

# A taxonomy of inductive problems

Charles Kemp · Alan Jern

© Psychonomic Society, Inc. 2013

**Abstract** Inductive inferences about objects, features, categories, and relations have been studied for many years, but there are few attempts to chart the range of inductive problems that humans are able to solve. We present a taxonomy of inductive problems that helps to clarify the relationships between familiar inductive problems such as generalization, categorization, and identification, and that introduces new inductive problems for psychological investigation. Our taxonomy is founded on the idea that semantic knowledge is organized into systems of objects, features, categories, and relations, and we attempt to characterize all of the inductive problems that can arise when these systems are partially observed. Recent studies have begun to address some of the new problems in our taxonomy, and future work should aim to develop unified theories of inductive reasoning that explain how people solve all of the problems in the taxonomy.

**Keywords** Induction · Semantic cognition · Generalization · Categorization · Discovery · Identification · Reasoning

Attempts to systematize knowledge have proven useful in several fields. Mendeleev presented a periodic table of the chemical elements that helped to clarify relationships between the known elements and that made predictions about the existence of new elements. Adelson and Bergen (1991) developed a “periodic table” of early vision that mapped out a space of visual and identified several that had previously received little attention. This article aims to make a similar contribution to the study of inductive reasoning. We describe a taxonomy of inductive problems that aims to clarify the relationships between familiar problems and to highlight problems that have previously been overlooked.

Inductive reasoning has been discussed by researchers from many fields (Hayes, Heit, & Swendsen, 2010; Heit,

2000; Holland, Holyoak, Nisbett, & Thagard, 1986; Vickers, 2012), and the term “induction” has been defined both broadly and narrowly (Colberg, Nester, & Trattner, 1985; Thagard, 2001). We will adopt a broad definition and will consider an inference to be inductive if the conclusion does not follow deductively from the premises (Chater, Oaksford, Hahn, & Heit, 2011; Holland et al., 1986; Skyrms, 1975). Inferences of this kind are sometimes called *ampliative*, because the conclusion goes beyond the information given and is at best likely rather than certain given the available evidence. An alternative tradition uses “induction” more narrowly to refer to any inference that moves from specific observations (e.g., Bob and Bill are mortal) to a general conclusion (e.g., all men are mortal; Quine & Ullian, 1978; Vickers, 2012). From this perspective, the set of ampliative inferences includes some that are inductive and others that are instances of *abduction* (Peirce, 1957) and of other kinds of reasoning. Although both broad and narrow definitions can be found in the psychological literature, contemporary work on induction tends to adopt the broad rather than the narrow view (Sloman, 2007).

The broad definition of inductive reasoning characterizes induction in opposition to deduction, and adopting this definition means that inductive problems form “a large and varied set” (Heit, 2008, p. 323) that includes “a vast number of argument forms and types” (Sloman, 2007, p. 329). The broad definition therefore raises the need for a taxonomy that provides a systematic characterization of the space of inductive problems (Bisanz, Bisanz, & Korpan, 1994). Inductive inference is relevant to just about every area of cognition; it takes place, for example, when humans predict the motion of an occluded object, assess the grammaticality of a novel sentence, or decide how to grasp an object that is encountered for the first time. We will not focus on vision, language, or motor control, but will instead focus on a cluster of problems from an area that has been called *semantic cognition* (Rogers & McClelland, 2004). Research in this area aims to capture knowledge about objects, features, categories, relationships between objects, and word meanings. The relevant literature includes studies of property induction (Gelman & Markman,

---

C. Kemp (✉) · A. Jern  
Department of Psychology, Carnegie Mellon University, 5000  
Forbes Avenue, Baker Hall 340T, Pittsburgh, PA 15213, USA  
e-mail: ckemp@cmu.edu

1986; Imai, 1995), categorization (both supervised [Nosofsky, 1986] and unsupervised [Pothos & Chater, 2002]), stimulus generalization (Shepard, 1987), identification (Nosofsky, 1986), and word learning (Markman, 1989). This article develops a taxonomy of inductive problems that includes all of these classic problems, along with others that have received little attention.

Taxonomies of inductive reasoning can be based on at least two different perspectives. Following Heit (2007), we refer to these perspectives as the *problem view* and the *process view*.<sup>1</sup> Our taxonomy adopts the problem view and attempts to characterize the space of inductive problems that people are able to solve. Each individual problem is characterized by describing the input available to the reasoner and the output that the reasoner must generate. Importantly, characterizations of these problems can be provided without specifying the psychological processes that allow the reasoner to convert the input into the output (Marr, 1982). Instead of focusing on the space of inductive problems, the process view proposes that taxonomies of inductive reasoning should aim to characterize the psychological processes that support induction. Sternberg (1986) presented a taxonomy along these lines that proposes that inductive reasoning involves three processes: selective encoding, selective comparison, and selective combination.

At first it might seem that psychological research should focus exclusively on the process view of induction, but we believe that the problem view is a necessary precursor to the process view. Any psychologist who sets out to understand inductive reasoning will need to consider data gathered from a variety of tasks, and characterizing the space of inductive problems is necessary in order to decide which tasks are relevant. For example, Sternberg (1986) presumably developed his process-based taxonomy by selecting several problems that are broadly representative of the space of inductive problems (the three that he discusses are analogical reasoning, series completion, and classification) and reflecting on the processes that enable these problems to be solved. Although some characterization of the problem space is essential, this characterization could be informal and pretheoretical. For example, a researcher might implicitly adopt a problem space that corresponds to the set of all problems that are commonly studied in the literature on inductive reasoning. From this perspective, the real question is not whether some characterization of the problem space is necessary, but whether it is possible to improve upon the informal characterization that is implicit in much psychological research.

We propose that a systematic characterization of the space of inductive problems can contribute to the field in at least four respects. First, a taxonomy of inductive problems can reveal the similarities and differences between problems that have been discussed in different parts of the literature. The popularity of inductive reasoning as a research area has led to a fragmentation of the literature that has made it difficult to understand the relationships between the problems that have been studied thus far. One symptom of this fragmentation is a proliferation of inconsistent terminology (Reber & Reber, 2001). For example, a reader might reasonably assume that “feature induction” and “feature learning” are two different names for the same problem, and that “categorical induction” refers to a distinct problem. In reality, “feature induction” (Murphy, 1993) and “categorical induction” (Sloman & Lagnado, 2005) are different names for the same problem, and “feature induction” (Murphy, 1993) and “feature learning” (Austerweil & Griffiths, 2011) are similar names for different problems. We will return to these problems later, but the point for now is that a successful taxonomy should help to avoid terminological confusion.

Second, a taxonomy of inductive problems can help to resolve theoretical disputes that turn on the nature of the problems posed by a given task. For example, consider a task in which participants learn that an object has a hidden feature and are subsequently exposed to a second, similar-looking object. There are at least two possible reasons to believe that the second object has the hidden feature (Brown, 1965). Some participants might understand that the second object is different from the first, but might infer that both objects have the hidden feature. Other participants might think that the second object has the hidden feature because they mistakenly identify it as the first object. In terms of the taxonomy that we will develop, the first explanation focuses on the problem of *generalization*, and the second explanation focuses on the problem of *identification*. The failure to acknowledge both of these explanations has led to some confusion in the literature (Chater, Vitanyi, & Stewart, 2001). For example, there has been some debate about whether generalization gradients are closer to Gaussian functions (Nosofsky, 1986) or to exponential functions (Shepard, 1987). One proposed resolution is that pure generalization curves are exponential, but that inferences about highly confusable objects include an identification component that produces near-Gaussian generalization curves (Ennis, 1988; Nosofsky, 1988; Shepard, 1986). Examples of this kind suggest that thinking carefully about the problems posed by a given task is a useful first step before attempting to characterize the underlying psychological processes.

Third, a taxonomy of inductive problems can reveal novel problems that can be explored by future empirical studies. Even though core inductive problems such as generalization and identification have been studied for many years, other important problems have been neglected. To illustrate this

<sup>1</sup> Heit (2007) uses these terms to refer to two proposals about how the relationship between induction and deduction should be characterized. Here we adopt his terminology to refer to two proposals about how inductive reasoning should be characterized.

point, we will describe recent studies of feature identification (Kemp, Chang, & Lombardi, 2010), category generation (Jern & Kemp, 2013), and simultaneous object and feature generalization (Kemp, Shafto, & Tenenbaum, 2012). All three studies address novel inductive problems that were explored while developing our taxonomy.

Fourth, a systematic characterization of the space of inductive problems is a useful step toward developing theories that can explain how people solve all of these problems. Some previous theories have been able to handle multiple inductive problems—for example, exemplar models have been used to account for problems including categorization and identification (Estes, 1994; Nosofsky, 1992). No current theory, however, comes close to handling all of the problems in our taxonomy. Because we aim to characterize inductive problems rather than to describe the psychological processes that allow them to be solved, we hope that our taxonomy will be useful to researchers from many different traditions, including modelers who pursue probabilistic, exemplar-based, or connectionist approaches. Toward the end of the article, however, we will argue that a probabilistic approach provides an especially natural way to work toward a unified theory that can address all of the problems that we consider.

Accounts of semantic cognition differ in many respects, but most of them rely on objects, features, categories, and relations. Our taxonomy takes these basic notions as a starting point and attempts to chart the space of inductive problems that can be posed given a commitment to these notions. We will develop the taxonomy by characterizing *semantic systems* of objects, features, categories, and relations that capture the state of some part of the world. We will then consider problems in which a reasoner receives incomplete information about a system and must make inductive inferences about unobserved aspects of the system. We will consider three basic problems, which we refer to as *generalization*, *discovery*, and *identification*. These three problems take different forms when defined over different kinds of semantic systems, and can be combined with each other to generate additional problems. As a result, the taxonomy that we describe includes a large number of distinct problems.

### Semantic systems


We begin by describing how semantic systems of objects, features, categories, and relations can be characterized. Figure 1 shows examples of five systems. Figure 1a is a system that includes four objects, three features (size, color, and texture), and a category that includes the large objects. Figure 1b is a system that includes four objects and a directed binary relation. For example, the objects could be baboons, and the relation could indicate which baboons dominate each other. Figure 1c is a molecule of methane, which includes five atoms and an

undirected relation that represents chemical bonds between some pairs of atoms. The system includes categories that specify the kind of each atom, and may also include features that specify the mass and other properties of each atom. Figure 1d is a system that includes multiple relations and features. The system is a nuclear family, and the objects in the system are seven individuals. The features, categories, and relations in the system are not shown in full in Fig. 1d, and may include a category that specifies the gender of each individual and kinship relations such as *parent(·,·)* and *spouse(·,·)*.

Figures 1a–d illustrate how objects, features, categories, and relations can be combined to construct systems. These systems in turn can be treated as “compound objects” over which features and relations can be subsequently defined. Figure 1e shows an example that includes eight systems, each of which includes three objects (i.e., three gray squares). The higher-level system in Fig. 1e includes a category defined over these systems that includes all systems with two or more slashes. Event categories such as “armed robbery” are real-world examples of categories that can be formulated similarly. Each robbery can be viewed as a semantic system that specifies relationships between components that include the robber, the victim, and the item that was stolen (Gentner & Kurtz, 2005). The category of “armed robberies” includes all systems that indicate that a weapon was used to commit the crime. Higher-level systems like Fig. 1e can also include features and relations defined over lower-level systems, and can in turn be treated as compound objects that are used to construct systems at an even higher level.

Figures 1a and e are examples in which categories correspond to classes of items. Categories, however, can also be viewed as abstract entities that can bear features. For example, generic statements such as “the dodo has a beak” and “the dodo is extinct” appear to correspond to claims about the features of the category named “dodo.” Statements of this kind cannot always be paraphrased as claims that all members of the category have the feature in question. For example, “is extinct” is a feature that can sensibly be applied to a category, but not to any individual member of the category (Carlson, 2010). As a result, our taxonomy includes problems in which categories are treated as first-class entities that can bear features in their own right. We distinguish between problems of this kind and problems in which categories are simply treated as classes of objects.

We have introduced the notion of a semantic system relatively informally, but Kemp (2012) shows how this notion can be formally captured using set-theoretic machinery. If desired, the same formal approach could be used to characterize the space of semantic systems that serves as the foundation of our taxonomy. Here, however, our primary goal is to characterize the space of inductive problems, and an informal characterization of the space of semantic systems will suffice for this purpose.



**Fig. 1** Semantic systems. **(a)** A system with four objects, three binary features (color, size, and texture), and a category that includes the two large objects. **(b)** A system with four objects and a single, directed binary relation. **(c)** A molecule of methane can be viewed as a system that includes five objects (atoms), categories that indicate the kind of each atom, and a binary relation that indicates which pairs of atoms are bonded to each other. **(d)** The nuclear family shown is a system that includes six relatives of an individual, labeled as Ego: his mother (M),

father (F), younger and elder sisters (Zy and Ze), and younger and elder brothers (By and Be). The system may include a category that specifies the sex of each individual, as well as several relations, including kinship relations like *parent*( $\cdot, \cdot$ ) **(e)** A high-level semantic system defined over eight compound objects, each of which is a three-object semantic system in its own right. The high-level system includes a category, shown as a dashed rectangle, that includes all compound objects that have two or more objects with slashes

## A taxonomy of inductive problems

In the previous section, we described how semantic systems can be constructed by combining objects, features, categories, and relations. An inductive problem arises when a reasoner is given incomplete information about a system. This section describes three basic ways in which knowledge about a system might be incomplete. These three kinds of incompleteness lead to three inductive problems that we will refer to as *generalization*, *discovery*, and *identification*.

The three basic problems are very general, and versions of these problems can arise with respect to any semantic system. Our initial discussion of these problems will focus on simple semantic systems that can be visualized as matrices. The rows of each matrix correspond either to individual objects (e.g., a specific elephant and a specific rhino) or to categories (e.g., elephants and rhinos). The columns correspond either to features (e.g., “is large,” “has a tail,” “moves slowly”) or to categories (e.g., “is a mammal,” “is an animal”). Figure 2a shows example matrices that specify information about five objects, three features, and two categories. The black cells in the first matrix indicate which features the objects have. For example, object  $o_1$  has feature  $f_1$  and feature  $f_2$ . The black cells in the second matrix indicate which categories the objects belong to. For example, object  $o_1$  belongs to category  $c_1$ , and object  $o_4$  belongs to category  $c_2$ .


The problems of generalization, discovery, and identification each come in many forms. We will introduce each problem by focusing on cases in which there is uncertainty about a single component of a semantic system—for example, uncertainty about a single feature, a single category, or a

single object. In general, however, the three inductive problems may involve uncertainty about multiple objects, features, and categories, and we will describe one such example for each of the three problems.

Many of the problems that we will describe have been studied by psychologists, and this section includes many references to previous work. We will not exhaustively survey the literature on inductive reasoning, but will focus instead on studies that have been especially influential and on studies that illustrate the range of different labels that have been used for the problems that we consider. Many of the articles that we will cite have presented both empirical results and theoretical approaches to inductive reasoning, but this section will not survey these results and theoretical approaches. The primary goal here is to provide a unified account of the space of inductive problems, and the next section will discuss the prospect of developing a unified theoretical account of induction.

### Generalization

Generalization problems arise when a reasoner knows about a semantic system including objects, features, and categories, but does not know the feature values or category assignments for all objects in the system. Figure 3a shows an example in which the semantic system is represented as two matrices. Black squares in the first matrix indicate cases in which an object is known to have a certain feature, and white squares indicate cases in which an object is known not to have a certain feature. Gray squares indicate feature values that are unknown and must be inferred. Similarly, black and white squares in the second matrix indicate cases in which an



**Fig. 2** (a) A semantic system that includes objects, features, and categories can be visualized using two matrices. The  $(i, j)$  cell in the first matrix is black if object  $o_i$  has feature  $f_j$ . The  $(i, j)$  cell in the second matrix is black if object  $o_i$  is a member of category  $c_j$ . (b) A system that includes several additional relations can be visualized by adding several

square matrices of objects by objects to the matrices of objects by features and objects by categories. Three square matrices are shown, one for each relation. The  $(i, j)$  cell in matrix  $k$  is black if relation  $r_k$  holds between objects  $o_i$  and  $o_j$

object is known to belong or not to belong to a category, and gray squares indicate cases in which the category assignment of an object is unknown.

*Feature generalization.* Feature generalization is a problem in which a reasoner makes inferences about one feature at a time. Figure 3a shows an example in which a reasoner observes that object  $o_1$  has feature  $f_7$ , and must infer which of the remaining objects  $o_2$  through  $o_5$  have this feature. For example, suppose that you learn that Elmer the elephant has enzyme X132 in his blood, and you need to predict whether each remaining animal in the same zoo will test positive for the enzyme. You might predict that Ronald the rhino is more likely to have this enzyme than Samuel the skunk. Figure 3a shows an example in which only one object is known to have the novel feature, but in general multiple observations may be available—for example, you might know that objects  $o_1$  and  $o_2$  have a novel feature but that  $o_3$  does not, and might have to infer whether or not  $o_4$  has the novel feature.

Figure 3 shows 24 problems in total and includes a label for each of these problems. The label for the feature generalization problem in Fig. 3 is  $Gn(O, f)$ .  $Gn$  is short for generalization, and the  $O$  and the  $f$  indicate that the problem requires an inference about a matrix of objects by features. We will use lowercase letters to indicate problems that focus on inferences about a single element in a semantic system. For example, the  $f$  in  $Gn(O, f)$  indicates that this problem focuses on inferences about a single feature. Uppercase letters do not constrain the number of elements involved—for example, the  $O$  in  $Gn(O, f)$  indicates that the problem involves inferences about one or more objects.

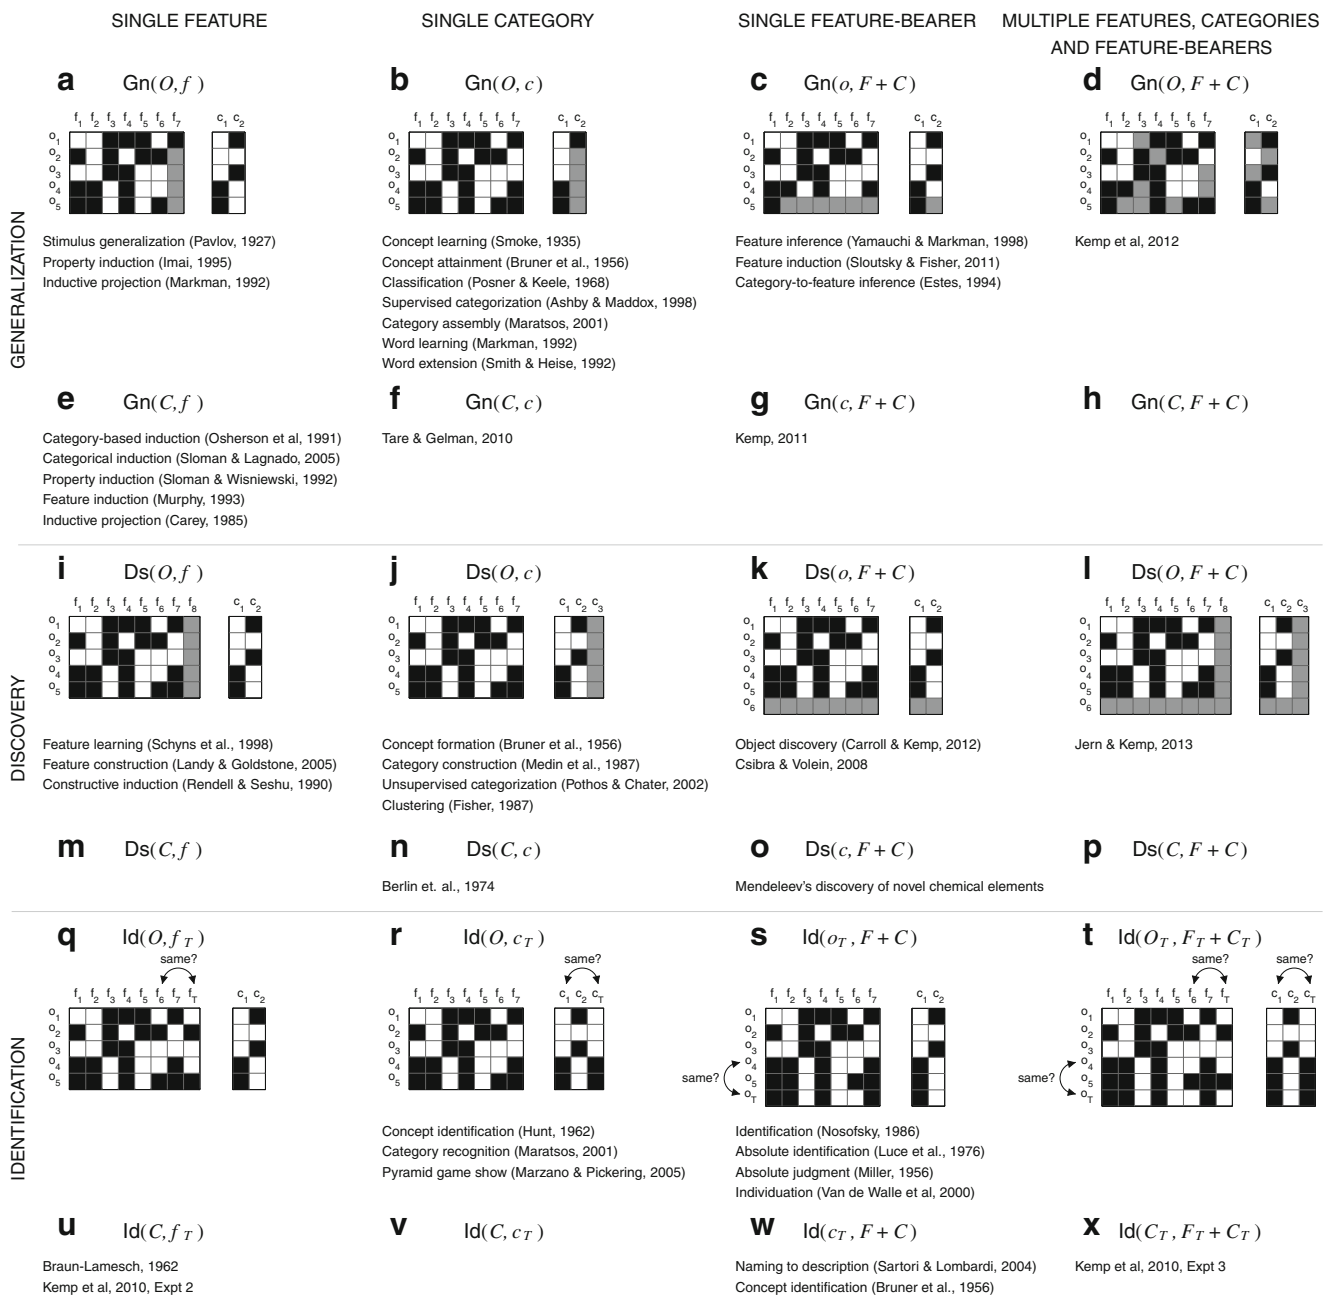
Problem  $Gn(O, f)$  has been extensively studied and is discussed, for example, in the literature on “stimulus generalization” (Gluck, 1991; Guttman & Kalish, 1956; Hull, 1943; Pavlov, 1927; Shepard, 1957, 1987; Skinner, 1938; Tenenbaum & Griffiths, 2001). A typical experiment in this literature might present a reasoner with a single object that has a desirable feature—for example, a reasoner might be given a berry that

tastes good. The reasoner is then presented with another berry that might be different in size, shape, or color, and is required to decide whether this second berry also tastes good.

Problem  $Gn(O, f)$  has also been discussed in the literature on conceptual development (Gelman & Markman, 1986, 1987) and is known there as the “problem of property induction” (Imai, 1995) or “inductive projection” (Markman, 1992). For example, Gelman and Markman (1987) used a task in which children were shown a picture of a small blue bird and told that the bird “gives its baby mashed up food.” The children were then asked to decide whether the property applied to other objects, including a blackbird and a small blue butterfly.

Problem  $Gn(O, f)$  asks a reasoner to project a feature from one object to other objects, and problem  $Gn(C, f)$  in Fig. 3e is a closely related problem in which a reasoner must project a feature across categories. For example, given that elephants have enzyme X132 in their blood, a reasoner might be asked to judge how likely it is that rhinos have enzyme X132 in their blood. Problem  $Gn(C, f)$  has been studied using both children (Carey, 1985) and adults (Rips, 1975), and goes by various names including “category-based induction” (Gelman, 2003; Osherson, Smith, Wilkie, Lopez, & Shafir, 1990), “categorical induction” (Sloman & Lagnado, 2005), “property induction” (Sloman & Wisniewski, 1992), “feature induction” (Hadjichristidis, Sloman, Stevenson, & Over, 2004; Murphy, 1993), and “inductive projection” (Carey, 1985; Rogers & McClelland, 2004). Figure 3e does not include a visual representation of the problem, but this representation can easily be created by adjusting the matrices in Fig. 3a so that the rows represent categories rather than individual objects.

*Category generalization.* Category generalization is similar to feature generalization, but involves inferences about a partially observed category rather than a partially observed feature. Figure 3b shows an example in which a reasoner learns that object  $o_1$  belongs to category  $c_2$  and must infer which other objects belong to this category. The label for this inductive problem is  $Gn(O, c)$ , in which the lowercase  $c$



**Fig. 3** Generalization, discovery, and identification problems that arise when reasoning about systems of objects, features, and categories. Black cells indicate features or category labels that are known to apply: For example, in (a) object  $o_1$  has  $f_3$  and is a  $c_2$ . White cells indicate features or category labels that are known not to apply: For example, in (a) object  $o_1$  does not have  $f_1$  and is not a  $c_1$ . Gray cells indicate entries with unknown

indicates that the problem focuses on inferences about a single category.

Problem  $Gn(O, c)$  has been widely studied and goes by many names including “concept learning” (Smoke, 1935), “concept attainment” (Bruner, Goodnow, & Austin, 1956), “classification” (Posner & Keele, 1968), “supervised categorization” (Ashby & Maddox, 1998), “category assembly” (Maratsos, 2001), “word learning” (Markman, 1992), and “word extension” (Smith &

Heise, 1992). A typical experiment in this literature might present a reasoner with several objects that are all said to be “wugs.” The reasoner is then shown a novel object and asked to decide whether or not it is also a wug. For present purposes, two aspects of this *category generalization* problem are especially important. First, category labels are provided for some objects, which means that the reasoner has direct evidence of the existence of the category in question. Problems in which labels are available

status. The columns of panels show four versions of each problem: a version that focuses on a single feature, a version that focuses on a single category, a version that focuses on a single feature-bearer (i.e., a single object or category), and a version that focuses on multiple features, categories, and feature-bearers. Selected references and problem names are included for problems that have been discussed by previous researchers

can be distinguished from *category discovery* problems (Fig. 3j), in which reasoners spontaneously form categories by noticing coherent clusters of objects. Second, category generalization problems require reasoners to learn novel categories, not just novel labels for categories that they already know. Problems involving novel categories can be distinguished from *category identification* problems (Fig. 3r), in which reasoners learn novel labels for preexisting categories.

The distinctions just described between category generalization, category discovery, and category identification are consistent with distinctions that have been proposed by previous researchers (Bruner et al., 1956; Maratsos, 2001). These distinctions, however, are by no means universally accepted. The literature on categorization includes many different names for inductive problems, and these names are often used in inconsistent ways. For example, Reber and Reber (2001) pointed out that the literature “abounds with terms,” including “concept acquisition, concept development, concept discovery, concept identification, concept use, concept attainment, and concept induction,” and that there is “precious little agreement about terminology” (p. 141). The references in Fig. 3b indicate that some researchers have used names such as “concept learning” and “concept attainment” to refer to problem  $Gn(O, c)$ , but this should not be taken to suggest that these names always refer to problem  $Gn(O, c)$ .

Problem  $Gn(C, c)$  is a version of the problem in Fig. 3b in which the rows of the matrices represent categories rather than objects. For example, a child might be told that mice, dogs, and elephants are “feps,” and then asked to decide which other categories belong to the category of feps. Problem  $Gn(C, c)$  has received relatively little attention, but a study by Tare and Gelman (2010) comes close to addressing a version of the problem. In one of their conditions, Tare and Gelman informed participants that “apples are feps,” and then showed them pictures of a balloon, a bunch of grapes, and a knife and asked them to “point to another fep.” Because the task asked for inferences about three specific objects, it does not qualify as an example of problem  $Gn(C, c)$ . The task, however, could be converted into a genuine example of problem  $Gn(C, c)$  if pictures had not been used and participants had simply been asked whether balloons are feps, whether grapes are feps, and whether knives are feps.

*Object generalization and category generalization.* The generalization problems described so far have involved inferences about a single feature or a single category. Object generalization is a companion problem that focuses on inferences about a single object. Some of the feature values and category assignments for the object are known, and the reasoner must infer the values of all remaining features and categories. We refer to this problem as  $Gn(o, F+C)$ , in which the lowercase  $o$  and the uppercase  $F+C$  indicate that the problem focuses on inferences about the features and category assignments of a single object.

Figure 3c shows that object generalization can be visualized as a matrix completion problem in which the goal is to complete a single row. For example, suppose that you learn that scientists have discovered a new creature, but all that you know so far is that the creature flies. You might be able to predict some of the other features that the creature would have; for example, it probably has wings, it probably does not have gills, and there is a good chance that the creature is a bird. The example in Fig. 3c specifies exactly two observations concerning the novel object—a feature value and a category label—but examples involving one, two, or more than two observations along the bottom row are all valid instances of object generalization.

Object generalization has been studied by many researchers and is often described as the problem of “feature inference” (Anderson, Ross, & Chin-Parker, 2002; Rehder & Burnett, 2005; Sweller & Hayes, 2010; Yamauchi & Markman, 1998). In a typical feature inference task, reasoners observe one or more features of a novel object and then predict which other features the object is likely to have (Hayes & Thompson, 2007; McCarrell & Callanan, 1995; Murphy & Ross, 2010). The same problem is occasionally described as “feature induction” (Sloutsky & Fisher, 2002), although this name is more commonly used to refer to the problem in Fig. 3e.

Object generalization has also been addressed by studies that explore inferences about novel nouns that are used in context. For example, a child who hears that “Mommy feeds the ferret” has been given some features of the referent of “ferret” (e.g., that it eats) and may be able to infer additional features (e.g., that the referent is animate; Goodman, McDonough, & Brown, 1998). Similarly, a child who hears that the boojum is hungry might be able to infer which other features the boojum is likely to have (it probably has a mouth) and which categories it belongs to (it probably is an animal; Keil, 1979).

Object generalization has also been explored in depth in the literature on categorization (Anderson, 1991; Anderson & Fincham, 1996). In one version of the problem, reasoners observe some features of a novel object and must decide which category labels apply to the object. For example, a reasoner might observe that a novel creature has wings and a beak and might categorize it as a bird. In a second version of the problem, reasoners are given the category label of a novel object and must decide which features the object is likely to have. For example, upon being told that a novel creature is a bird, a reasoner may infer that it has wings and a beak. In cases in which the category label of the novel object is known (Markman, 1989), object generalization is sometimes described as the problem of “category-to-feature inference” (Estes, 1994).

The object generalization problem in Fig. 3g is a version of the problem in Fig. 3c in which the rows of the matrix represent categories rather than objects. We refer to this problem as  $Gn(c, F+C)$ , because it involves inferences about the features and category assignments of a single category. Problem  $Gn(c, F+C)$  can be informally called “category

generalization,” but this name is ambiguous because the same name was previously used for problem  $Gn(O, c)$  in Fig. 3b. The labels for these two problems, however, indicate how they are different. Problem  $Gn(O, c)$  is a case in which the category of interest is treated as a class, and the task is to decide which objects belong to this class. For example, given a collection of specimens in a museum, a biology student may need to decide which specimens are dodos. Problem  $Gn(c, F+C)$  is a case in which the category of interest is treated as a feature-bearer, and the task is to decide which features apply to this category and which superordinate categories it belongs to. For example, upon learning that dodos have beaks, a child may be able to infer that dodos have wings and that dodos are birds.

Problem  $Gn(c, F+C)$  has received less attention than problem  $Gn(o, F+C)$ , but it has been addressed by at least one study. Kemp (2011) considered a version of the problem in which participants learn two features of a novel category (e.g., “wugs fly” and “wugs have no legs”) and must decide which other features the category is likely to have (e.g., wugs probably have wings).

*Simultaneous object, feature, and category generalization.* The generalization problems discussed so far have involved inferences about a single element of a semantic system. Most psychological studies of generalization have considered problems of this form, but generalization problems can also involve uncertainty about multiple elements of a semantic system. Figure 3d shows one such problem that can be visualized as a matrix completion problem in which the unobserved entries are scattered across the entire matrix. We refer to this problem as  $Gn(O, F+C)$ , where the uppercase  $O$ ,  $F$ , and  $C$  indicate that the problem does not focus on inferences about a single row or column of the matrix.

Problems like  $Gn(O, F+C)$  arise in many real-world settings. Consider, for example, two parents who are deciding where to send their child for college. Their knowledge about the choices available can be captured using a matrix in which the rows are colleges and the columns capture features including tuition price, climate, academic reputation, sporting prowess, and so on. The parents only know some of the entries in this matrix: For example, they may know that College A is expensive but have no idea whether its sports program is strong, and they may know that College B has a strong football team but have no idea what tuition at the college costs. In order to decide which colleges to investigate more closely, the parents will need to make inferences about unobserved values of the features that matter most to them. For example, they may infer that College C is likely to be expensive because it has a strong academic reputation, and because it is similar to College A, which they know to be expensive.

Although  $Gn(O, F+C)$  is a common real-world problem, it is discussed relatively rarely in the psychological literature. One recent study explored a version of the problem in which

participants were given sparsely observed object–feature matrices like the example in Fig. 4a and asked to fill in the gaps (Kemp et al., 2012). Before completing this task, participants were told about causal relationships between the features—for example, feature  $f_1$  tends to cause feature  $f_2$ , which tends to cause feature  $f_3$ . The results suggested that participants were able to generalize across both objects and features. For example, after learning that the mouse had  $f_1$ , Fig. 4a shows that participants were relatively confident that the mouse had  $f_2$  (object generalization) and that the rat had  $f_1$  (feature generalization). Participants were also able to make inferences that relied on generalization across objects and features—for example, their inferences that the rat had  $f_2$  were above baseline.

Problem  $Gn(C, F+C)$  in Fig. 3h is a version of the matrix completion problem in Fig. 3d in which the rows of the matrix represent categories rather than objects. The study shown in Fig. 4a asked participants to reason about four individual animals—a mouse, a rat, a squirrel, and a sheep—but similar studies could be conducted that ask for inferences about four categories (mice, rats, squirrels, and sheep).

## Discovery


All of the generalization problems in Fig. 3a–h were created by concealing some of the entries in a semantic system, and in each case the observed entries included information about each row and each column in the matrix. In general, however, there may be no observations for a given row in the matrix, and a reasoner may have no direct evidence that the corresponding object exists. Similarly, there may be no observations for a given column, and the reasoner may have no direct evidence that the corresponding feature or category exists. We will refer to problems in which the reasoner must infer the existence of unobserved elements as *discovery* problems. Figures 3i–l show four examples in which the objects, features, and categories to be discovered are shown as gray columns or rows.

*Feature discovery.* Feature discovery is a problem in which a reasoner is given an initial set of objects and features and subsequently constructs or infers the existence of a new feature that did not belong to the initial set. Figure 3i shows an example in which the new feature corresponds to the final column in the object–feature matrix. The label of this problem is  $Ds(O, f)$ . The  $Ds$  indicates that the problem is a discovery problem, the  $O$  and  $f$  indicate that the problem involves a matrix of objects by features, and the lowercase  $f$  indicates that the problem requires the reasoner to discover a single feature.


Feature discovery has been extensively discussed by psychologists, who sometimes refer to the problem as “feature learning” or “feature construction” (Austerweil & Griffiths, 2009; Landy & Goldstone, 2005; Schyns, Goldstone, & Thibaut, 1998; Wisniewski & Medin, 1994), and by machine-learning researchers, who often refer to the problem as




**a** Gn( $O, F$ ): Simultaneous object and feature generalization



**b** Ds( $O, F$ ): Simultaneous object and feature discovery



**c** Id( $C_T, F_T$ ): Simultaneous category and feature identification



**Fig. 4** Novel problems that were identified while developing the taxonomy in Fig. 3. Each problem combines two of the basic problems in Fig. 3. **(a)** (i) A generalization problem in which participants observe one entry in an object–feature matrix and then make inferences about all remaining entries in the matrix. (ii) Human inferences are sensitive to both the taxonomic relationships between the objects (e.g., the rat is judged more likely to have  $f_1$  than the sheep) and the causal relationships between the features (e.g., the judged probability that the mouse has  $f_2$  is above baseline). **(b)** (i) A discovery problem in which

“constructive induction” (Rendell & Seshu, 1990). A typical study of feature learning might use a task in which participants view objects that are constructed from hundreds of oriented line segments. By tracking which of these segments tend to appear together, participants might discover that each object is built from several parts that each correspond to a constellation of line segments. These parts can be described as features, and discovering the existence of these parts therefore qualifies as an instance of feature discovery. In other settings, feature discovery may rely more heavily on rich background knowledge than on observations of low-level features such as line segments (Wisniewski & Medin, 1994). In general, feature discovery can therefore be characterized as a problem that requires preexisting features to be combined with background knowledge in order to discover new features.

Problem Ds( $C, f$ ) in Fig. 3m is a version of the problem in Fig. 3i in which the rows of the matrix represent categories

participants observe eight objects and then must draw a ninth object that has not been observed. Each object is created by combining a horizontal piece with a vertical piece, and the gray cells in the design matrix represent the eight combinations observed during training. (ii) Participants tend to generate new objects by combining horizontal and vertical pieces that have not previously been combined. **(c)** (i) An identification problem in which participants must infer the identities of feature F and animal A. (ii) The most common response is that F is “white” and that A is “polar bear”

rather than features. Given information about several categories, a reasoner may be able to construct new features by combining existing features. For example, a reasoner who hears about several animal species that “come in many colors” and “come in many sizes” might combine these features to create a new feature that might be glossed as “variable in appearance.” To our knowledge, however, problem Ds( $C, f$ ) has not been addressed in the psychological literature.

*Category discovery.* Figure 3j shows a problem in which the column in gray represents an unobserved category that must be discovered. The new category need not receive a verbal label—for example, an experimental participant might construct two new categories by sorting objects into two piles and thinking of them as “the ones over here” and “the ones over there,” without ever assigning a verbal label to either group.

The problem of category discovery goes by several names, including “concept formation” (Bruner et al., 1956), “category construction” (Ahn & Medin, 1992; Medin, Wattenmaker, & Hampson, 1987), and “clustering” (Fisher, 1987). It is also known as the problem of “unsupervised categorization”—unsupervised, because none of the category labels is provided by a supervisor or teacher (Pothos & Chater, 2002). For example, the first European explorers to visit Australia were able to organize the animals that they saw into categories without needing a supervisor to provide category labels. Unsupervised categorization has received less attention than supervised categorization, but has nevertheless been extensively studied (Anderson, 1991; Inhelder & Piaget, 1964; Love, 2002; Pothos et al., 2011).

Figure 3n is a problem in which a reasoner begins with an existing set of categories and constructs a new higher-order category that includes some of these categories. For example, a European explorer might organize some of the categories that he discovered into the higher-order category of “marsupials.” Inferences of this kind have been studied using pile-sorting tasks in which the items to be sorted are the names of categories. For example, Berlin, Breedlove, and Raven (1974) gave people a set of plant categories and asked them to organize these categories into higher-order categories.

*Object discovery and category discovery.* Object discovery is a problem in which all relevant features and categories are known but a reasoner must infer the existence of one or more unobserved objects. Figure 3k shows an example in which the existence of the object that occupies the final row in the matrix must be inferred. For example, before the planet Neptune was directly observed, Le Verrier inferred the existence of this object on the basis of perturbations in the orbit of Uranus. Le Verrier’s discovery can be viewed as an inductive inference about a partially observed matrix in which the rows represent heavenly bodies and the columns represent features of these objects, including their masses and positions. There were at least two possible explanations of the available data: Perhaps an unobserved planet was affecting the orbit of Uranus, or perhaps Newton’s laws did not apply as expected (Leverington, 2003). The first explanation was correct in the case of Neptune, but a subsequent episode demonstrated that the second explanation was a genuine contender. On the basis of perturbations in the orbit of Mercury, Le Verrier inferred the existence of a planet, Vulcan, that lay between Mercury and the sun. Le Verrier’s conclusion was plausible but wrong, and the anomalies in Mercury’s orbit were eventually explained when Newton’s laws were superseded by Einstein’s general theory of relativity.

Object discovery has received relatively little attention in the psychological literature, but some simple versions of the problem have been explored empirically (Csibra & Volein, 2008; Kemp, Jern, & Xu, 2009). In one study that was inspired by the discovery of Neptune, participants were

asked to make inferences about unobserved charged particles that determined the trajectory of an observed particle (Carroll & Kemp, 2012). Inferences of this kind can be viewed as inductive inferences about a partially observed matrix in which the rows represent particles and the columns represent features of the particles, including their charges and positions. In a second line of work, developmental researchers have explored the conditions that lead infants to infer the existence of a hidden object (Csibra & Volein, 2008; Saxe, Tenenbaum, & Carey, 2005). For example, Csibra and Volein provided evidence that infants are able to infer the existence of a hidden object by following the gaze of another person. The studies just described explore some of the simplest possible versions of the problem of object discovery, but it is likely that other versions of the problem would repay empirical investigation.

Figure 3o shows a category discovery problem that is a version of the matrix completion problem in Fig. 3k in which the rows of the matrix represent categories rather than objects. Mendeleev solved a version of this problem when he used his periodic table to infer the existence of elements that had not yet been observed. Mendeleev’s discoveries can be viewed as inductive inferences about a matrix in which the rows represent elements and the columns represent features of these elements, including their atomic weights, specific heats, and melting points. Although Mendeleev achieved some striking successes, he also made predictions about several elements that do not exist in reality, including two elements that were proposed to be lighter than hydrogen. As for Le Verrier, Mendeleev’s successful and not-so-successful predictions can both be viewed as inductive inferences, or inferences to conclusions that were plausible but not certain, given the available data.

Computational models of human learning are typically not designed to address the object and category discovery problems in Fig. 3k and o. These problems, however, have been addressed by several computational models of scientific discovery (Langley, Simon, Bradshaw, & Zytkow, 1987; Valdés-Peréz, Żytkow, & Simon, 1993). For example, the MENDEL system has been used to model the discovery of genes, and the GELL-MANN system has been used to model the discovery of subatomic particles such as quarks (Fischer & Zytkow, 1992). Researchers have argued that the inferences made by some of these models are consistent with historical evidence about the human discoveries that inspired these approaches, but as yet few empirical studies have evaluated these models as psychological proposals.

*Simultaneous object, feature, and category discovery.* The discovery problems in Fig. 3i–k focus on a single feature, object, or category, but in general a reasoner may need to simultaneously discover multiple objects, features, and categories. Figure 3l shows an instance of this problem in which a reasoner must discover an unobserved object, an unobserved feature, and an unobserved category. The problem in Fig. 3p is

similar, except the rows of the matrix in Fig. 31 now correspond to categories rather than objects.

Because object discovery is discussed only rarely in the literature, it is not surprising that problems that include object discovery as a component have received almost no attention. Problems of this kind, however, are sometimes encountered in the real world. For example, given the genomes of several present-day species, scientists may discover new features that correspond to genetic sequences that are conserved across species, and they may use these features to infer the existence of ancestral species that have never directly been observed. Jern and Kemp (2013) recently developed a study that explored a very simple version of this problem. Participants in the study were shown objects like the example in Fig. 4b. These objects were described as the genomes for different kinds of flu viruses, and each one was constructed by combining a horizontal pair of letters (e.g.,  $QJ$  in the second row in the figure) with a vertical pair (e.g.,  $DB$  in the first column). Participants were not informed that the individual letters were paired in this way, and recognizing that the pairs existed was an instance of feature discovery. After studying the observed genomes, participants were asked to draw a genome for a virus that had not yet been observed but was likely to exist. The results in Fig. 4b show that the majority of participants combined a known horizontal pair with a known vertical pair to create a new genome, suggesting that they were able to use the features that they had discovered to make inferences about unobserved objects. Because participants were specifically asked to draw an unobserved genome, the experiment did not address the problem of object discovery in its purest form. We expect, however, that similar results would be achieved if participants were first asked to decide whether unobserved genomes were likely to exist, and were only asked to draw such a genome if they answered in the affirmative.

## Identification

Imagine that the objects in a set are repeatedly encountered, and that every time a reasoner encounters an object, he or she observes a feature of the object that he or she had not previously noticed. So far we have assumed that each observation of this kind allows the reasoner to fill in one entry in a matrix of objects by features. In general, however, a reasoner may encounter an object multiple times without identifying it as the same object each time. In other words, a reasoner may encounter multiple object *tokens* without knowing that all of these tokens are instances of the same object. As in the previous sections, the information available to the reasoner can still be viewed as incomplete information about a matrix of objects by features. Allowing for uncertainty about token identity, however, introduces a new set of inductive problems.

All of the generalization problems discussed previously have counterparts involving tokens. For example,  $Gn(O_T, f)$

is a version of the generalization problem in Fig. 3a in which the rows of the matrix represent object tokens rather than objects. Although the problems in Fig. 3b–h also have counterparts involving tokens, discussing each problem in detail would add relatively little to our previous discussion of generalization. Instead, we will focus on the inductive problem of identification, which is qualitatively different from the problems previously considered in this article.

Identification problems arise when a reasoner observes object tokens, feature tokens, or category tokens and must infer which tokens are instances of the same entity. For simplicity, we begin by considering problems in which a single token is observed, and a reasoner must decide whether this token is an instance of a previously observed object, feature, or category.

*Feature identification.* Feature identification is a problem in which a reasoner must infer the identity of a feature token after observing that the token applies to a certain set of objects. For example, if told that Sarah Ferguson, Julianne Moore, and Conan O'Brien all have feature  $F$ , a reasoner might guess that feature  $F$  is “red hair.” We label this problem as  $Id(O, f_T)$ , in which  $Id$  is short for *identification*, the  $O$  indicates that the problem concerns a matrix in which the rows represent objects, and the lowercase  $f_T$  indicates that the problem focuses on an inference about a single feature token. Figure 3q shows an example in which the double-headed arrow indicates that a reasoner must decide whether feature token  $f_T$  is an instance of feature  $f_6$ . The matrix in Fig. 3q is fully observed, but in general the matrix may contain gray elements that represent entries that have not been observed.

Problem  $Id(C, f_T)$  in Fig. 3u is a version of the identification problem in Fig. 3q in which the rows of the matrix represent categories rather than objects. Kemp, Chang, et al. (2010) studied this problem using a task in which participants were given several categories that have a certain feature (e.g., swans, polar bears, and doves all have feature  $F$ ) and were then asked to guess the feature in question ( $F$  might be a token of the feature “white”). Feature identification has also been studied using the “word context task” (Werner & Kaplan, 1952), which requires participants to infer the meaning of a nonsense word after hearing it used in context. For example, after hearing that “elephants are big and zazy,” a child may be able to infer that “zazy” is another word for “strong” (Braun-Lamesch, 1962).

As the word context task suggests, feature identification is a problem that regularly arises when learning new words. Any student learning a second language will frequently encounter new labels for familiar features. For example, a student learning German may need to infer that *rot* and *schwarz* correspond to features that would be labeled “red” and “black” in English. In cases of this kind, *rot* and *schwarz* can be viewed as tokens of features that the reasoner already knows. The same general phenomenon arises during first language learning, when

learners are exposed to linguistic labels for features that they have already noticed. For example, a child may have noticed that tomatoes, radishes, and cherries all share a certain feature. When she later hears the word “red,” she may be able to recognize that this novel label is a token of a feature that she already knows. Many researchers have proposed that word learning often involves mapping linguistic labels onto preexisting concepts (Bloom, 2000; Fodor, 1975; MacNamara, 1972; Merriman, Schuster, & Hager, 1991; Mervis, 1987; Snedeker & Gleitman, 2004), and in our taxonomy this mapping problem corresponds to a problem of identification.

*Category identification.* The category identification problems in Fig. 3r and v are similar to the feature identification problems just discussed, but they involve inferences about category tokens rather than feature tokens. For example, an English speaker who sees a German friend point at a colorful insect and call it a *Schmetterling* might infer that *Schmetterling* is a label for the category of butterflies. Similarly, if the friend mentions that kangaroos, koalas, and wombats are all *Beuteltiere*, the English speaker might infer that *Beuteltiere* is a label for the category of marsupials.

Category identification has been discussed by previous researchers, and it is sometimes called “concept identification” (Hunt, 1962) or “category recognition” (Maratsos, 2001). Experiments in which participants learn the meanings of novel words are often viewed as studies of category generalization (Fig. 3b), but some of these studies are perhaps better viewed as cases in which participants learn new labels for preexisting categories (Bloom, 2000; Maratsos, 2001). For example, researchers have documented cases of “fast mapping” in which children learn novel words after hearing them used on a single occasion (Carey & Bartlett, 1978; Heibeck & Markman, 1987). It is difficult to understand how a child could acquire a novel category given a single exemplar, but easier to understand how a single example could allow a child to map a novel label onto a preexisting category.

Although category identification may play a role in many experimental studies of word learning, we are aware of few studies that have been explicitly designed as studies of category identification. One relevant task is based on a television game show called *Pyramid*. Contestants in this show are given several exemplars of a category (e.g., cymbals, glockenspiel, and timpani) and then asked to guess the category from which the exemplars are drawn (e.g., percussion instruments). Guessing games of this kind have been used in classrooms as vocabulary-building exercises (Marzano & Pickering, 2005), but to our knowledge have not yet been used in studies of inductive reasoning.

*Object identification and category identification.* Object identification is a problem in which a reasoner must infer that an object token is an instance of a previously observed

object. Figure 3s shows an example in which the object token appears as the final row of the matrix. A well-known historical example of object identification is the inference that the morning star and the evening star are the same object. In terms of our framework, the relevant inductive problem can be characterized using a matrix like the one in Fig. 3s, in which one row represents the bright star that is observed before sunrise, another row represents the bright star that is observed after sunset, and the columns represent the features (e.g., brightness and position) of these stars. The Greeks originally used data of this kind to infer that the two stars were different objects that they called Phosphorus and Hesperus. Later, however, they came to believe that Phosphorus and Hesperus were one and the same.

Object identification has been discussed in detail in the psychological literature. One family of studies refers to the problem as “absolute judgment” (Miller, 1956) or “absolute identification” (Luce, Green, & Weber, 1976), and is based on objects that are simple perceptual stimuli such as lines or tones (Brown, Marley, Donkin, & Heathcote, 2008; Estes, 1994; Nosofsky, 1986). A typical experiment might use  $n$  lines of different lengths, each of which is associated with an identifying label. On each trial, participants observe a token of one of the lines and are required to provide the identifying label for the token. In a study of this kind, uncertainty typically arises because of perceptual noise and memory failures—for example, because a participant cannot accurately detect and remember the length of a line. Uncertainty can also arise in the absence of perceptual noise if some features of the object tokens are unobserved, and absolute identification has also been studied in this setting (Kemp et al., 2009).

A second family of studies has focused on visual object perception. For example, participants might observe one object token (e.g., a square) moving behind an occluder and a second object token (e.g., a rectangle) emerging from the other side (Burke, 1952). Depending on the perceptual features of the two objects, participants may infer that the two tokens are glimpses of the same object or glimpses of two different objects. Similar studies have been carried out with infants (Bower, 1974; Spelke, Kestenbaum, Simons, & Wein, 1995; Van de Walle, Carey, & Prevor, 2000; Xu, 2005) and nonhuman animals (Mendes, Rakoczy, & Call, 2008). The literature on this topic often refers to the problem of “object individuation” rather than “object identification.” Some researchers have distinguished between individuation and identification: For example, Tremoulet, Leslie, and Hall (2000) proposed that individuation involves setting up an object representation, and that identification involves deciding “which, if any, previously individuated object is presently encountered” (p. 499). Most of the literature on individuation, however, is relevant to the problem that we have called *identification*.

Instead of discerning the identity of an object token, reasoners may simply be asked to indicate whether or not they

have encountered the object before. The resulting problem is often called *recognition*. Although recognition does not appear in our taxonomy, it is closely related to the problem of identification, and the two problems are often considered together (Estes, 1994). Several studies have also established connections between recognition and the problem that we have called *generalization* (Hayes, Fritz, & Heit, 2013; Heit, Rotello, & Hayes, 2012; Sloutsky & Fisher, 2004).

Figure 3w shows a category identification problem that corresponds to a version of the problem in Fig. 3s in which the rows of the matrix represent categories rather than objects. This problem has been studied in the literature on “naming to description” (Lambon Ralph, Graham, Ellis, & Hodges, 1998; Lombardi & Sartori, 2007; Sartori & Lombardi 2004). Experiments in this literature are similar to guessing games. For example, a participant might be told that a certain kind of animal has whiskers and catches mice, and might guess that the animal is a cat. From our perspective, solving this problem requires the inference that the category mentioned in the task is a token of the category of cats.

Category identification has also been studied using the “word context task” described in a previous section (Werner & Kaplan, 1952). For example, after hearing statements that include “the painter used a corplum to mix his paints” and “you can make a corplum smooth with sandpaper,” a reasoner might be able to identify “corplum” as a label for the category of sticks (Werner & Kaplan, 1952). Because “corplum” is a novel label for a familiar concept, problems of this kind are sometimes called “concept identification” problems (Bruner et al., 1956).

*Simultaneous object, feature, and category identification.* The identification problems in Fig. 3q–s require inferences about a single object token, feature token, or category token. In general, however, identification problems may require reasoners to reason about multiple tokens of different kinds. Figure 3t shows an example involving an object token  $o_T$ , a feature token  $f_T$ , and a category token  $c_T$ , and the problem in Fig. 3x is similar, except that the rows of the matrix represent categories rather than objects.

Kemp, Chang, et al. (2010) recently explored a version of the problem in Fig. 3x in which the categories were animal categories and the features were perceptual, behavioral, and anatomical features. Participants were given trios of statements like the example in Fig. 4c. Each trio included a token of an unidentified category (e.g., “animal A”) and a token of an unidentified feature (e.g., “feature F”), and participants were required to identify the category and the feature. For example, the most common response to the problem in Fig. 4c indicated that feature F was “white” and that animal A was a “polar bear.” The task in Fig. 4c can be viewed as a simplified version of a real-world problem faced by language learners (Kemp, Chang, et al., 2010). For example, suppose that a student watching a German nature program hears that a *Schmetterling*

is *bunt*, and also hears *Schmetterling* and *bunt* used in several other contexts. Combining all of this information may allow the student to identify *Schmetterling* as a label for the category “butterfly” and *bunt* as a label for the feature “colorful.”

#### Combining generalization, discovery, and identification

The previous sections illustrated how problems that focus on features, problems that focus on categories, and problems that focus on objects can be combined to create problems that require inferences about features, categories, and objects. For example, Fig. 3d shows how feature generalization, category generalization, and object generalization can be combined. The combined problems in the rightmost column of Fig. 3 all include multiple versions of the same basic problem: for example, multiple generalization problems, multiple discovery problems, or multiple identification problems. The three basic problems, however, can also be combined to generate additional problems.

Discovery and generalization are combined in settings in which a reasoner must discover new objects and make inferences about their properties. For example, Mendeleev was able not only to infer the existence of novel elements, but to make predictions about unobserved properties of these elements. Discovery and identification are combined in settings in which reasoners encounter object tokens, and in which some of these tokens are instances of familiar objects but others are instances of novel objects. For example, if you are shown a series of family photographs you might identify the first individual as your grandfather but decide that you have never previously seen the second individual. Finally, identification and generalization are combined in settings in which reasoners are asked to make inferences about unobserved properties of confusable objects. For example, suppose that you know that Tim has diabetes, and you now encounter a person who is either Tim or his identical twin Tom. You might infer that this person probably has diabetes either because this person is Tim (identification) or because this person is similar to Tim who has diabetes (generalization).

The compound problems just considered are all cases in which multiple inductive problems arise and in which these problems are coupled in some way. For example, the diabetes example couples identification and generalization because your final conclusion depends on both identification (the person might be Tim, the known diabetic) and generalization (even if the person is Tim’s twin Tom, there is still a good chance that he has diabetes). In other cases, multiple inductive problems arise, but these problems are not coupled in any important way. For example, if shown a photograph of a person wearing glasses, you might recognize the person as your great-grandfather (identification) and might infer that the person is probably long-sighted (generalization). Here, however, the identification judgment and the generalization

need not be combined in any way, and we therefore view this example as a case involving two distinct problems instead of as an instance of a compound problem.

Recognizing that inductive problems can incorporate two or more basic problems can help to explain experimental results that would otherwise be puzzling. Consider a feature generalization problem (Fig. 3a) in which a reasoner observes that a source object has a certain feature and is then asked to decide how likely it is that a target object has the feature. Many studies of this kind have measured generalization curves that indicate how generalizations decay as a function of the dissimilarity between the source object and a target object. Some researchers disagree about whether these generalization gradients are closer to Gaussian functions or exponential functions (Nosofsky, 1986; Shepard, 1987). As we mentioned earlier, one proposed resolution is that pure feature generalization problems produce exponential curves, but that inferences about highly confusable objects such as Tim and Tom involve an identification component, and therefore produce near-Gaussian generalization curves (Ennis, 1988; Nosofsky, 1988; Shepard, 1986). The literature on this topic suggests that some inductive tasks require two or more basic problems to be solved, and that it is important to think clearly about the inductive problems posed by any given task.

#### Beyond object–feature matrices

Our taxonomy is organized around the three problems of generalization, discovery, and identification. We introduced these problems using semantic systems constructed from objects, features, and categories, and Fig. 3 is based on simple systems of this kind. Figure 3, however, shows only part of our taxonomy. The full taxonomy includes all problems of generalization, discovery, and identification that arise with respect to any semantic system, including systems that incorporate relations and systems that include subsystems as components. The next two sections provide examples of generalization, discovery, and identification problems that involve relations and higher-level semantic systems.

*Problems involving relations.* The semantic systems in Fig. 3 incorporate two special kinds of relations. The first relation holds between feature-bearers and features, and might be called the *has* relation. For example, in Fig. 3a object  $o_1$  *has*  $f_3$ . The second relation holds between feature-bearers and categories, and might be called the *is a* relation. For example, in Fig. 3a object  $o_1$  *is a*  $c_2$ .

In addition to these two special relations, semantic systems may include relations that hold between feature-bearers and other feature-bearers. Consider first a system that includes a set of objects, a set of binary features defined over these objects, and a set of binary relations defined over these objects. As shown in Fig. 2b, the system can be visualized by

adding several square matrices of objects by objects to a matrix of objects by features. For example, the objects might be a set of individuals, and the relations might include *taller*( $\cdot$ ;  $\cdot$ ), *older*( $\cdot$ ;  $\cdot$ ), *sister of*( $\cdot$ ;  $\cdot$ ), and *friends with*( $\cdot$ ;  $\cdot$ ), among other examples.

Generalization problems arise when some entries in the square relational matrices are unobserved. One simple case occurs when the value of a given relation is observed only for some pairs of objects, and a reasoner must decide whether or not the relation applies to the remaining pairs of objects (Kemp, Tenenbaum, Niyogi, & Griffiths, 2010). For example, a reasoner might be told that Alice is the sister of Betsy, and that Chloe is the sister of Betsy, and might then have to infer whether or not Alice is the sister of Chloe. Four-term analogy problems provide a second example of relational generalization (Rumelhart & Abrahamson, 1973; Sternberg & Gardner, 1983). For example, a reasoner might be given the problem Alice:Betsy :: Daphne:?, which can be glossed as “Alice is to Betsy as Daphne is to whom?” This problem is an example in which a reasoner learns that some relation  $R(\cdot, \cdot)$  applies to the pair (Alice, Betsy), and the reasoner needs to find some person  $X$  such that  $R(\text{Daphne}, X)$ . The examples so far have focused on generalization problems involving a single relation and multiple pairs of objects, but other generalization problems might require a reasoner to make inferences about multiple relations that apply to a single pair of objects. For example, a reasoner might learn that Alice is younger than Betsy, and then might have to decide whether a different relation holds between the pair (e.g., is Alice smaller than Betsy?).

The problem of discovery can also apply to systems involving relations. For example, Dalton’s discovery of atoms was enabled by relational data that captured the outcome of chemical reactions. A discovery problem may also require a reasoner to infer the existence of unobserved relations. For example, if Alice and Bob both simultaneously come down with a very rare illness, we might infer that the two have recently come into contact in some way. Discovering this contact relation is an instance of the problem of “relation discovery.”

Finally, the problem of identification also applies to systems involving relations. Consider, for example, a problem in which an individual is learning English as a second language and is told that “Alice dislikes Betsy” and that “Chloe admires Daphne.” Even if the learner has not previously encountered the words “dislikes” and “admires,” he or she may be able to identify these words as labels for relations that he or she already knows. Problems of this kind qualify as instances of “relation identification.”

*Problems involving higher-level systems.* As we described earlier, semantic systems can be combined in order to create a higher-level system that includes features, categories, or

relations defined over the component systems. The problems of generalization, discovery, and identification can all arise when reasoning about these higher-level systems. This section describes examples of each problem in which the component systems are molecules, and the higher-level system is a set of molecules along with features of and relations between these molecules. For example, one of the features might indicate whether a molecule relieves pain, and one of the relations might pick out pairs of molecules that react with each other.

Generalization problems arise when some of the features that apply to the systems are partially observed. For example, suppose that a reasoner observes that one of the molecules in the set relieves pain. Predicting whether or not the other molecules are likely to relieve pain is an example of generalizing across semantic systems. Problems involving generalization across features can also arise. For example, a reasoner who learns that a certain molecule stimulates mu-receptors might infer that the molecule relieves pain.

Generalization problems involving multiple systems of objects, features, and relations have been explored extensively in the literature on relational categorization and analogical reasoning (Gentner, 1983; Gentner & Kurtz, 2005; Holyoak, Lee, & Lu, 2010; Holyoak & Thagard, 1989). A classic inductive problem considered in this literature requires a reasoner to use what is known about one relational system (e.g., the solar system) to make inferences about unobserved aspects of a second system (e.g., an atom). Analogical inferences of this kind rely on the insight that the two systems resemble each other in some respect. In terms of our framework, this resemblance can be formalized using the idea that both systems belong to a higher-level category (Christie & Gentner, 2010; Gick & Holyoak, 1983; Kemp & Jern, 2009). Recognizing the existence of this higher-level category provides the basis for generalizations about the objects, features, and relations within each individual system.

Discovery problems arise when reasoning about a higher-level system that includes lower-level systems or features of these lower-level systems that were not initially observed. For example, a reasoner might solve the problem of system discovery by inferring the existence of new kinds of molecules, such as buckminsterfullerene. Similarly, a reasoner might notice that certain molecules contain a ring of carbon atoms, and thereby discover a feature that was not part of the original description of the system.

Identification problems arise when a reasoner encounters tokens of systems and must decide which tokens are instances of the same molecule. For example, suppose that a chemist has created a microscopic pen that includes two molecules of methane: One molecule includes a carbon-12 atom, and the other includes a carbon-13 atom. If the reasoner has access to a device that allows individual molecules

to be viewed, he or she may be required to identify the molecule currently being viewed as either the carbon-12 or the carbon-13 molecule. Problems of this kind qualify as instances of “system identification.”

For simplicity, this section has focused on problems involving multiple low-level systems (e.g., multiple molecules) in which each object (e.g., each atom) belongs to a single low-level system. We can also formulate problems in which a given object can belong to multiple low-level systems. For example, when thinking about animals and their properties, it may be useful to distinguish between a taxonomic system that includes categories such as “mammal” and “reptile” and an ecological system that includes categories such as “predator” and “sea creature” (Heit & Rubinstein, 1994; López, Atran, Coley, Medin, & Smith, 1997; Shafto, Kemp, Mansinghka, & Tenenbaum, 2011). Any given animal will appear in both systems, and some inductive problems draw on both of these systems. For example, a reasoner might combine taxonomic knowledge and ecological knowledge to make the plausible inference that dolphins have especially large lungs. Our taxonomy allows for high-level semantic systems that include overlapping or cross-cutting low-level systems, but inferences about these high-level systems do not appear to introduce problems that are qualitatively different from the inductive problems already discussed.

### Unified accounts of inductive reasoning

An important motivation for developing a taxonomy of inductive reasoning is to characterize the space of inferences that psychological theories should aim to explain. Our taxonomy is intended to be relatively theory-neutral, and we hope that researchers of different theoretical persuasions will agree that accounting for the full space of inferences is a worthy challenge for theories of human reasoning. For example, researchers who work with probabilistic models, exemplar models, and connectionist models may be able to agree that the phenomena in the taxonomy need to be explained, even if they disagree about the most promising way to develop a unified theory of inductive reasoning.

Our taxonomy highlights three challenges that must be addressed by any unified theory of inductive reasoning. We will refer to these challenges as the *representation challenge*, the *knowledge challenge*, and the *inference challenge*. The taxonomy is based on inferences about systems of objects, features, categories, and relations, and these systems can take many forms. For example, we argued earlier that semantic systems can include higher-level features, categories, and relations defined over lower-level systems, and this process of introducing higher-level features, categories, and relations can continue indefinitely. In order to address the representation

challenge, a unified account of inductive reasoning must be able to represent the full collection of possible semantic systems.

The taxonomy characterizes inductive problems without specifying the knowledge required to solve these problems. Inductive reasoning, however, depends critically on background knowledge, and for the problems in our taxonomy, the relevant background knowledge is knowledge about systems of objects, features, categories, and relations. This knowledge can take many forms, and may include causal knowledge (e.g., large animals tend to be heavy), associative knowledge (animals with hooves tend to eat plants), taxonomic knowledge (bats are mammals), knowledge about similarity (horses and zebras are similar), and intuitive or scientific theories (e.g., theories about how parents pass on properties to their offspring). In order to address the knowledge challenge, a unified account of inductive reasoning must be able to capture the many different kinds of background knowledge that people bring to inductive problems.

The taxonomy aims to include all problems in which a reasoner is given an incomplete specification of a semantic system and is asked to infer a complete specification of the system. Given any fully specified system of objects, features, categories, and relations, there are many ways to create incomplete specifications by concealing different aspects of the system. In order to address the inference challenge, a unified account of inductive reasoning must be able to handle queries about unobserved aspects of any incompletely specified system.

No existing computational model accounts for all of the problems in our taxonomy, but probabilistic models, exemplar models, and connectionist models have all been able to account for human inferences about multiple inductive problems. This section briefly considers all three approaches, and argues that the probabilistic approach offers the most direct route toward a unified account that addresses the three challenges just described.

Probabilistic models have been applied to many of the problems in Fig. 3, including generalization (Heit, 1998; Kemp et al., 2012; Shepard, 1987), discovery (J. R. Anderson, 1991; Austerweil & Griffiths, 2011; Kemp et al., 2009), and identification (Kemp, Chang, et al., 2010; Kemp et al., 2009). These models take many forms, but all of them specify a prior distribution over a space of semantic systems and update this distribution using probabilistic inference when evidence is observed. The probabilistic approach is compatible with many different representational proposals, including the idea that semantic systems are mentally represented in a compositional representation language such as predicate logic. Compositionality is the classic solution to the representation challenge: A compositional representation language includes symbols for objects, features, categories, and relations, and a vast number of semantic systems can be represented by

combining these symbols in different ways. The probabilistic approach highlights the role of background knowledge, and the prior distribution required by a probabilistic model provides a flexible way to capture many kinds of background knowledge. For example, researchers have worked with priors that capture causal knowledge (Glymour, 2001; Holyoak & Cheng, 2011; Sloman, 2005), associative knowledge (Y. Xu & Kemp, 2010), taxonomic knowledge (Kemp & Tenenbaum, 2009), knowledge about similarity (Shepard, 1987), and theories, both intuitive (Tenenbaum, Griffiths, & Kemp, 2006) and scientific (Howson & Urbach, 1993). Finally, probabilistic inference provides a solution to the inference challenge. Given a prior distribution over a space of semantic systems and a set of assumptions about how observations are generated, general-purpose probabilistic inference can be used to reason about any unobserved aspect of a given system. The probabilistic approach is therefore capable of addressing all three of the challenges previously identified, and it offers a promising path toward a unified account of inductive reasoning.

Exemplar- or instance-based models have also been used to account for multiple inductive problems, including generalization and identification (Estes, 1994; Nosofsky, 1992). The key idea that motivates the exemplar-based approach is that inferences about novel objects, features, categories, relations, or semantic systems are based on similar, previously encountered objects, features, categories, relations, or systems. An influential version of the exemplar-based approach proposes that objects are represented as points in a continuous similarity space (Nosofsky, 1986), but exemplar models can also be defined over representations constructed in a compositional representation language (Aamodt & Plaza, 1994). As a result, the exemplar-based approach can address the representational challenge in just the same way as the probabilistic approach. Compared with the probabilistic approach, however, the exemplar-based approach offers a less complete solution to the knowledge challenge. The background knowledge used by an exemplar model is carried by previously observed objects, features, categories, relations, or systems, and specific observations of this kind seem incapable of capturing the abstract knowledge that shapes some inductive inferences. For example, Newtonian mechanics is a classic example of a body of abstract knowledge, and this knowledge has supported inferences about generalization problems (Newton inferred the mass of the sun), discovery problems (Le Verrier inferred the existence of Neptune), and identification problems (Halley inferred that three comets observed in 1531, 1607, and 1682 were actually the same object). It is far from clear how an exemplar-based approach might account for inferences of this kind, as well as for everyday inferences that rely on intuitive rather than scientific theories (Carey, 1985; Gopnik & Meltzoff, 1997). Finally, the exemplar-based approach is



capable of addressing the inference challenge. Given a semantic system with several hidden elements, comparisons with previously encountered systems provide a general-purpose method for filling in the hidden elements.

The connectionist approach provides a third computational perspective on inductive reasoning, and connectionist models have been used to account for multiple inductive problems, including generalization (Rogers & McClelland, 2004), identification (Lacouture & Marley, 1991), and relational learning (Doumas, Hummel, & Sandhofer, 2008). Connectionist models can take many different forms, but our discussion will focus on the work of Rogers and McClelland, who provided a comprehensive treatment of semantic cognition and inductive reasoning. Their connectionist theory proposes that semantic knowledge is carried by the connection weights in a network of simple units and that inductive inferences arise from a process in which each unit communicates with its neighbors. Although Rogers and McClelland showed that their network accounts for certain kinds of generalization problems, their approach falls short of a unified account of induction. Their approach does not address the representation challenge in full, because it is not clear how the distributed representations that they advocate can capture the full collection of possible semantic systems, including systems with higher-level features, categories, and relations defined over lower-level systems. Their approach also fails to address the knowledge challenge in full, because it is not clear how patterns of weights in a network can capture all of the kinds of knowledge that shape inductive reasoning, including scientific and intuitive theories. Finally, the specific models implemented by Rogers and McClelland fail to address the inference challenge. These models are feedforward networks in which the units representing features appear among the output units. As a result, these models do not handle inference problems in which some of the features of a novel object are observed and the task is to predict which other features the object might have. Rogers and McClelland acknowledged this limitation but suggested that the feedforward nature of their networks is a simplification, and that their underlying theoretical approach would be better captured by a recurrent network. Recurrent networks address the inference challenge, because these networks can make inferences about any subset of variables that happens to be unobserved. Overall, then, the connectionist approach is able to address the inference challenge, but it seems less capable of addressing the representation and knowledge challenges.

Although we have argued that the probabilistic approach provides the most promising path toward a unified theory of inductive reasoning, probabilistic inference per se can contribute only a small part of this theory. The greater part of the theory will need to be a set of principles that address the knowledge challenge and characterize the background knowledge that reasoners bring to inductive problems. As we

suggested earlier, this knowledge may take many forms which range from simple associative links to complex scientific theories, and it seems unlikely that a single formalism will be able to elegantly capture all of these varieties of background knowledge. As a result, aiming for a single, unified probabilistic model of inductive reasoning may be a mistake. A better research strategy may be to aim for a family of models that rely on a single inference strategy—namely, probabilistic inference—but that incorporate priors induced by a variety of different knowledge structures. For example, probabilistic models have been applied to many of the problems in Fig. 3, but these models have relied on problem-specific prior distributions, and it is hard to see how all of these distributions could be interpreted as special cases of a general-purpose method for specifying prior distributions.

Just as inductive inferences may draw on many qualitatively different kinds of background knowledge, inductive inferences may be carried out by multiple qualitatively different processes. For example, classical conditioning can be viewed as a kind of inductive learning, and the processes responsible for classical conditioning may be rather different from the processes that support scientific discovery (although see Holland et al., 1986). Statisticians and computer scientists have demonstrated that probabilistic inference can be implemented using many different processes, and it may turn out that all of the different psychological processes that support inductive reasoning can be usefully characterized as forms of probabilistic inference. For example, researchers have developed probabilistic models that span the range of inferences from classical conditioning (Courville, Daw, & Touretzky, 2006) to scientific discovery (Glymour, 2001). Even if the probabilistic approach succeeds in providing a unifying perspective on inductive reasoning, demonstrating that a given inference is consistent with probabilistic reasoning is at best a starting point. To turn this demonstration into a fully specified psychological account, it will be necessary to characterize the knowledge structures involved and the inference processes that operate over these structures. It seems likely that many qualitatively different structures and processes will need to be invoked in order to explain how humans solve all of the problems in our taxonomy.

### Acquiring mental representations

Our taxonomy deliberately focuses on problems that can be formulated extensionally as inferences about unobserved elements of a semantic system. As we mentioned in the previous section, our hope is that researchers from multiple theoretical persuasions can agree on our characterization of these problems, even if they hold very different opinions about the nature of the mental representations that support solutions to these problems. For example, categorization researchers may

agree that people often need to decide whether a given category label applies to a novel object, even if they disagree about whether the intension of the category corresponds to a rule, a prototype, or a set of exemplars.

Although we have focused on problems that can be formulated without referring to mental representations, the need to acquire these representations leads to additional inductive problems. Consider, for example, the problem of learning the intension of a rule-based category. A reasoner who has solved this problem might be able to report that an object is a wug if and only if it is green with red stripes. As a second example, consider a reasoner who is able to predict that a novel flying animal probably has wings. The mental representation that supports this ability may correspond to a causal network over features, and learning this causal network is an inductive problem.

Psychologists continue to debate the psychological reality of structured representations such as rules and causal networks (Hahn & Chater, 1998; Rogers & McClelland, 2004; Smith, Langston, & Nisbett, 1992). As a result, we believe that it is impossible to develop a general account of representation learning that will be acceptable to the broad community of researchers who study semantic cognition. Researchers who postulate the existence of structured representations, however, may find it useful to explore the probabilistic approach to representation learning. In the previous section, we argued that probabilistic inference helps to explain how structured representations are used for inductive inference, and the same general approach can also explain how these representations are acquired (Kemp & Tenenbaum, 2009; Tenenbaum, Kemp, Griffiths, & Goodman, 2011). For example, researchers have developed probabilistic models that help to explain how logical rules (Goodman, Tenenbaum, Feldman, & Griffiths, 2008) and causal networks (Devereitt & Kemp, 2012; Griffiths & Tenenbaum, 2005) are acquired.

### The problem view and the process view

As we described in the introduction, our taxonomy adopts the problem view of induction and attempts to characterize the space of inductive problems. Given that inductive reasoning is likely to be supported by many different processes, it will also be useful to develop taxonomies that adopt the process view and that assign two inferences to the same class if they are carried out by the same underlying process. Our taxonomy does not meet this criterion, because two inferences that are supported by very different processes may be assigned to the same cell in Fig. 3. For example, generalizations about unfamiliar objects may rely on the rapid computation of visual similarity: If object  $X$  has feature  $F$  and object  $Y$  looks similar to  $X$ , then  $Y$  is likely to have feature  $F$ . Generalizations about familiar objects may involve the explicit formation and

evaluation of hypotheses: For example, if salmon have  $F$ , maybe grizzly bears have  $F$  if feature  $F$  can be transmitted from prey to predator (Medin, Coley, Storms, & Hayes, 2003; Shafto, Kemp, Bonawitz, Coley, & Tenenbaum, 2008). Although the processes involved in these generalizations may be rather different, our taxonomy classifies both inferences as instances of feature generalization.

In other cases, problems assigned to different cells in our taxonomy may be addressed using similar or identical processes. For example, we previously suggested that Le Verrier solved a discovery problem when he inferred the existence of Neptune, and that Halley solved an identification problem when he inferred that three comets were one and the same. Both problems were solved using a process that involved explicit mathematical computation. More generally, scientific problems appear in all of the cells in our taxonomy, and the same basic processes of hypothesis formation and evaluation may be relevant to many of these problems.

Although our taxonomy is not organized around psychological processes, in the introduction we described four contributions that problem-based taxonomies can make to the psychological literature on induction. Here we review two of those contributions and suggest that they are especially relevant to researchers who develop process models of induction. First, our taxonomy lays out the problems that process models will need to address, including several problems that have previously received little attention. Identifying novel problems may not be especially useful if those problems are contrived or theoretically unilluminating. The novel problems in our taxonomy, however, are all in the neighborhood of more familiar problems, and many of them raise theoretical questions that follow naturally from existing work on induction. For example, the novel problems in Fig. 3a and c highlight questions about how multiple sources of knowledge are integrated in order to make inductive inferences (Kemp, Chang, et al., 2010; Kemp et al., 2012). Second, our taxonomy highlights relationships between inductive problems, and therefore motivates theoretical accounts that handle two or more of these problems. For example, the problems of feature generalization, category generalization, and object generalization are conceptually similar according to our taxonomy, which suggests that it may be useful to develop process models that can handle all three problems.

### The origin of semantic systems

Our taxonomy is founded on the idea that inductive problems can be characterized as inferences about partially observed semantic systems. We have taken these semantic systems as a starting point, but it is important to consider how these systems might arise. The systems discussed in this article have included objects, features, categories, and relations, and some

of these conceptual elements can be acquired by solving discovery problems. For example, we have suggested that solving discovery problems allows learners to infer the existence of objects that they have not directly observed, and to recognize novel categories and features. Some conceptual elements, however, must be acquired by means other than discovery. In particular, the very first discovery problem encountered by a learner requires an inference about a partially observed semantic system that is not acquired via discovery.

Aside from discovery, the elements of semantic systems can arise in at least two distinct ways. Some of these elements may be provided by the perceptual system. For example, the visual system may supply information about objects and their features that can serve as the basis for subsequent inductive inferences (Spelke, 1990). Other conceptual elements may be supplied by direct instruction. For example, a child may be told directly that dolphins are mammals, and that dirt has tiny bugs inside that make people sick.

Both of these pathways to semantic knowledge involve inductive inferences that are not captured by our taxonomy. For example, inductive inferences are needed to infer the shape and color of an object if it is partially occluded and the spectrum of the illuminating light is not known. Similarly, accepting a teacher's claim that dolphins are mammals may require an inductive inference that the teacher is telling the truth. Semantic systems are a useful starting point when considering the kind of inductive problems that are discussed in the literature on semantic cognition, but a complete account of inductive reasoning will need to explain in full how these systems are constructed.

### Complexity of inductive problems

Our taxonomy includes a variety of problems, and it is natural to ask whether some of these problems are intrinsically more difficult than others. Each problem requires inferences about unobserved elements of a semantic system, and in this section we discuss whether some semantic systems are more difficult than others to think about and whether some ways of concealing the elements of a semantic system make for especially difficult inductive problems. Of necessity, proposals about the cognitive complexity of an inductive problem will depend on assumptions about cognitive processing. The proposals in this section therefore illustrate the kind of connections that might exist between the problem and process views of inductive reasoning.

Semantic systems like those in Fig. 1 are built from objects, categories, features, and relations, and some of these systems appear to be intrinsically more complex than others. For example, the system in Fig. 1e includes a category defined over compound objects, and it is therefore more complex than the system in Fig. 1a, which does not include compound

objects. Similarly, systems that include higher-order features, relations, or categories will tend to be more complex than systems that do not incorporate higher-order elements. Complex semantic systems will tend to be difficult to fit into working memory, and will therefore be relatively difficult to reason about. Halford, Wilson, and Phillips (1998) have developed a complexity metric that captures this idea, and have argued that young children and nonhuman primates find it difficult to think about systems that are high in complexity according to their metric.

Most of the examples discussed by Halford et al. (1998) involve deductive tasks rather than inductive tasks, but some evidence suggests that a similar complexity ordering applies to inductive tasks. For example, Gentner, Rattermann, Markman, and Kotovsky (1995) developed a task in which children were shown two sets of three objects and were asked to pick a target object in Set 2 that matched a source object from Set 1. Some problems could be solved using a relation between objects; for example, children could pick a target object from Set 2 that was identical to the source object from Set 1. Other problems could only be solved using a higher-order relation between relations. For example, the source and target objects might both be the largest objects in their respective sets. Gentner et al. found that the ability to think about higher-order relations increases with age, and similar results have been reported by other psychologists (Piaget, Montangero, & Billeter, 1977; Sternberg & Nigro, 1980). Comparative psychologists have provided converging evidence that suggests that higher-order relations are intrinsically more difficult to think about than relations between objects. For example, Penn, Holyoak, and Povinelli (2008) suggested that there is no compelling evidence that nonhuman animals can reason about higher-order relations in a systematic way.

In addition to considering the relative complexity of different semantic systems, we can consider the relative complexity of different inductive problems that are formulated with respect to the same semantic system. Our taxonomy is organized around the problems of generalization, discovery, and identification, and discovery appears to be the most challenging of the three. Developmental studies suggest that generalization problems (Cohen, Gelber, & Lazar, 1971) and identification problems (Spelke et al., 1995) can be solved before the age of four months, but we know of no studies that have suggested that discovery problems can be solved before the age of eight months (Csibra & Volein, 2008). Comparative psychologists have provided converging evidence that discovery is more difficult than generalization or identification. For example, nonhuman primates can solve generalization (Sigala, Gabbiani, & Logothetis, 2002) and identification (Mendes et al., 2008) problems, but Premack (2010) wrote that "there is no suggestion that chimpanzees are capable of inferring unobserved objects" (p. 27).

Although discovery problems often seem to be more challenging than generalization or identification problems, this difference in complexity is at best a rule of thumb. Some discovery problems are relatively simple, and some generalization and identification problems are relatively challenging. For example, a discovery problem in which infants use gaze direction to infer the existence of an unobserved object (Csibra & Volein, 2008) seems simpler than the identification problem solved by Halley. Deciding which of two problems is easier will typically require consideration of factors that go beyond our taxonomy, including the background knowledge that reasoners bring to the problems and the ways in which this background knowledge must be put to use.

We have argued that humans are able to solve all of the problems included in our taxonomy, but it is possible that inductive problems exist outside of this taxonomy that are too complex for humans to solve. There will certainly be problem *instances* that are impossible for humans to solve without external assistance—for example, generalization problems that require a reasoner to draw together a body of evidence that is simply too large for any one person to learn and remember. We propose, however, that there are no general *classes* of inductive problems that humans are unable to solve. Inductive reasoning is shaped by factors that operate across many cognitive domains, including constraints on working memory and other aspects of executive function. These constraints, however, are not specific to induction, and we propose that no induction-specific constraints limit the class of problems that people are able to solve. Consistent with this view, our taxonomy is motivated by the goal of capturing *all* inductive problems that can be formulated as inferences about unobserved aspects of a semantic system.

### Extending the taxonomy

Our taxonomy includes a wide range of problems but will need to be extended to cover the full space of inductive inferences that humans are able to make. Three extensions in particular seem especially important. First, the semantic systems in this article are static, and an obvious extension would be to allow for semantic systems that vary over time. Second, the semantic systems in this article all capture the state of the world, but an extended taxonomy would allow for inferences about what another person believes or knows about the world. One way to accommodate these inferences would be to add “agents” to the set of conceptual building blocks that currently includes objects, features, categories, and relations. Each agent maintains a semantic system that captures its knowledge and beliefs about the world, and agents can also make inferences about the semantic systems possessed by other agents. Finally, this article has focused on inferences about semantic systems that capture the *actual*

state of the world. An extended taxonomy that includes the notion of possible worlds could accommodate modal inferences (Kemp, Han, & Jern, 2011) and inferences about how the world might have been under various counterfactual scenarios (Pearl, 2000; Rips, 2010).

### Conclusion

Psychologists dream of developing unified theories of cognition (Newell, 1989), and our long-term goal is only slightly more modest: We dream of a unified theory of inductive reasoning. In order to reach this goal, it will be necessary to understand the space of inductive problems that people are able to solve. We have taken a step in this direction by providing a taxonomy of inductive problems that arise within the domain of semantic cognition.

Our taxonomy is founded on two core ideas. We began by proposing that semantic knowledge consists of knowledge about objects, features, categories, and relations, and that these conceptual elements can be organized into semantic systems. We then proposed that inductive problems can be characterized as inferences about partially observed semantic systems. All of the problems in our taxonomy are closely related, and it is surprising that some of them have received little previous attention. Future work can aim to explore all of these problems in detail.

This article focused on characterizing inductive problems instead of describing how they can be solved, but we suggested that a probabilistic approach provides a promising way to work toward a formal framework that addresses all of the problems in our taxonomy. Previous researchers have described probabilistic models that can address many of the individual problems in our taxonomy, and developing a family of models that collectively address all of these problems is a worthy challenge for future research.

**Author note** An early version of this work was presented at the 31st Annual Conference of the Cognitive Science Society. We thank John Anderson, Kai-min Chang, Anna Fisher, Michael Lee, Tania Lombrozo, Mike Oaksford, David Plaut, and Momme von Sydow for valuable suggestions. This work was supported by NSF Grant No. CDI-0835797 and by the Pittsburgh Life Sciences Greenhouse Opportunity Fund.

### References

- Aamodt, A., & Plaza, E. (1994). Case-based reasoning: Foundational issues, methodological variations, and system approaches. *AI Communications*, *7*, 39–59.
- Adelson, E. H., & Bergen, J. R. (1991). The plenoptic function and the elements of early vision. In M. Landy & J. A. Movshon (Eds.), *Computational models of visual processing* (pp. 3–20). Cambridge, MA: MIT Press.

- Ahn, W.-K., & Medin, D. L. (1992). A two-stage model of category construction. *Cognitive Science*, *16*, 81–121.
- Anderson, A. L., Ross, B. H., & Chin-Parker, S. (2002). A further investigation of category learning by inference. *Memory & Cognition*, *30*, 119–128. doi:10.3758/BF03195271
- Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review*, *98*, 409–429. doi:10.1037/0033-295X.98.3.409
- Anderson, J. R., & Fincham, J. M. (1996). Categorization and sensitivity to correlation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*, 259–277. doi:10.1037/0278-7393.22.2.259
- Ashby, F. G., & Maddox, W. T. (1998). Stimulus categorization. In M. H. Birnbaum (Ed.), *Handbook of perception and cognition: Judgment, decision making and measurement* (pp. 251–301). San Diego, CA: Academic Press.
- Austerweil, J. L., & Griffiths, T. L. (2009). Analyzing human feature learning as nonparametric Bayesian inference. In D. Koller, D. Schuurmans, Y. Bengio, & L. Bottou (Eds.), *Advances in neural information processing systems 21* (pp. 97–104). Cambridge, MA: MIT Press.
- Austerweil, J. L., & Griffiths, T. L. (2011). A rational model of the effects of distributional information on feature learning. *Cognitive Psychology*, *4*, 173–209. doi:10.1016/j.cogpsych.2011.08.002
- Berlin, B., Breedlove, D. E., & Raven, P. H. (1974). *Principles of Tzeltal plant classification: An introduction to the botanical ethnography of a Mayan-speaking people of highland Chiapas*. New York: Academic Press.
- Bisanz, J., Bisanz, G. L., & Korpan, C. A. (1994). Inductive reasoning. In R. J. Sternberg (Ed.), *Thinking and Problem Solving*. Academic Press.
- Bloom, P. (2000). *How children learn the meanings of words*. Cambridge, MA: MIT Press.
- Bower, T. G. R. (1974). *Development in infancy*. San Francisco: W. H. Freeman.
- Braun-Lamesch, M. M. (1962). Le rôle du contexte dans la compréhension du langage chez l'enfant. *Psychologie Française*, *7*, 180–89.
- Brown, J. S. (1965). Generalization and discrimination. In D. I. Mostofsky (Ed.), *Stimulus generalization* (pp. 7–23). Stanford, CA: Stanford University Press.
- Brown, S. D., Marley, A. A. J., Donkin, C., & Heathcote, A. (2008). An integrated model of choices and response times in absolute identification. *Psychological Review*, *115*, 396–425. doi:10.1037/0033-295X.115.2.396
- Bruner, J. A., Goodnow, J. S., & Austin, G. J. (1956). *A study of thinking*. New York, NY: Wiley.
- Burke, L. (1952). On the tunnel effect. *Quarterly Journal of Experimental Psychology*, *4*, 121–138.
- Carey, S. (1985). *Conceptual change in childhood*. Cambridge, MA: MIT Press.
- Carey, S., & Bartlett, E. (1978). Acquiring a single new word. *Papers and Reports on Child Language Development*, *15*, 17–29.
- Carlson, G. (2010). Generics and concepts. In F. J. Pelletier (Ed.), *Kinds, things, and stuff: Mass terms and generics* (pp. 16–35). New York, NY: Oxford University Press.
- Carroll, C. D., & Kemp, C. (2012). Object discovery and inverse physical reasoning. In N. Miyake, D. Peebles, & R. P. Cooper (Eds.), *Proceedings of the 34th Annual Conference of the Cognitive Science Society* (pp. 180–185). Austin, TX: Cognitive Science Society.
- Chater, N., Oaksford, M., Hahn, U., & Heit, E. (2011). Inductive logic and empirical psychology. In D. M. Gabbay, S. Hartmann, & J. Woods (Eds.), *Handbook of the history of logic. Vol. 10: Inductive logic* (pp. 553–624). Amsterdam, The Netherlands: North Holland.
- Chater, N., Vitanyi, P. M. B., & Stewart, N. (2001). Universal generalization and universal inter-item confusability. *Behavioral and Brain Sciences*, *24*, 659–660.
- Christie, S., & Gentner, D. (2010). Where hypotheses come from: Learning new relations by structural alignment. *Journal of Cognition and Development*, *11*, 356–373.
- Cohen, L. B., Gelber, E. R., & Lazar, M. A. (1971). Infant habituation and generalization to differing degrees of stimulus novelty. *Journal of Experimental Child Psychology*, *11*, 379–389.
- Colberg, M., Nester, M. A., & Trattner, M. H. (1985). Convergence of the inductive and deductive models in the measurement of reasoning abilities. *Journal of Applied Psychology*, *70*, 681–694.
- Courville, A. C., Daw, N. D., & Touretzky, D. S. (2006). Bayesian theories of conditioning in a changing world. *Trends in Cognitive Sciences*, *10*, 294–300. doi:10.1016/j.tics.2006.05.004
- Csibra, G., & Volein, A. (2008). Infants can infer the presence of hidden objects from referential gaze information. *British Journal of Developmental Psychology*, *26*, 1–11.
- Deverett, B., & Kemp, C. (2012). Learning deterministic causal networks from observational data. In N. Miyake, D. Peebles, & R. P. Cooper (Eds.), *Proceedings of the 34th Annual Conference of the Cognitive Science Society* (pp. 288–293). Austin, TX: Cognitive Science Society.
- Doumas, L. A. A., Hummel, J. E., & Sandhofer, C. M. (2008). A theory of the discovery and predication of relational concepts. *Psychological Review*, *115*, 1–43. doi:10.1037/0033-295X.115.1.1
- Ennis, D. M. (1988). Confusable and discriminable stimuli: Comment on Nosofsky (1986) and Shepard (1986). *Journal of Experimental Psychology: General*, *117*, 408–411.
- Estes, W. K. (1994). *Classification and cognition*. New York, NY: Oxford University Press.
- Fischer, P., & Zytow, J. M. (1992). Incremental generation and exploration of hidden structure. In J. M. Zytow (Ed.), *Proceedings of the ML-92 Workshop on Machine Discovery* (pp. 103–110). Wichita, KS: National Institute for Aviation Research.
- Fisher, D. H. (1987). Knowledge acquisition via incremental conceptual clustering. *Machine Learning*, *2*, 139–172.
- Fodor, J. A. (1975). *The language of thought*. Cambridge, MA: Harvard University Press.
- Gelman, S. A. (2003). *The essential child: Origins of essentialism in everyday thought*. Oxford, UK: Oxford University Press.
- Gelman, S. A., & Markman, E. M. (1986). Categories and induction in young children. *Cognition*, *23*, 183–209.
- Gelman, S. A., & Markman, E. M. (1987). Young children's inductions from natural kinds: The role of categories and appearances. *Child Development*, *58*, 1532–1541.
- Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, *7*, 155–170.
- Gentner, D., & Kurtz, K. (2005). Relational categories. In W. Ahn, R. L. Goldstone, B. C. Love, A. B. Markman, & P. Wolff (Eds.), *Categorization inside and outside the lab: Essays in honor of Douglas L. Medin* (pp. 151–175). Washington, DC: American Psychological Association.
- Gentner, D., Rattermann, M. J., Markman, A. B., & Kotovsky, L. (1995). Two forces in the development of relational similarity. In T. J. Simon & G. S. Halford (Eds.), *Developing cognitive competence: New approaches to process modeling* (pp. 263–313). Hillsdale, NJ: Erlbaum.
- Gick, M. L., & Holyoak, K. J. (1983). Schema induction and analogical transfer. *Cognitive Psychology*, *15*, 1–38. doi:10.1016/0010-0285(83)90002-6
- Gluck, M. A. (1991). Stimulus generalization and representation in adaptive network models of category learning. *Psychological Science*, *2*, 50–55.
- Glymour, C. (2001). *The mind's arrows: Bayes nets and graphical causal models in psychology*. Cambridge, MA: MIT Press.
- Goodman, J. C., McDonough, L., & Brown, N. N. (1998). The role of semantic context and memory in the acquisition of novel nouns. *Child Development*, *69*, 1330–1344.
- Goodman, N. D., Tenenbaum, J. B., Feldman, J., & Griffiths, T. L. (2008). A rational analysis of rule-based concept learning. *Cognitive Science*, *32*, 108–154.

- Gopnik, A., & Meltzoff, A. N. (1997). *Words, thoughts, and theories*. Cambridge, MA: MIT Press.
- Griffiths, T. L., & Tenenbaum, J. B. (2005). Structure and strength in causal induction. *Cognitive Psychology*, *51*, 354–384. doi:10.1016/j.cogpsych.2005.05.004
- Guttman, N., & Kalish, H. I. (1956). Discriminability and stimulus generalization. *Journal of Experimental Psychology*, *51*, 79–88.
- Hadjichristidis, C., Sloman, S., Stevenson, R., & Over, D. (2004). Feature centrality and property induction. *Cognitive Science*, *28*, 45–74.
- Hahn, U., & Chater, N. (1998). Similarity and rules: Distinct? exhaustive? empirically distinguishable? *Cognition*, *65*, 197–230.
- Halford, G. S., Wilson, W. H., & Phillips, W. (1998). Processing capacity defined by relational complexity: Implications for comparative, developmental and cognitive psychology. *Behavioral and Brain Sciences*, *21*, 803–831.
- Hayes, B. K., Fritz, K., & Heit, E. (2013). The relationship between memory and inductive reasoning: Does it develop? *Developmental Psychology*, *49*, 848–860.
- Hayes, B. K., Heit, E., & Swendsen, H. (2010). Inductive reasoning. *Wiley Interdisciplinary Reviews: Cognitive Science*, *1*, 278–292. doi:10.1002/wcs.44
- Hayes, B. K., & Thompson, S. P. (2007). Causal relations and feature similarity in children's inductive reasoning. *Journal of Experimental Psychology: General*, *136*, 470–484.
- Heibeck, T., & Markman, E. (1987). Word learning in children: An examination of fast mapping. *Child Development*, *58*, 1021–1024.
- Heit, E. (1998). A Bayesian analysis of some forms of inductive reasoning. In M. Oaksford & N. Chater (Eds.), *Rational models of cognition* (pp. 248–274). Oxford, UK: Oxford University Press.
- Heit, E. (2000). Properties of inductive reasoning. *Psychonomic Bulletin & Review*, *7*, 569–592. doi:10.3758/BF03212996
- Heit, E. (2007). What is induction and why study it? In A. Feeney & E. Heit (Eds.), *Inductive reasoning: Experimental, developmental and computational approaches* (pp. 1–24). Cambridge, UK: Cambridge University Press.
- Heit, E. (2008). Models of inductive reasoning. In R. Sun (Ed.), *The Cambridge handbook of computational psychology* (pp. 322–338). Cambridge, UK: Cambridge University Press.
- Heit, E., Rotello, C. M., & Hayes, B. K. (2012). Relations between memory and reasoning. In B. H. Ross (Ed.), *The psychology of learning and motivation (Vol. 57)* (pp. 57–101). San Diego, CA: Academic Press.
- Heit, E., & Rubinstein, J. (1994). Similarity and property effects in inductive reasoning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 411–422.
- Holland, J. H., Holyoak, K. J., Nisbett, R. E., & Thagard, P. R. (1986). *Induction: Processes of inference, learning and discovery*. Cambridge, MA: MIT Press.
- Holyoak, K. J., & Cheng, P. W. (2011). Causal learning and inference as a rational process: The new synthesis. *Annual Review of Psychology*, *62*, 135–163. doi:10.1146/annurev.psych.121208.131634
- Holyoak, K. J., Lee, H. S., & Lu, H. (2010). Analogical and category-based inference: A theoretical integration with Bayesian causal models. *Journal of Experimental Psychology: General*, *139*, 702–727.
- Holyoak, K. J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science*, *13*, 295–355.
- Howson, C., & Urbach, P. (1993). *Scientific reasoning: The Bayesian approach*. Chicago, IL: Open Court.
- Hull, C. (1943). *Principles of behavior*. New York, NY: Appleton-Century-Crofts.
- Hunt, E. B. (1962). *Concept learning: An information processing problem*. New York, NY: Wiley.
- Imai, M. (1995). Asymmetry in the taxonomic assumption: Word learning versus property induction. In E. V. Clark (Ed.), *Proceedings of the 27th Annual Child Language Research Forum* (pp. 157–167). Stanford, CA: Stanford University, Center for the Study of Language and Information.
- Inhelder, B., & Piaget, J. (1964). *The early growth of logic in the child*. London, UK: Routledge & Kegan Paul.
- Jern, A., & Kemp, C. (2013). A probabilistic account of exemplar and category generation. *Cognitive Psychology*, *66*, 85–125.
- Keil, F. C. (1979). *Semantic and conceptual development*. Cambridge, MA: Harvard University Press.
- Kemp, C. (2011). Inductive reasoning about chimeric creatures. In J. Shawe-Taylor, R. Zemel, P. Bartlett, F. C. N. Pereira, & K. Q. Weinberger (Eds.), *Advances in neural information processing systems 24* (pp. 316–324). Cambridge, MA: MIT Press.
- Kemp, C. (2012). Exploring the conceptual universe. *Psychological Review*, *119*, 685–722.
- Kemp, C., Chang, K. K., & Lombardi, L. (2010a). Category and feature identification. *Acta Psychologica*, *133*, 216–233.
- Kemp, C., Han, F., & Jern, A. (2011). Concept learning and modal reasoning. In L. Carlson, C. Hölscher, & T. F. Shipley (Eds.), *Proceedings of the 33rd Annual Conference of the Cognitive Science Society* (pp. 513–518). Austin, TX: Cognitive Science Society.
- Kemp, C., & Jern, A. (2009). Abstraction and relational learning. In Y. Bengio, D. Schuurmans, J. Lafferty, C. K. I. Williams, & A. Culotta (Eds.), *Advances in neural information processing systems 22* (pp. 934–942). Cambridge, MA: MIT Press.
- Kemp, C., Jern, A., & Xu, F. (2009). Individuation, identification and object discovery. In Y. Bengio, D. Schuurmans, J. Lafferty, C. K. I. Williams, & A. Culotta (Eds.), *Advances in neural information processing systems 22* (pp. 925–933). Cambridge, MA: MIT Press.
- Kemp, C., Shafto, P., & Tenenbaum, J. B. (2012). An integrated account of generalization across objects and features. *Cognitive Psychology*, *1–2*, 35–73.
- Kemp, C., & Tenenbaum, J. B. (2009). Structured statistical models of inductive reasoning. *Psychological Review*, *116*, 20–58. doi:10.1037/a0014282
- Kemp, C., Tenenbaum, J. B., Niyogi, S., & Griffiths, T. L. (2010b). A probabilistic model of theory formation. *Cognition*, *114*, 165–196. doi:10.1016/j.cognition.2009.09.003
- Lacouture, Y., & Marley, A. A. J. (1991). A connectionist model of choice and reaction time in absolute identification. *Connection Science*, *3*, 401–433.
- Lambon Ralph, M. A., Graham, K. S., Ellis, A. W., & Hodges, J. R. (1998). Naming in semantic dementia—what matters? *Neuropsychologia*, *36*, 775–784.
- Landy, D., & Goldstone, R. L. (2005). How we learn about things we don't already understand. *Journal of Experimental and Theoretical Artificial Intelligence*, *17*, 343–369.
- Langley, P., Simon, H. A., Bradshaw, G. L., & Zytkow, J. M. (1987). *Scientific discovery: Computational explorations of the creative process*. Cambridge, MA: MIT Press.
- Leverington, D. (2003). *Babylon to Voyager and beyond: A history of planetary astronomy*. Cambridge, UK: University of Cambridge.
- Lombardi, L., & Sartori, G. (2007). Models of relevant cue integration in name retrieval. *Journal of Memory and Language*, *57*, 101–125.
- López, A., Atran, S., Coley, J. D., Medin, D., & Smith, E. E. (1997). The tree of life: Universal and cultural features of folkbiological taxonomies and inductions. *Cognitive Psychology*, *32*, 251–295.
- Love, B. C. (2002). Comparing supervised and unsupervised category learning. *Psychonomic Bulletin & Review*, *9*, 829–835. doi:10.3758/BF03196342
- Luce, R. D., Green, D. M., & Weber, D. L. (1976). Attention bands in absolute identification. *Perception & Psychophysics*, *20*, 49–54. doi:10.3758/BF03198705
- MacNamara, J. (1972). Cognitive basis of language learning in infants. *Psychological Review*, *79*, 1–13. doi:10.1037/h0031901

- Maratsos, M. (2001). How fast does a child learn a word? *Behavioral and Brain Sciences*, *24*, 1111–1112.
- Markman, E. M. (1989). *Categorization and naming in children: Problems of induction*. Cambridge, MA: MIT Press.
- Markman, E. M. (1992). Constraints on word learning: Speculations about their nature, origins, and domain specificity. In M. R. Gunnar & M. Maratsos (Eds.), *Modularity and constraints in language and cognition* (pp. 59–101). Hillsdale, NJ: Erlbaum.
- Marr, D. (1982). *Vision*. San Francisco, CA: W. H. Freeman.
- Marzano, R. J., & Pickering, D. J. (2005). *Building academic vocabulary: Teacher's manual*. Alexandria, VA: Association for Supervision and Curriculum Development.
- McCarrell, N. S., & Callanan, M. A. (1995). Form–function correspondences in children's inference. *Child Development*, *66*, 532–546.
- Medin, D. L., Coley, J. D., Storms, G., & Hayes, B. L. (2003). A relevance theory of induction. *Psychonomic Bulletin & Review*, *10*, 517–532. doi:10.3758/BF03196515
- Medin, D. L., Wattenmaker, W. D., & Hampson, S. E. (1987). Family resemblance, conceptual cohesiveness and category construction. *Cognitive Psychology*, *19*, 242–279.
- Mendes, N., Rakoczy, H., & Call, J. (2008). Ape metaphysics: Object individuation without language. *Cognition*, *106*, 730–749.
- Merriman, W. E., Schuster, J. M., & Hager, L. (1991). Are names ever mapped onto preexisting categories? *Journal of Experimental Psychology: General*, *120*, 288–300.
- Mervis, C. B. (1987). Child-basic object categories and early lexical development. In U. Neisser (Ed.), *Concepts and conceptual development* (pp. 201–233). Cambridge, UK: Cambridge University Press.
- Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, *63*, 81–97.
- Murphy, G. L. (1993). Theories and concept formation. In I. Van Mechelen, J. Hampton, P. Theuns, & R. S. Michalski (Eds.), *Categories and concepts: Theoretical views and inductive data analysis* (pp. 173–200). London, UK: Academic Press.
- Murphy, G. L., & Ross, B. H. (2010). Category vs. object knowledge in category-based induction. *Journal of Memory and Language*, *63*, 1–17.
- Newell, A. (1989). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.
- Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, *115*, 39–57.
- Nosofsky, R. M. (1988). On exemplar-based exemplar representations: Reply to Ennis (1988). *Journal of Experimental Psychology: General*, *117*, 412–414.
- Nosofsky, R. M. (1992). Exemplar-based approach to relating categorization, identification and recognition. In F. G. Ashby (Ed.), *Multidimensional models of perception and cognition* (pp. 363–393). Hillsdale, NJ: Erlbaum.
- Osherson, D. N., Smith, E. E., Wilkie, O., Lopez, A., & Shafir, E. (1990). Category-based induction. *Psychological Review*, *97*, 185–200.
- Pavlov, I. P. (1927). *Conditioned reflexes: An investigation of the physiological activities of the cerebral cortex (G. V. Anrep, Trans.)*. London, UK: Oxford University Press.
- Pearl, J. (2000). *Causality: Models, reasoning and inference*. Cambridge, UK: Cambridge University Press.
- Peirce, C. S. (1957). The logic of abduction. In V. Tomas (Ed.), *Peirce's essays in the philosophy of science*. New York, NY: Liberal Arts Press.
- Penn, D. C., Holyoak, K. J., & Povinelli, D. J. (2008). Darwin's mistake: Explaining the discontinuity between human and nonhuman minds. *Behavioral and Brain Sciences*, *31*, 109–129.
- Piaget, J., Montangero, J., & Billeter, J. (1977). La formation des correlats. In J. Piaget (Ed.), *Recherches sur l'abstraction réfléchissante I* (pp. 115–129). Paris, France: Presses Universitaires de France.
- Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, *77*, 353–363.
- Pothos, E. M., & Chater, N. (2002). A simplicity principle in unsupervised human categorization. *Cognitive Science*, *26*, 303–343.
- Pothos, E. M., Perlman, A., Bailey, T. M., Kurtz, K., Edwards, D. J., Hines, P., & McDonnell, J. V. (2011). Measuring category intuitiveness in unconstrained categorization tasks. *Cognition*, *121*, 83–100. doi:10.1016/j.cognition.2011.06.002
- Premack, D. (2010). Why humans are unique: Three theories. *Perspectives on Psychological Science*, *5*, 22–32.
- Quine, W. V. O., & Ullian, J. (1978). *The web of belief*. New York, NY: Random House.
- Reber, A. S., & Reber, E. S. (2001). *The Penguin dictionary of psychology*. London, UK: Penguin.
- Rehder, B., & Burnett, R. (2005). Feature inference and the causal structure of categories. *Cognitive Psychology*, *50*, 264–314.
- Rendell, L., & Seshu, R. (1990). Learning hard concepts through constructive induction: Framework and rationale. *Computational Intelligence*, *6*, 247–270.
- Rips, L. J. (1975). Inductive judgments about natural categories. *Journal of Verbal Learning and Verbal Behavior*, *14*, 665–681.
- Rips, L. J. (2010). Two causal theories of counterfactual conditionals. *Cognitive Science*, *34*, 175–221. doi:10.1111/j.1551-6709.2009.01080.x
- Rogers, T. T., & McClelland, J. L. (2004). *Semantic cognition: A parallel distributed processing approach*. Cambridge, MA: MIT Press.
- Rumelhart, D. E., & Abrahamson, A. A. (1973). A model for analogical reasoning. *Cognitive Psychology*, *5*, 1–28.
- Sartori, G., & Lombardi, L. (2004). Semantic relevance and semantic disorders. *Journal of Cognitive Neuroscience*, *16*, 439–452.
- Saxe, R., Tenenbaum, J. B., & Carey, S. (2005). Secret agents: Inferences about hidden causes by 10- and 12-month-old infants. *Psychological Science*, *16*, 995–1001.
- Schyns, P. G., Goldstone, R. L., & Thibaut, J.-P. (1998). The development of features in object concepts. *Behavioral and Brain Sciences*, *21*, 1–17. doi:10.1017/S0140525X98520109
- Shafto, P., Kemp, C., Bonawitz, E. B., Coley, J. D., & Tenenbaum, J. B. (2008). Inductive reasoning about causally transmitted properties. *Cognition*, *109*, 175–192.
- Shafto, P., Kemp, C., Mansinghka, V., & Tenenbaum, J. B. (2011). A probabilistic model of cross-categorization. *Cognition*, *120*, 1–25.
- Shepard, R. N. (1957). Stimulus and response generalization: A stochastic model relating generalization to distance in psychological space. *Psychometrika*, *22*, 325–345.
- Shepard, R. N. (1986). Discrimination and generalization in identification and classification: Comment on Nosofsky. *Journal of Experimental Psychology: General*, *115*, 58–61.
- Shepard, R. N. (1987). Towards a universal law of generalization for psychological science. *Science*, *237*, 1317–1323. doi:10.1126/science.3629243
- Sigala, N., Gabbiani, F., & Logothetis, N. K. (2002). Visual categorization and object representation in monkeys and humans. *Journal of Cognitive Neuroscience*, *14*, 187–198.
- Skinner, B. F. (1938). *The behavior of organisms: An experimental analysis*. New York, NY: Appleton-Century.
- Skyrms, B. (1975). *Choice and chance: An introduction to inductive logic*. Belmont, CA: Dickenson.
- Sloman, S. A. (2005). *Causal models: How people think about the world and its alternatives*. Oxford: Oxford University Press.
- Sloman, S. A. (2007). Taxonomizing induction. In A. Feeney & E. Heit (Eds.), *Inductive reasoning: Experimental, developmental and computational approaches* (pp. 328–343). Cambridge, UK: Cambridge University Press.

- Sloman, S. A., & Lagnado, D. A. (2005). The problem of induction. In R. Morrison & K. Holyoak (Eds.), *Cambridge handbook of thinking and reasoning* (pp. 95–116). New York, NY: Cambridge University Press.
- Sloman, S. A., & Wisniewski, E. (1992). Extending the domain of a feature-based model of property induction. In J. K. Kruschke (Ed.), *Proceedings of the 14th Annual Conference of the Cognitive Science Society* (pp. 355–359). Hillsdale, NJ: Erlbaum.
- Sloutsky, V. M., & Fisher, A. V. (2002). The development of categorization. In B. H. Ross (Ed.), *The psychology of learning and motivation* (Vol. 54, pp. 142–166). New York, NY: Academic Press.
- Sloutsky, V. M., & Fisher, A. V. (2004). Induction and categorization in young children: A similarity based model. *Journal of Experimental Psychology: General*, *133*, 166–188.
- Smith, E. E., Langston, C., & Nisbett, R. E. (1992). The case for rules in reasoning. *Cognitive Science*, *16*, 1–40.
- Smith, L. B., & Heise, D. (1992). Perceptual similarity and conceptual structure. In B. Burns (Ed.), *Percepts, concepts and categories*. Amsterdam, The Netherlands: Elsevier.
- Smoke, K. L. (1935). The experimental approach to concept learning. *Psychological Review*, *42*, 274–279.
- Snedeker, J., & Gleitman, L. (2004). Why it is hard to label our concepts. In D. G. Hall & S. R. Waxman (Eds.), *Weaving a lexicon* (pp. 257–293). Cambridge, MA: MIT Press.
- Spelke, E. S. (1990). Principles of object perception. *Cognitive Science*, *14*, 29–56.
- Spelke, E. S., Kestenbaum, R., Simons, D. J., & Wein, D. (1995). Spatiotemporal continuity, smoothness of motion and object identity in infancy. *British Journal of Developmental Psychology*, *13*, 113–142.
- Sternberg, R. J. (1986). Toward a unified theory of human reasoning. *Intelligence*, *10*, 281–314.
- Sternberg, R. J., & Gardner, M. K. (1983). Unities in inductive reasoning. *Journal of Experimental Psychology: General*, *112*, 80–116.
- Sternberg, R. J., & Nigro, G. (1980). Developmental patterns in the solution of verbal analogies. *Child Development*, *51*, 27–38.
- Sweller, N., & Hayes, B. K. (2010). More than one kind of inference: Re-examining what's learned in feature inference and classification. *Quarterly Journal of Experimental Psychology*, *63*, 1568–1589. doi:10.1080/17470210903438547
- Tare, M., & Gelman, S. A. (2010). Determining that a label is kind-referring: Factors that influence children's and adults' novel word extensions. *Journal of Child Language*, *37*, 1007–1026.
- Tenenbaum, J. B., & Griffiths, T. L. (2001). Generalization, similarity, and Bayesian inference. *Behavioral and Brain Sciences*, *24*, 629–640. doi:10.1017/S0140525X01000061
- Tenenbaum, J. B., Griffiths, T. L., & Kemp, C. (2006). Theory-based Bayesian models of inductive learning and reasoning. *Trends in Cognitive Sciences*, *10*, 309–318.
- Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: Statistics, structure, and abstraction. *Science*, *331*, 1279–1285. doi:10.1126/science.1192788
- Thagard, P. (2001). Induction. In R. A. Wilson & F. C. Keil (Eds.), *The MIT encyclopedia of the cognitive sciences*. Cambridge, MA: MIT Press.
- Tremoulet, P. D., Leslie, A. M., & Hall, D. G. (2000). Infant individuation and identification of objects. *Cognitive Development*, *15*, 499–522.
- Valdés-Peréz, R. E., Żytkow, J. M., & Simon, H. A. (1993). Scientific model-building as search in matrix spaces. In R. Fikes & W. G. Lehnert (Eds.), *Proceedings of the Eleventh National Conference on Artificial Intelligence* (pp. 472–478). Cambridge, MA: MIT Press.
- Van de Walle, G. A., Carey, S., & Prevor, M. (2000). Bases for object individuation in infancy: Evidence from manual search. *Journal of Cognition and Development*, *1*, 249–280.
- Vickers, J. (2012). The problem of induction. In E. N. Zalta (Ed.), *Stanford encyclopedia of philosophy (Fall 2012 ed.)*. Stanford, CA: Stanford University, Metaphysics Research Lab.
- Werner, H., & Kaplan, E. (1952). The acquisition of word meanings: A developmental study. *Monographs of the Society for Research in Child Development*, *15*, 1–119.
- Wisniewski, E. J., & Medin, D. L. (1994). On the interaction of theory and data in concept learning. *Cognitive Science*, *18*, 221–281.
- Xu, F. (2005). Categories, kinds, and object individuation in infancy. In L. Gershkoff-Stowe & D. H. Rakison (Eds.), *Building object categories in developmental time* (pp. 63–89). Mahwah, NJ: Erlbaum.
- Xu, Y., & Kemp, C. (2010). Inference and communication in the game of Password. In J. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. Zemel, & A. Culotta (Eds.), *Advances in neural information processing systems 23* (pp. 2514–2522). Cambridge, MA: MIT Press.
- Yamauchi, T., & Markman, A. B. (1998). Category learning by inference and classification. *Journal of Memory and Language*, *39*, 124–148.