree that the abstract properties of the tasks and models studied in this literature can be mapped on to educationally relevant learning tasks, such analyses may offer insight into which environments are best suited for self-directed learning. Second, this work has emphasized how self-directed information sampling might be combined with other types of learning (such as unsupervised learning; Dasgupta, 2010). Finally, research in this area has designed “sampling norms” or choice utility functions that assign value to future observations on the basis of their potential for revealing information (MacKay, 1992; Seung, Opper, & Sompolinsky, 1992). Such proposals provide a framework for studying how humans evaluate different sources of information and make sampling decisions.

It is also possible to use the pitfalls from the machine learning literature to better understand the limitations of self-directed learning as a pedagogical strategy. For example, a well-recognized weakness of many active learning algorithms is that if the learning model is incorrectly specified for the

domain (i.e., the space of possible hypotheses or representations within the model does not encompass the to-be-learned concept or assigns it a very low a priori probability), the information samples acquired will be severely biased (MacKay, 1992). This may result in a nefarious feedback loop in which an incorrect early impression of a problem leads to biased information gathering, which in turn reinforces or fails to correct early impressions, similar to confirmation bias in human psychology (Nickerson, 1998; Wason, 1960). In this instance, something that might at first seem like an “irrational” bias on the part of a learner may instead be thought of as a misspecified model of the learning task or environment.

In addition, the focus on minimizing costs or uncertainty in artificial systems may fail to capture the information that is most useful for human learners. For example, active machine classifiers often preferentially explore the boundaries of a concept (i.e., the “margin”), but the same borderline cases may be less useful to a human learner because they are the ones most likely to be associated with inconsistent feedback (Ashby, Boynton, & Lee, 1994). In one example, Lang and Baum (1992) developed a handwritten digit recognition system that could synthesize novel characters and could ask a human oracle for feedback. The system quickly began testing borderline examples that were difficult for the human assistant to classify (e.g., those that resemble both a “3” and an “8”), leading to inconsistent trial-to-trial feedback that confused the system. Although borderline or ambiguous items may be highly valuable in terms of information, they can often be nonrepresentative of the overall concept. However, human learners often benefit from training on mixtures of extreme, typical, and borderline examples (Avrahami et al., 1997; Clapper & Bower, 2002; Elio & Anderson, 1984; see also Dasgupta, 2010, for a discussion in the context of machine learning). One interesting question is whether human learners themselves prefer to gather information that is representative or primarily diagnostic of a target concept (e.g., G. L. Murphy & Ross, 2005). These issues become particularly important when evaluating whether these machine learning algorithms might be used as assistive training tools for human learners.

Overall, this area of machine learning research shares many of the same goals as research on self-directed learning in humans and provides a useful set of formal tools for analyzing such behaviors. In the following section, we highlight three examples in which analyses of learning tasks with machines have shown advantages for self-direction, and we attempt to distill relevant psychological insights.

Example 1: Active sampling for generalization

As noted earlier, in data-rich, feedback-sparse learning environments, it is beneficial to maximize the informativeness of each training experience. Machine learning researchers have repeatedly shown that learning algorithms that can strategically select their own training data can reduce the number of training exposures needed to reach a particular level of error

compared with equivalent models without this capability (e.g., Angluin, 1988; Castro et al., 2008; Cohn et al., 1994; Dasgupta, 2010; Dasgupta et al., 2005; Lang & Baum, 1992; MacKay, 1992; K. P. Murphy, 2001; Seung et al., 1992; Tong & Koller, 2001). One way to quantify the advantages of self-directed learning is depicted in the so-called banana curve, which plots expected accuracy as a function of the number of training trials (analogous to a standard human learning curve; see Fig. 4, left panel). The difference between the curves shows the advantage given by active, self-directed learning over the passive training sequence. The characteristic pattern is that self-directed learning can lead to similar levels of performance with less training even within the same learning architecture. For example, in Figure 4 (left panel), the self-directed model needs only one third as much training to achieve the same level of performance as the passive model.

The following example provides some intuition for why this works (Dasgupta, 2010; Settles, 2009). Consider a simple unidimensional classification task in which the goal is to accurately estimate some unknown threshold, θ , below which items are in Category A and above which items are in Category B (e.g., learning the building height at which a structure becomes a “skyscraper”). If the acceptable error in our estimate of the true threshold is ϵ , formal learning theory shows it is enough to “passively” sample $O(1/\epsilon)$ random (uniformly distributed) training items with labels (Valiant, 1984).¹ In this case, more data help better identify the threshold, but many of those observations will be uninformative (the threshold always will lie between the tallest negative example of a skyscraper and the shortest building called a skyscraper). Instead, if we can query particular items (e.g., ask questions about buildings

at a specific height), the threshold estimation problem becomes equivalent to a binary search (dividing the remaining search interval in half with each query) and error ϵ can be achieved with $O(\log(1/\epsilon))$ queries, an exponential reduction in the number of labeled examples. In other words, an idealized learner would be much faster by asking targeted questions than by randomly sampling examples to query. Recent studies of self-directed category learning show similar advantages for human learners (Castro et al., 2008; Markant & Gureckis, 2010). What is important is that these analyses highlight how advantages for self-directed learning might arise simply as a consequence of effective information gathering, independent of deeper encoding or attentional effects.

The extension of this basic idea to more complex learning tasks has been established in numerous domains (Settles, 2009). Of course, not all problems reduce to simple threshold estimation. Typically, in more complex learning problems, decisions about the potential value of new observations are based on the current “beliefs” of the learning agent. For example, the learning system might request the label for items that it is more uncertain about how to classify (a bit like asking for help). This uncertainty can be quantified or computed in various ways. For example, one way to make information sampling decisions is to maintain a “committee” of different hypotheses or mental models of the situation and to use the disagreement among them as a measure of potential for information (cf. Seung et al., 1992). The intuition is that disagreement between the various hypotheses are places in which targeted information would be most helpful. This idea has a long history in the philosophy of science wherein “crucial experiments” that decide between alternative theories should

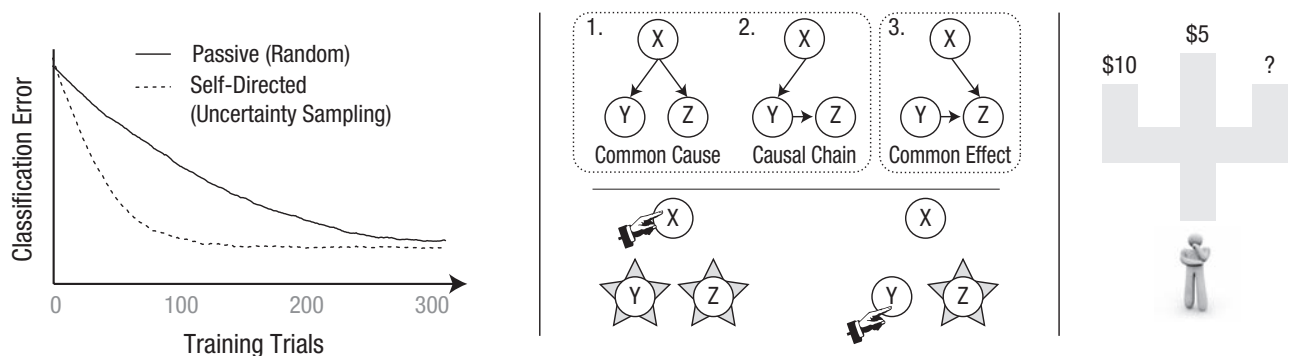


Fig. 4. Left: Example of a “banana curve” typical of computational active learning applications. The curves compare expected error as a function of training episodes for both self-directed and passive learners (Settles, 2009). Uncertainty-driven self-directed sampling can often achieve the same classification performance with an order of magnitude of less training. For example, the passive learning algorithm reaches asymptote at Trial 300, whereas the self-directed learning algorithm reaches a similar point at Trial 100 (one third as much training). Middle: The top row shows three possible causal relationships between a set of three variables (X, Y, Z). The arrows indicate the direction of causality (e.g., in Structure 1, X causes both Y and Z). The structures are grouped with a dotted line reflecting Markov equivalence classes. Structures grouped in a Markov equivalence class are indistinguishable by passive observation alone (e.g., ignoring time, the pattern of resulting effects for 1 and 2 are identical). However, by actively intervening on the variables, the learner can disambiguate the causal structure. The bottom row shows a set of possible interventions on the structure (the hand represents a variable whose value is set by the learner; the variables with stars are “turned on” following the intervention). In the first example, the learner manipulates X and finds that both Y and Z turn on. As a result, the learner might rule out Structure 3 (for simplicity, assume deterministic causes and no background causes). In contrast, in the second example, the learner manipulates the value of Y and observes that only Z turns on. This effectively rules out Structure 1. Right: The explore–exploit dilemma. Two options are available with known values, but a third has not been sampled. The learner must decide whether to “exploit” the known source of reward (the path leading to the \$10 option) or to “explore” the unknown and potentially more valuable option leading to the question mark.

allow more effective discovery (Bacon, 1620/1902; Platt, 1964; Popper, 1935/1959). Such uncertainty-driven information sampling also has a natural framing in information-theoretic terms using concepts like entropy or Kullback-Leibler divergence (MacKay, 1992), ideas that closely parallel early work on hypothesis testing and information selection in cognitive psychology (Oaksford & Chater, 1994). One major difference is the scale at which these approaches may apply. Early work on hypothesis testing in psychology considered tasks where there were only a few possible hypotheses, whereas recent work in machine learning applies to real-world problems with possibly millions of alternative hypotheses.

In summary, recent machine learning research provides quantitative support for the idea that active information selection can improve the rate of learning and generalization. In addition, research in this area has explored formal ways a learning agent might assign value to future information samples on the basis of the agent's current uncertainty. Each of these advances go beyond early work on hypothesis testing in cognitive psychology while potentially providing novel insights into self-directed learning in humans.

Example 2: Intervention-based causal learning

Self-directed learning can also be critical in the discovery of causal relations. As described above, causal relationships establish which variables or events in the world are the consequence of other variables or events (e.g., do cell phones cause cancer?). The top row of Figure 4 (middle panel) shows three possible causal relationships between a set of simple variables. To make the discussion concrete, Variable X might represent "cell phone usage," Variable Y might be "industrial chemicals production," and Variable Z might represent "cancer rates." Structure 1 thus represents a scenario in which cell phones cause both industrial chemicals and cancer (a common cause structure). Structure 2 represents the causal relation in which cell phones cause the production of industrial chemicals, which in turn cause cancer (a causal chain). Ignoring temporal ordering, these two structures (1 and 2) are indistinguishable through observation alone because both produce a pattern of intercorrelation between the variables X, Y, and Z (i.e., increases in cell phones usage will be associated with increases in chemicals and cancer). Formally speaking, these two structures represent a "Markov equivalence class" (Pearl, 2000). However, a learner who actively intervenes on the system by fixing the value of different variables can begin to disambiguate the causal structure (see bottom row of Fig. 4, middle panel). For example, after "manipulating" or "experimenting" with Variable Y and observing no effect on Z (independent of the possible base-rate occurrence of Z), the learner can rule out the causal chain structure. Active, intervention-driven causal learning can be viewed as a form of uncertainty reduction within the space of possible causal models that might relate the variables. In Bayesian terms, each possible causal structure in the top row of Figure 4 (middle panel) might represent

a hypothesis, and the goal is to reduce the uncertainty about the true structure by making maximally informative interventions. In machine learning, computational systems have been developed to predict which intervention or experiment to perform in order to best learn about the causal system (K. P. Murphy, 2001; Tong & Koller, 2001). Such approaches may greatly inform our understanding of how people interact with causal systems in order to learn. For example, Steyvers et al. (2003) compared a similar model with the behavior of human participants in an intervention-based causal learning task and found that human intervention decisions were well predicted by the goal of reducing entropy (i.e., uncertainty) over the space of possible causal models.

Example 3: Exploration–exploitation trade-offs and learning from reinforcement

Active information acquisition also figures prominently in theories of computational reinforcement learning (CRL; Sutton & Barto, 1998). In CRL, learning is viewed as an attempt to maximize the reward that an agent can receive from its environment.² In pursuit of this goal, learners must exploit resources that are known to be productive (Fig. 4, right panel). However, given that the full distribution of rewards in the world is usually unknown, agents must balance the desire to exploit options that have been productive in the past with the need to explore (i.e., gather information about) relatively unknown outcomes. For example, imagine trying to decide the best restaurant at which to eat in a large city. Often we face a dilemma in that we could continue eating at our local favorite, but are always unsure whether a better restaurant exists nearby. Thus, our challenge is to balance the desire to "exploit" options that are known to be good against possible gains from "exploring" (i.e., sampling) new alternatives.

Outside of a few domains (Gittins & Jones, 1979), the optimal solution to this exploration–exploitation trade-off is computationally intractable. As a result, many CRL algorithms developed in machine learning assume that learners engage in stochastic, random exploration (e.g., the "Softmax" rule proposed by Sutton & Barto, 1998), and this default assumption often carries over into human research that uses CRL as a theoretical model of human learning (Daw, O'Doherty, Seymour, Dayan, & Dolan, 2006).

However, other CRL research has considered alternative exploration strategies, such as "novelty bonuses," which reward selections of actions that have not been sampled recently (Kakade & Dayan, 2002) or which use cumulative absolute prediction error as a signal for guiding exploration (Schmidhuber, 1991). The intuition is that exploration should be devoted to areas of the environment that are unfamiliar or have not been sampled recently (especially when the environment is continually changing). Consistent with these ideas, there is evidence that novelty may serve as a neurobiological signal that engages exploratory behaviors in human learners (Wittmann, Daw, Seymour, & Dolan, 2008). Alternatively,

prediction error (i.e., the inability to correctly predict outcomes) can be a signal about the need for additional learning about particular parts of the environment. For example, if a person frequently has unpredictable eating experiences at Italian restaurants, it could be a cue that more learning is needed about this category of restaurant. An agent that can monitor its own prediction failures (a type of “metacognition” or performance monitoring) can use this information to guide exploration. Alternatively, at some point, the learner might decide that the variance in Italian restaurants is irreducible (i.e., it is due to intrinsic variability of the restaurants rather than a lack of understanding; see Yu & Dayan, 2003, for a discussion of expected and unexpected uncertainty). Finally, exploration may also be allocated to reduce uncertainty about the possible states of the environment as in the active learning and causal learning research discussed above (Kaelbling, Littman, & Cassandra, 1998).

Open Questions and Future Challenges

We conclude our review by discussing some of the open questions and future challenges confronting the study of self-directed learning.

How do people make information collection decisions?

Relative to what is known about how people make decisions in economic contexts, the question of how people make decisions in order to gain information is less well understood. As outlined above, there are a variety of ways that human learners might evaluate the potential information content of future observations. At one extreme, human information sampling behavior could be effectively random or loosely guided by past experience (i.e., stochastic exploration, as in Sutton & Barto, 1998). Conversely, information sampling decisions could reflect a belief-driven process that gathers specific information in an attempt to reduce uncertainty (Knox, Otto, Stone, & Love, 2012; Kruschke, 2008; J. D. Nelson, 2005; J. D. Nelson, McKenzie, Cottrell, & Sejnowski, 2010; Steyvers et al., 2003) or costs (Gureckis & Markant, 2009; Juni et al., 2011).

A key focus of ongoing research is to better distinguish these alternatives. For example, recent work by Nelson and colleagues has looked for an information sampling “norm” that best describes human information search in a variety of tasks (J. D. Nelson, 2005; J. D. Nelson et al., 2010). Many of these information sampling norms parallel those developed in the machine learning and statistics research on active information sampling. Markant and Gureckis (2012b) examined self-directed learning in a relatively complex rule learning task that gave participants the ability to “design and test” stimuli they wanted to learn about. On a subset of trials, participants were asked to report their uncertainty about how to classify the item they had just designed. Using a computer model-based analysis, we found that people tended to prefer testing items that

discriminated between two hypotheses or categories at a time rather than information that reduced the “global” uncertainty across all categories (somewhat related to Tweney et al.’s [1980] observation that people are effective at testing between two alternative hypotheses). The models tested in Markant and Gureckis (2012b) were motivated by various information sampling norms originally proposed in the machine learning literature (Settles, 2009). Relatedly, Kincannon and Spellman (2003) showed that people prefer gathering different kinds of evidence when generalizing a hypothesis to all members of a category compared with limiting a hypothesis to only members of a category, suggesting that framing effects may also influence information sampling. Of course, a further challenge is that people might use multiple strategies or sampling norms at various stages of learning. Developing a complete account of human information sampling behavior will thus require a fuller understanding of how basic drives (like curiosity or novelty preferences; e.g., Berlyne, 1960) interact with belief-driven, uncertainty-reducing sampling processes.

Contemporary machine learning research and computational modeling may prove critical in answering these questions. For example, one lesson from the machine learning literature is that understanding what information an agent will seek at any point in time requires a moment-to-moment model of the agent’s current state of knowledge or belief. Along these lines, we have attempted to fit learning and decision models to the trial-by-trial information gathering decisions of human participants (Gureckis & Markant, 2009; Markant & Gureckis, 2012a). By tracking the learning sequence for each participant (along with a plausible model of how people update their beliefs), these models may allow for a stronger test of various hypotheses about how people evaluate the potential information to be gained from future observations.

What environments are most likely to show advantages for self-directed learning?

Interfaces between basic cognitive research and machine learning may provide insight into which tasks or environments are best suited to self-directed learning (an issue with implications for education policy). As one example, Markant and Gureckis (2010) examined self-directed learning across a variety of different category structures. They found that participants’ prior biases about the target structure strongly guided their search behavior and that self-directed learning was most effective when a participant’s prior hypothesis space was congruent with that of the target concept. These results echo research in machine learning that suggests that active information sampling is less effective when the learner’s hypothesis space does not encompass the target concept or assigns it a low probability of being correct (MacKay, 1992). These results highlight a potential paradox. If people (or machines) use their current beliefs or understanding to drive their information acquisition, and those beliefs are incorrect, sampling will tend to be biased and learning will suffer. In the most extreme case,

if the learner has no way to represent a particular concept, then there is no way that self-directed learning will lead them to “discover” it, and their information gathering decisions will appear consistently suboptimal.

Predicting whether a particular problem is amenable to self-directed learning thus depends critically on an understanding of the learner’s representation of the task. In particular, self-direction is likely to speed learning in situations in which there is a large space of possible things to learn but the learner has a proper understanding or representation of this space. A failure of self-directed learning may reflect errors in the learner’s conception of the problem domain rather than poor information gathering. The critical question is thus whether self-directed learning allows people to overcome default strategies and prior biases when required by the problem they face. Such analyses may help to predict when self-directed learning will be more effective when used in concert with other instructional methods. For example, situations in which a teacher provides helpful or informative examples can correct early misconceptions, help elaborate the space of concepts or stimulus properties that are relevant, and then allow more effective self-directed learning within that space. Such ideas have interesting parallels with the tension between “direct instruction” and “discovery learning” in the education literature (e.g., Klahr & Nigam, 2004).

Coming from a memory optimization perspective, Son and Sethi (2010) presented a mathematical analysis of learning environments in which adaptive allocation of effort is most likely to lead to more effective learning. They pointed out that the effectiveness of various self-directed study-time allocation strategies depends critically on the “uptake” function that relates time spent studying to performance. Figure 5 shows two example functions. The first (denoted by the solid line) shows an item for which time spent studying monotonically increases memory strength (a concave function). The second (denoted by the dashed line) shows an item for which early effort results in little performance changes, but after a certain threshold of effort, large gains are realized (an S-shaped

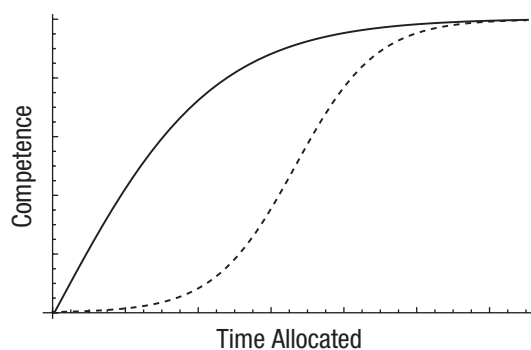


Fig. 5. Two learning tasks that vary in the function relating time spent to competence (modeled after Son & Sethi, 2010). In the solid curve, early time investments lead to immediate gains in competence. In the second environment (dotted line), there is an S-shaped curve such that early time investment leads to small gains, but larger gains are experienced later.

function). These differences might be related to the nature of the material being studied (some information one struggles with for some time before it starts making sense). One study strategy a learner might adopt is to allocate effort to items that seem to be returning the greatest improvement in competence per unit time spent. For environments in which the uptake function is nonconcave (e.g., S-shaped), this particular strategy will tend to neglect poorly learned items (particularly those for which an additional “investment” of study time is needed to increase the rate of learning, as in the dotted line in Fig. 5). As a result, self-directed study time allocation will tend to be suboptimal in environments that contain items with this property. Research on this question is undoubtedly in an early stage, but the potential payoff is a better understanding of how learning strategy (self-directed or passive) interacts with the informational structure of learning problems.

A unified view of self-directed learning

Despite the many advantages described in this review, self-directed learning can also lead to many learning disadvantages (a point gathering renewed attention in education; Mayer, 2004). For example, a now famous quote from Atkinson (1972a) reads:

One way to avoid the challenge and responsibility of developing a theory of instruction is to adopt the view that the learner is the best judge of what to study, when to study, and how to study. I am alarmed by the number of individuals who advocate this position despite a great deal of negative evidence. My data, and the data of others, indicate that the learner is not a particularly effective decision maker. (p. 930)

This observation is echoed in the extensive work showing that people are often biased in how they select new information when learning (e.g., failing to differentiate between information that would confirm or disconfirm their current beliefs; Klayman & Ha, 1987; Skov & Sherman, 1986; Wason, 1960). Biased information collection can skew understanding about true patterns in the environment and may underlie a variety of cognitive biases (Denrell, 2007; Fiedler, 2000; Fiedler & Juslin, 2006; Hertwig et al., 2004; Juslin et al., 2007; March, 1996). In addition, people frequently exhibit metacognitive illusions (e.g., Koriat & Bjork, 2005, 2006) that can lead to deficiencies in self-directed learning. For example, students often confuse the sense of perceptual fluency gained from massed practice with true long-term learning, leading to suboptimal preferences for massed practice over spaced practice (Son & Kornell, 2009).

How do we reconcile all these potential disadvantages for self-directed learning with the results reviewed in this article suggesting the opposite? We argue that research on this topic may benefit from a more unified theoretical perspective offered by the parallel work in machine learning. Leveraging a computational-level understanding of these basic issues has a

number of advantages. To date, the field has yet to arrive at a comprehensive understanding of when self-directed learning will or will not succeed. For example, although in some situations people show a robust tendency to seek confirmation, other work has found evidence that people seek disconfirmatory evidence (e.g., Gorman, Stafford, & Gorman, 1987; Markant & Gureckis, 2010; Nelson et al., 2010). Up to this point, it has been difficult to integrate these disparate results independent of the specific task or context in which behavior has been studied. In contrast, computational models can potentially expose general principles of learning that apply across many situations. In addition, although the tasks considered in this review are relatively simple, the applications of the machine learning approaches have extended to complex, real-world problems. This suggests that aspects of the theoretical approach may scale gracefully beyond laboratory tasks.

From self-directed learning to assistive training

Given that people are not always optimal self-directed learners, one promising avenue for future research is to use insight gained from the study of active information sampling (in both human and machines) to develop assistive training methods. Instead of predicting what information people will choose on their own to solve a task, cognitive models can be used to determine what information would be most helpful to the individual (given the nature of the task and measures of prior learning). This inversion of formal models into an assistive learning device may help shortcut the time required to develop perceptual or conceptual expertise in a domain and may be used to tailor learning experiences to the strengths and weaknesses of the individual learners (an approach that has been used with some success in the memory literature; Atkinson, 1972a, 1972b; Pavlik & Anderson, 2008). For example, Atkinson (1972b) compared recall performance for a set of 84 German–English vocabulary translations following three different study techniques: a self-directed condition in which participants made their own choices about which pair to study on each trial, random selection of study pairs, and a model-based optimization technique that attempted to predict at each point in time which study item would maximize the number of “permanently learned” items. The finding was that the model-based adaptive training technique led to better subsequent memory than the other two strategies (79% recall for the model training sequence versus 58% recall for self-directed decisions). Similar gains have been demonstrated by Pavlik and Anderson (2008) using the ACT-R model of memory. Relatedly, Castro et al. (2008) reported faster learning of a simple binary classification problem with a procedure in which a computational model selected training examples for the learner. Similar model-based approaches have been extended to predict human information needs in tasks as rich as a complex video game (e.g., Love, Jones, Tomlinson, & Howe, 2008).

In conclusion, although the concept of self-directed learning has had a widespread influence on education research, this idea has received less attention in basic studies of learning and memory. As our review highlights, the study of self-directed learning opens new, relatively underexplored avenues for psychological research. However, progress on these issues will require experimenters to relinquish the control they are accustomed to exerting over the learning process and let individuals freely explore and sample information in their environment. In combination with recent advances in machine learning, it is increasingly possible to make sense of such highly individualized learning sequences. In addition, machine learning research provides new quantitative tools for analyzing the effectiveness of self-directed learning and how it might vary across learning environments. A more complete understanding of the psychological processes underlying self-directed information sampling behavior may help bridge the gap between basic cognitive research and education research.

Acknowledgments

The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Intelligence Advanced Research Projects Activity, the Department of the Interior, or the U.S. Government. We thank Sean Kang, Lisa Son, and Barbara Spellman for helpful comments. In addition, we'd like to thank Stephanie Chen, Kari Kretch, Burr Settles, and Dylan Simon for reading an earlier version of this article.

Declaration of Conflicting Interests

The authors declared that they had no conflicts of interest with respect to their authorship or the publication of this article.

Funding

This work was supported by the Intelligence Advanced Research Projects Activity via Department of the Interior Contract D10PC20023.

Notes

1. The $O()$ is a common notation in computer science that reflects how the running time of a program changes as the size of its input increases. In this example, with passive learning, the training time of the simple threshold estimation algorithm would scale with $1/\epsilon$ (the length of processing would increase for smaller values of ϵ).
2. A key difference between active learning research and CRL centers on the nature of what is being optimized. In CRL, the rewards in the environment are unknown, and learning is about discovering both where the reward is in the environment and a decision policy that can optimize that reward. In contrast, active learning involves optimizing a known utility function (typically model uncertainty). However, both critically involve active information gathering in the pursuit of learning.

References

- Adolph, K., & Eppler, M. (1998). Development of visually guided locomotion. *Ecological Psychology, 10*, 303–321. doi:10.1080/10407413.1998.9652687
- Angluin, D. (1988). Queries and concept learning. *Machine Learning, 2*, 319–342. doi:10.1007/BF00116828
- Ashby, F. G., Boynton, G., & Lee, W. W. (1994). Categorization response time with multidimensional stimuli. *Perception & Psychophysics, 55*(1), 11–27. doi:10.1037/0278-7393.14.1.33
- Ashby, F. G., & Gott, R. E. (1988). Decision rules in the perception and categorization of multidimensional stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 14*(1), 33–53. doi:10.1037/0278-7393.14.1.33
- Atkinson, R. (1972a). Ingredients for a theory of instruction. *American Psychologist, 27*, 921–931. doi:10.1037/h0033572
- Atkinson, R. (1972b). Optimizing the learning of a second-language vocabulary. *Journal of Experimental Psychology, 96*, 124–129. doi:10.1037/h0033475
- Avrahami, J., Kareev, Y., Bogot, Y., Caspi, R., Dunaevsky, S., & Lerner, S. (1997). Teaching by examples: Implication for the process of category acquisition. *The Quarterly Journal of Experimental Psychology, Section A, 50*(3), 586–606. doi:10.1080/713755719
- Bacon, F. (1902). *Novum organum*. (J. Dewey, Ed.). New York, NY: P. F. Collier. (Original work published 1620)
- Berlyne, D. E. (1960). *Conflict, arousal, and curiosity*. New York, NY: McGraw-Hill.
- Berlyne, D. E., & Frommer Frances, D. (1966). Some determinants of the incidence and content of children's questions. *Child Development, 37*(1), 177–189. doi:10.2307/1126438
- Blair, M. R., Watson, M. R., Calen Walshe, R., & Maj, F. (2009). Extremely selective attention: Eye-tracking studies of the dynamic allocation of attention to stimulus features in categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 35*(5), 1196–1206. doi:10.1037/a0016272
- Bruner, J. S. (1961). The act of discovery. *Harvard Educational Review, 31*(1), 21–32.
- Bruner, J. S., Goodnow, J. J., & Austin, G. A. (1956). *A study of thinking*. New York, NY: Wiley.
- Bruner, J. S., Jolly, A., & Sylva, K. (1976). *Play: Its role in development and evolution*. New York, NY: Basic Books.
- Castro, R., Kalish, C., Nowak, R., Qian, R., Rogers, T., & Zhu, X. (2008). Human active learning. In D. Koller, Y. Bengio, D. Schuurmans, L. Bottou, & A. Culotta (Eds.), *Advances neural information processing systems* (Vol. 21, pp. 241–248). Cambridge, MA: MIT Press.
- Chi, M. T. H. (2009). Active-constructive-interactive: A conceptual framework for differentiating learning activities. *Topics in Cognitive Science, 1*, 73–105. doi:10.1111/j.1756-8765.2008.01005.x
- Chouinard, M. (2007). Children's questions: A mechanism for cognitive development. *Monographs of the Society for Research in Child Development, 72*, 1–126. doi:10.1111/j.1540-5834.2007.00412.x
- Clapper, J. P., & Bower, G. H. (2002). Adaptive categorization in unsupervised learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 28*(5), 908–923. doi:10.1037/0278-7393.28.5.908
- Cohn, D., Atlas, L., & Ladner, R. (1994). Improving generalization with active learning. *Machine Learning, 15*(2), 201–221. doi:10.1023/A:1022673506211
- Crutcher, R. J., & Healy, A. F. (1989). Cognitive operations and the generation effect. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15*(4), 669–675. doi:10.1037/0278-7393.15.4.669
- Dasgupta, S. (2010). Two faces of active learning. *Theoretical Computer Science, 412*(19), 1767–1781. doi:10.1016/j.tcs.2010.12.054
- Dasgupta, S., Kalai, A., & Monteleoni, C. (2005). Analysis of perception-based active learning. In P. Auer & R. Meir (Eds.), *Proceedings of 18th Annual Conference on Learning Theory (COLT)* (pp. 249–263). Heidelberg, Germany: Springer. doi:10.1007/11503415_17
- Daw, N. D., O'Doherty, J. P., Seymour, B., Dayan, P., & Dolan, R. J. (2006). Cortical substrates for exploratory decisions in humans. *Nature, 441*, 876–879. doi:10.1038/nature04766
- Dempster, F. N. (1989). Spacing effects and their implications for theory and practice. *Educational Psychology Review, 1*, 309–330. doi:10.1007/BF01320097
- Denrell, J. (2005). Why most people disapprove of me: Experience sampling in impression formation. *Psychological Review, 112*(4), 951–978. doi:10.1037/0033-295X.112.4.951
- Denrell, J. (2007). Adaptive learning and risk taking. *Psychological Review, 114*, 177–187. doi:10.1037/0033-295X.114.1.177
- Dunlosky, J., & Hertzog, C. (1998). Training programs to improve learning in later adulthood: Helping older adults educate themselves. In D. J. Hacker, J. Dunlosky, & A. C. Graesser (Eds.), *Metacognition in educational theory and practice* (pp. 249–276). Mahwah, NJ: Erlbaum.
- Edwards, W. (1965). Optimal strategies for seeking information: Models for statistics, choice reaction times, and human information processing. *Journal of Mathematical Psychology, 2*, 312–329. doi:10.1016/0022-2496(65)90007-6
- Elio, R., & Anderson, J. R. (1984). The effects of information order and learning mode on schema abstraction. *Memory & Cognition, 12*, 20–30. doi:10.3758/BF03196994
- Fazio, R., Eiser, J., & Shook, N. (2004). Attitude formation through exploration: Valence asymmetries. *Journal of Personality and Social Psychology, 87*(3), 293–311. doi:10.1037/0022-3514.87.3.293
- Fiedler, K. (2000). Beware of samples! A cognitive-ecological sampling approach to judgment biases. *Psychological Review, 107*, 659–676. doi:10.1037/0033-295X.107A659
- Fiedler, K., & Juslin, P. (Eds.). (2006). *Information sampling and adaptive cognition*. Cambridge, England: Cambridge University Press.
- Gibson, E. (1988). Exploratory behavior in the development of perceiving, active, and the acquiring of knowledge. *Annual Review of Psychology, 39*, 1–41. doi:10.1146/annurev.ps.39.020188.000245
- Gittins, J. C., & Jones, D. M. (1979). A dynamic allocation index for the discounted multiarmed bandit problem. *Biometrika, 66*(3), 561–565. doi:10.1093/biomet/66.3.561
- Glenberg, A. M. (1979). Component-levels theory of the effect of spacing of repetitions on recall and recognition. *Memory & Cognition, 7*, 95–112. doi:10.3758/BF03197590

- Gopnik, A., Glymour, C., Sobel, D., Schulz, L., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review, 111*(1), 1–31. doi:10.1037/0033-295X.111.1.3
- Gorman, M., Stafford, A., & Gorman, M. (1987). Disconfirmation and dual hypotheses on a more difficult version of Wason's 2-4-6 task. *The Quarterly Journal of Experimental Psychology, 39*(1), 1–28. doi:10.1080/02724988743000006
- Griffiths, T. L., & Tenenbaum, J. B. (2001). Generalization, similarity, and Bayesian inference. *Behavioral and Brain Sciences, 24*(4), 629–640. doi:10.1017/S0140525X01000061
- Gureckis, T. M., & Markant, D. M. (2009). Active learning strategies in a spatial concept learning task. In N. Taatgen, H. van Rijn, L. Schomaker, & J. Nerbonne (Eds.), *Proceedings of the 31st Annual Conference of the Cognitive Science Society* (pp. 3145–3150). Austin, TX: Cognitive Science Society.
- Gweon, H., Tenenbaum, J., & Schultz, L. E. (2010). Infants consider both the sample and the sampling process in inductive generalization. *Proceedings of the National Academy of Sciences, USA, 107*(20), 9066–9071. doi:10.1073/pnas.1003095107
- Harman, K. L., Humphrey, G. K., & Goodale, M. A. (1999). Active manual control of object views facilitates visual recognition. *Current Biology, 9*, 1315–1318. doi:10.1016/S0960-9822(00)80053-6
- Hau, R., Pleskac, T. J., Kiefer, J., & Hertwig, R. (2008). The description-experience gap in risky choice: The role of sample size and experienced probabilities. *Journal of Behavioral Decision Making, 21*, 493–518. doi:10.1002/bdm.598
- Hebb, D. (1955). Drives and the CNS (Central nervous system). *Psychological Review, 62*, 243–254. doi:10.1037/h0041823
- Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choices. *Psychological Science, 15*, 534–539. doi:10.1111/j.0956-7976.2004.00715.x
- Hills, T., & Hertwig, R. (2010). Information search and decisions from experience: Do our patterns of sampling foreshadow our decisions? *Psychological Science, 21*, 1787–1792. doi:10.1177/0956797610387443
- Huttenlocher, J. (1962). Effects of manipulation of attributes on efficiency of concept formation. *Psychological Reports, 10*, 503–509. doi:10.2466/PR.10.2.503-509
- Jacoby, L. L. (1978). On interpreting the effects of repetition: Solving a problem versus remembering a solution. *Journal of Verbal Learning and Verbal Behavior, 17*(6), 649–668. doi:10.1016/S0022-5371(78)90393-6
- Juni, M. Z., Gureckis, T. M., & Maloney, L. T. (2011). “Don't stop 'till you get enough”: Adaptive information sampling in a visuo-motor estimation task. In L. Carlson, C. Holscher, & T. Shipley (Eds.), *Proceedings of the 33rd Annual Conference of the Cognitive Science Society* (pp. 2854–2859). Austin, TX: Cognitive Science Society.
- Juslin, P., Winman, A., & Hansson, P. (2007). The naive intuitive statistician: A naive sampling model of intuitive confidence intervals. *Psychological Review, 114*, 678–703. doi:10.1037/0033-295X.114.3.678
- Kaelbling, L., Littman, M. L., & Cassandra, A. R. (1998). Planning and acting in partially observable stochastic domains. *Artificial Intelligence Journal, 101*, 99–134. doi:10.1016/S0004-3702(98)00023-X
- Kakade, S., & Dayan, P. (2002). Dopamine: Generalization and bonuses. *Neural Networks, 15*, 549–559. doi:10.1016/S0893-6080(02)00048-5
- Kaplan, R., Doeller, C. F., Barnes, G. R., Litvak, V., Duzel, E., Bandettini, P. A., & Burgess, N. (2012). Movement-related theta rhythm in humans: Coordinating self-directed hippocampal learning. *PLoS Biology, 10*(2), e1001267. doi:10.1371/journal.pbio.1001267
- Kemler Nelson, D. G., Egan, L. C., & Holt, M. (2004). When children ask, “What is it?” What do they want to know about artifacts? *Psychological Science, 15*, 384–389. doi:10.1111/j.0956-7976.2004.00689.x
- Kimball, D., & Metcalfe, J. (2003). Delaying judgments of learning affects memory, not metamemory. *Memory & Cognition, 31*(6), 918–929. doi:10.3758/BF03196445
- Kincannon, A., & Spellman, B. A. (2003). The use of category and similarity information in limiting hypotheses. *Memory & Cognition, 31*(1), 114–132. doi:10.3758/BF03196087
- Klahr, D., & Nigam, M. (2004). The equivalence of learning paths in early science instruction. *Psychological Science, 15*(10), 661–667. doi:10.1111/j.0956-7976.2004.00737.x
- Klayman, J., & Ha, Y.-W. (1987). Confirmation, disconfirmation, and information in hypothesis testing. *Psychological Review, 94*, 211–228. doi:10.1037/0033-295X.94.2.211
- Knox, W. B., Otto, A. R., Stone, P. H., & Love, B. C. (2012). The nature of belief-directed exploratory choice in human decision-making. *Frontiers in Psychology, 2*, 398. doi:10.3389/fpsyg.2011.00398
- Kolb, D. (1984). *Experiential learning: Experience as the source of learning and development*. Englewood Cliffs, NJ: Prentice Hall.
- Koriat, A., & Bjork, R. A. (2005). Illusions of competence in monitoring one's knowledge during study. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31*, 187–194. doi:10.1037/0278-7393.31.2.187
- Koriat, A., & Bjork, R. A. (2006). Illusions of competence during study can be remedied by manipulations that enhance learners' sensitivity to retrieval conditions at test. *Memory & Cognition, 34*, 959–972. doi:10.3758/BF03193244
- Kornell, N., & Bjork, R. A. (2007). The promise and perils of self-regulated study. *Psychonomic Bulletin & Review, 14*(2), 219–224. doi:10.3758/BF03194055
- Kornell, N., & Bjork, R. A. (2008). Learning concepts and categories: Is spacing the “enemy of induction”? *Psychological Science, 19*, 585–592. doi:10.1111/j.1467-9280.2008.02127.x
- Kornell, N., & Metcalfe, J. (2006). Study efficacy and the region of proximal learning framework. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 32*, 609–622. doi:10.1037/0278-7393.32.3.609
- Kruschke, J. (2008). Bayesian approaches to associative learning: From passive to active learning. *Learning & Behavior, 36*(3), 210–226. doi:10.3758/LB.36.3.210
- Kuhn, D., Black, J., Keselman, A., & Kaplan, D. (2000). The development of cognitive skills to support inquiry learning. *Cognition and Instruction, 18*(4), 495–523. doi:10.1207/S1532690XCI1804_3

- Kuhn, D., & Brannock, J. (1977). Development of the isolation of variables scheme in experimental and "natural experiment" contexts. *Developmental psychology, 13*(1), 9–14. doi:10.1037/0012-1649.13.1.9
- Kuhn, D., & Ho., V. (1980). Self-directed activity and cognitive development. *Journal of Applied Developmental Psychology, 1*, 119–133. doi:10.1016/0193-3973(80)90003-9
- Lagnado, D. A., & Sloman, S. (2004). The advantage of timely intervention. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 30*, 856–876. doi:10.1037/0278-7393.30.4.856
- Lang, K., & Baum, E. (1992). Query learning can work poorly when a human oracle is used. In *Proceedings of IEEE International Joint Conference on Neural Networks* (pp. 335–340). Washington, DC: IEEE Press.
- Leotti, L. A., Iyengar, S. S., & Oshsner, K. N. (2010). Born to choose: The origins and value of the need for control. *Trends Cognitive Science, 14*, 457–463. doi:10.1016/j.tics.2010.08.001
- Lombrozo, T. (2006). The structure and function of explanations. *Trends in Cognitive Sciences, 10*(10), 464–470. doi:10.1016/j.tics.2006.08.004
- Love, B. C., Jones, M., Tomlinson, M. T., & Howe, M. (2008). Predicting information needs: Adaptive display in dynamic environments. In B. C. Love, K. McRae, & V. M. Sloutsky (Eds.), *Proceedings of the 30th Annual Conference of the Cognitive Science Society* (pp. 875–880). Mahwah, NJ: Erlbaum.
- MacKay, D. J. C. (1992). Information based objective functions for active data selection. *Neural Computation, 4*(4), 589–603. doi:10.1162/neco.1992.4.4.590
- March, J. G. (1996). Learning to be risk averse. *Psychological Review, 103*, 309–319. doi:10.1037/0033-295X.103.2.309
- Markant, D., & Gureckis, T. M. (2010). Category learning through active sampling. In S. Ohlsson & R. Catrambone (Eds.), *Proceedings of the 32nd Annual Conference of the Cognitive Science Society* (pp. 248–253). Austin, TX: Cognitive Science Society.
- Markant, D., & Gureckis, T. M. (2012a). Does the utility of information influence sampling behavior? In N. Miyake, D. Peebles, & R. P. Cooper (Eds.), *Proceedings of the 34th Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.
- Markant, D., & Gureckis, T. M. (2012b). One piece at a time: Learning complex rules through self-directed sampling. In N. Miyake, D. Peebles, & R. P. Cooper (Eds.), *Proceedings of the 34th Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.
- Mayer, R. E. (2004). Should there be a three-strikes rule against pure discovery learning? The case for guided methods of instruction. *American Psychologist, 59*, 14–19. doi:10.1037/0003-066X.59.1.14
- Metcalf, J. (2002). Is study time allocated selectively to a region of proximal learning. *Journal of Experimental Psychology: General, 131*(3), 349–363.
- Metcalf, J. (2009). Metacognitive Judgments and control of study. *Current Directions in Psychological Science, 18*(3), 159–163. doi:10.1111/j.1467-8721.2009.01628.x
- Metcalf, J., & Kornell, N. (2003). The dynamics of learning and allocation of study time to a region of proximal learning. *Journal of Experimental Psychology: General, 132*(4), 530–542. doi:10.1037/0096-3445.132.4.530
- Mills, C., Legare, C. H., Bills, M., & Mejias, C. (2010). Preschoolers use questions as a tool to acquire knowledge from different sources. *Journal of Cognition and Development, 11*, 533–560. doi:10.1080/15248372.2010.516419
- Mill, J. (1950). *Philosophy of the scientific method*. New York, NY: Hafner. (Original work published 1843)
- Mitchell, T. (1997). *Machine learning*. New York, NY: McGraw-Hill.
- Montessori, M. (1964). *The Montessori method*. New York, NY: Schocken. (Original work published 1912)
- Murphy, G. L., & Ross, B. H. (2005). The two faces of typicality in category-based induction. *Cognition, 95*(2), 175–200. doi:10.1016/j.cognition.2004.01.009
- Murphy, K. P. (2001). *Active learning of causal Bayes net structure* (Technical Report). Berkeley: Department of Computer Science, University of California.
- National Research Council. (1999). *How people learn*. Washington, DC: National Academies Press.
- Nelson, J. D. (2005). Finding useful questions: On Bayesian diagnosticity, probability, impact, and information gain. *Psychological Review, 112*(4), 979–999. doi:10.1037/0033-295X.112.4.979
- 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. doi:10.1177/0956797610372637
- Nelson, T. O., & Leonesio, R. J. (1988). Allocation of self-paced study time and the "labor-in-vain effect." *Journal of Experimental Psychology: Learning, Memory, and Cognition, 14*, 676–686. doi:10.1037/0278-7393.14.4.676
- Nelson, T. O., & Narens, L. (1994). Why investigate metacognition? In J. Metcalfe & A. P. Shimamura (Eds.), *Metacognition: Knowing about knowing* (pp. 1–25). Cambridge, MA: MIT Press.
- Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology, 2*(2), 175–220. doi:10.1037/1089-2680.2.2.175
- Oaksford, M., & Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review, 101*(4), 608–631. doi: 10.1037/0033-295X.101.4.608
- Papert, S. (1980). *Mindstorms: Children, computers, and powerful ideas*. New York, NY: Basic Books.
- Pavlik, P. I., & Anderson, J. R. (2008). Using a model to compute the optimal schedule of practice. *Journal of Experimental Psychology: Applied, 14*(2), 101–117. doi:10.1037/1076-898X.14.2.101
- Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. Cambridge, England: Cambridge University Press.
- Payne, J. W., Braunstein, M. L., & Carroll, J. S. (1978). Exploring predecisional behavior: An alternative approach to decision research. *Organizational Behavior and Human Performance, 22*, 17–44.

- Pearl, J. (2000). *Causality: Models, reasoning, and inference*. Cambridge, England: Cambridge University Press.
- Piaget, J. (1930). *The child's conception of physical causality*. New York, NY: Harcourt, Brace.
- Platt, J. R. (1964). Strong inference. *Science*, *146*(3642), 347–353. doi:10.1126/science.146.3642.347
- Popper, K. R. (1959). *The logic of scientific discovery*. London, England: Hutchinson. (Original work published 1935)
- Rehder, B., & Hoffman, A. B. (2005). Eyetracking and selective attention in category learning. *Cognitive Psychology*, *51*, 1–41. doi:10.1037/0278-7393.31.5.811
- Rhodes, M., Gelman, S. A., & Brickman, D. (2010). Children's attention to sample composition in learning, teaching, and discovery. *Developmental Science*, *13*(3), 421–429. doi:10.1111/j.1467-7687.2009.00896.x
- Roscoe, R. D., & Chi, M. T. H. (2007). Understanding tutor learning: Knowledge-building and knowledge-telling in peer tutors' explanations and questions. *Review of Educational Research*, *77*(4), 534–574. doi:10.3102/0034654307309920
- Roscoe, R. D., & Chi, M. T. H. (2008). Tutor learning: The role of explaining and responding to questions. *Instructional Science*, *36*(4), 321–350.
- Rottman, B. M., & Keil, F. C. (2012). Causal structure learning over time: Observations and interventions. *Cognitive Psychology*, *64*, 93–125. doi:10.1016/j.cogpsych.2011.10.003
- Schmidhuber, J. (1991). Curious model-building control systems. In *Proceedings of IEEE International Joint Conference on Neural Networks* (Vol. 2, pp. 1458–1463). Washington, DC: IEEE Press.
- Schulz, L., & Bonawitz, E. (2007). Serious fun: Preschoolers engage in more exploratory play when evidence is confounded. *Developmental Psychology*, *43*, 1045–1050. doi:10.1037/0012-1649.43.4.1045
- Settles, B. (2009). *Active learning literature survey* (Computer Sciences Technical Report No. 1648). Madison: University of Wisconsin, Madison.
- Seung, H., Opper, M., & Sompolinsky, H. (1992). Query by committee. In *Proceedings of ACM Workshop on Computational Learning Theory* (pp. 287–294). New York, NY: ACM.
- Shafto, P., & Goodman, N. D. (2008). Teaching games: Statistical sampling assumptions for learning in pedagogical situations. In B. C. Love, K. McRae, & V. M. Sloutsky (Eds.), *Proceedings of the 30th Annual Conference of the Cognitive Science Society* (pp. 1632–1637). Austin, TX: Cognitive Science Society.
- Shafto, P., Goodman, N. D., & Frank, M. (in press). Learning from others: The consequences of psychological reasoning for human learning. *Perspectives on Psychological Science*.
- Shepard, R. N. (1987). Toward a universal law of generalization for psychological science. *Science*, *237*(4820), 1317. doi:10.1126/science.3629243
- Shepard, R. N., Hovland, C. I., & Jenkins, H. M. (1961). Learning and memorization of classifications. *Psychological Monographs: General and Applied*, *75*(13), 1–41. doi:10.1037/h0093825
- Simon, D. A., & Bjork, R. A. (2001). Metacognition in motor learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *27*, 907–912. doi:10.1037/0278-7393.27.4.907
- Skov, R., & Sherman, S. (1986). Information-gathering processes: Diagnosticity, hypothesis-confirmatory strategies, and perceived hypothesis confirmation. *Journal of Experimental Social Psychology*, *22*(2), 93–121. doi:10.1016/0022-1031(86)90031-4
- Sobel, D. M., & Kushnir, T. (2006). The importance of decision making in causal learning from interventions. *Memory & Cognition*, *34*, 411–419. doi:10.3758/BF03193418
- Son, L. K. (2004). Spacing one's study: Evidence for a metacognitive control strategy. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *30*, 601–604. doi:10.1037/0278-7393-30.3.601
- Son, L. K., & Kornell, N. (2008). Research on the allocation of study time: Key studies from 1890 to the present (and beyond). In J. Dunlosky & R. A. Bjork (Eds.), *Handbook of metamemory and memory* (pp. 333–351). New York, NY: Taylor & Francis.
- Son, L. K., & Kornell, N. (2009). Simultaneous decisions at study: Time allocation, ordering, and spacing. *Metacognition Learning*, *4*, 237–248. doi:10.1007/s11409-009-9049-1
- Son, L. K., & Metcalfe, J. (2000). Metacognitive and control strategies in study-time allocation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *26*, 204–221. doi:10.1037//0278-7393,26.1.204
- Son, L. K., & Sethi, R. (2010). Adaptive learning and the allocation of time. *Adaptive Behavior*, *18*, 132–140. doi:10.1177/1059712309344776
- Steyvers, M., Tenenbaum, J. B., Wagenmakers, E. J., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, *27*, 453–489. doi:10.1016/S0364-0213(03)00010-7
- Sutton, R., & Barto, A. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.
- Tenenbaum, J. B. (1999). Bayesian modeling of human concept learning. In M. Kearns, S. Solla, & D. Cohn (Eds.), *Advances in neural information processing systems* (Vol. 11, pp. 59–65). Cambridge, MA: MIT Press.
- Todd, P. M., & Dieckmann, A. (2005). Heuristics for ordering cue search in decision making. In L. K. Saul, Y. Weiss, L. Bottou (Eds.), *Advances in neural information processing systems* (Vol. 17, pp. 1393–1400). Cambridge, MA: MIT Press.
- Todd, P. M., Hills, T., & Robbins, T. (Eds.). (in press). *Cognitive search: Evolution, algorithms, and the brain, vol. 9* (Strüngmann Forum Reports). Cambridge, MA: MIT Press.
- Tong, S., & Koller, D. (2001, August). Active learning for structure in Bayesian networks. *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence* (pp. 863–869). Palo Alto, CA: AAAI Press.
- Tversky, A., & Edwards, W. (1966). Information versus reward in binary choices. *Journal of Experimental Psychology*, *71*(5), 680–683. doi:10.1037/h0023123
- Tweney, R. D., Doherty, M. E., Worner, W. J., Pliske, D. B., Mynatt, C. R., Gross, K. A., & Arkkelin, D. L. (1980). Strategies of rule discovery in an inference task. *Quarterly Journal of Experimental Psychology*, *32*, 109–123. doi:10.1080/0033558008248237
- Valiant, L. G. (1984). A theory of the learnable. *Communications of the ACM*, *27*(11), 1134–1142.

- Voss, J. L., Gonsalves, B. D., Federmeier, K. D., Tranel, D., & Cohen, N. H. (2011). Hippocampal brain-network coordination during volitional exploratory behavior enhances learning. *Nature Neuroscience, 14*, 115–120. doi:10.1038/nn.2693
- Vygotsky, L. S. (1987). *The collected works of L. S. Vygotsky: Vol. 1. Problems of general psychology including the volume thinking and speech*. New York, NY: Plenum Press.
- Wason, P. (1960). On the failure to eliminate hypotheses in a conceptual task. *Quarterly Journal of Experimental Psychology, 12*(3), 129–140. doi:10.1080/17470216008416717
- Wittmann, B. C., Daw, N. D., Seymour, B., & Dolan, R. J. (2008). Striatal activity underlies novelty-based choice in humans. *Neuron, 58*, 967–973. doi:10.1016/j.neuron.2008.04.027
- Xu, F., & Tenenbaum, J. B. (2007). Sensitivity to sampling in Bayesian word learning. *Developmental Science, 10*, 228–297. doi:10.1111/j.1467-7687.2007.00590.x
- Yu, A. J., & Dayan, P. (2003). Expected and unexpected uncertainty: ACh and NE in the neocortex. In S. Becker, S. Thrun, & K. Obermayer (Eds.), *Advances in neural information processing systems* (Vol. 15, pp. 173–180). Cambridge, MA: MIT Press.