UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

A taxonomy of inductive problems

Permalink

<https://escholarship.org/uc/item/6wb4n6t0>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 31(31)

ISSN

1069-7977

Authors

Jern, Alan
Kemp, Charles

Publication Date

2009

Peer reviewed

A Taxonomy of Inductive Problems

Charles Kemp and Alan Jern

{ckemp, ajern}@cmu.edu

Department of Psychology, Carnegie Mellon University

Abstract

Inductive inferences about objects, properties, categories, relations, and labels 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 that includes more than thirty inductive problems. The taxonomy helps to clarify the relationships between familiar problems such as identification, stimulus generalization, and categorization, and introduces several novel problems including property identification and object discovery.

Keywords: induction; semantic cognition; generalization; categorization; 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 known elements and that made predictions about the existence of new elements. Adelson and Bergen [1] developed a “periodic table” of early vision that maps out a space of visual computations and identifies several that had previously received little attention. This paper aims to make a similar contribution to the study of inductive inference. We describe a taxonomy of inductive problems that aims to clarify the relationships between different problems and to highlight problems that have previously been overlooked.

An inductive inference goes beyond the information given and reaches a conclusion that is likely but not certain given the available evidence. Inferences of this kind are relevant to almost every area of cognition, and take place, for example, when humans predict the motion of an occluded object, guess the meaning of a novel word, or decide how to grasp an object that is encountered for the first time. We will not discuss vision, language, or motor control, but instead focus on a cluster of problems from an area that has been called semantic cognition. Research in this area aims to capture knowledge about objects and their properties, categories or collections of objects, relationships between objects, and word meanings. The relevant literature includes studies of property induction [7], categorization (both supervised [12] and unsupervised [2]), stimulus generalization [15], identification [12], and word learning [17].

Accounts of semantic cognition differ in many respects but most of them rely on six basic notions: objects, properties, categories, relations, labels, and truth values. Our taxonomy takes these six notions as a starting point and attempts to chart the space of inductive problems that can be posed given a commitment to these notions. Two familiar problems that belong to this space are categorization and property induction, or deciding whether an object has an unobserved property.

Most psychological work on inductive inference focuses on a single inductive problem, but some existing frameworks ad-

dress multiple problems [7]. For example, exemplar models that formalize objects as points in a multidimensional space have been used to account for several problems including identification, stimulus generalization, recognition, and categorization [13]. Our taxonomy includes all of these problems along with many others. Since we aim to characterize inductive problems rather than to describe the psychological mechanisms 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, logical, or connectionist approaches. The taxonomy we describe can serve as a guide for future work, and future models and experiments can address the problems that it contains.

A semantic repository

Our approach proposes that semantic knowledge can be captured in terms of objects, relations, labels and truth-values. Our goal is to characterize all inductive problems that can be formulated in terms of these notions.

We assume that knowledge about objects, relations and labels can be organized into a semantic repository. Let O be a set of objects, L be a set of labels, and T be the set $\{1, 0\}$ that includes two truth values. For most cases that we consider, set O will include individuals such as dogs, people, and chairs, and set L will include strings of phonemes such as “Fido,” “dog,” and “brown.” Here we discuss a running example where O and L include the four people and the seven labels shown in Figure 1.

Sets O , L and T correspond to primitive types, and relations are constructed out of these types. Any property can be formalized as a unary relation $r : O \rightarrow T$ that assigns a truth value (1 or 0) to each object depending on whether it has the property. Figure 1 shows a property $r_1(\cdot)$ that includes three of the four objects and can be glossed as bearded($\cdot$). A category can also be formalized as a unary relation $r : O \rightarrow T$ where the truth values now indicate whether a given object belongs to the category. Figure 1 shows a category $r_2(\cdot)$ that can be glossed as Sikh($\cdot$). Binary relations of the form $r : O \times O \rightarrow T$ assign a truth value to each pair of objects. In Figure 1, for example, relation $r_3(\cdot, \cdot)$ can be glossed as parent($\cdot, \cdot$), and assigns value 1 to all pairs (o_i, o_j) such that o_i is the parent of o_j . Relations with three or more places can also be considered, but here we focus on unary and binary relations.

Both objects and relations can be associated with labels. Object labels can be captured by a relation $r : O \times L \rightarrow T$ that indicates for each pair (o_i, l_j) whether l_j is a label of object o_i . Figure 1 shows, for example, that each of the four objects in the repository has a unique label. Labels for the unary

Figure 1: A small semantic repository that includes six relations with the types shown in Table 1. The set of relations includes unary relations that can be glossed as $\text{bearded}(\cdot)$ and $\text{Sikh}(\cdot)$, a binary relation that can be glossed as $\text{parent}(\cdot, \cdot)$, and three $\text{name}(\cdot, \cdot)$ relations that specify labels of the objects, the unary relations, and the binary relation.

Type	Description
$r : O \rightarrow T$	property
$r : O \rightarrow T$	category
$r : O \times O \rightarrow T$	binary relation
$n : O \times L \rightarrow T$	object labels
$n : (r : O \rightarrow T) \times L \rightarrow T$	property/category labels
$n : (r : O \times O \rightarrow T) \times L \rightarrow T$	binary relation labels

Table 1: Six examples of the many kinds of relations that can appear in the semantic repository. Each relation is built from the three primitive types (O , L and T). Examples of each kind of relation are shown in Figure 1.

relations are captured by a relation $n : (r : O \rightarrow T) \times L \rightarrow T$, and labels for the binary relations are captured by a relation $n : (r : O \times O \rightarrow T) \times L \rightarrow T$. The three $\text{name}(\cdot, \cdot)$ relations in Figure 1 have different types and are therefore distinct, but we will overload our notation and use $\text{name}(\cdot, \cdot)$ to refer to each of them.

A semantic repository captures what is true about the world. A very general problem faced by humans is to make inferences about the contents of this repository given partial and noisy data. This paper discusses three instances of this general problem. We first consider two problems—generalization and discovery—that arise when the available data specify a partially observed repository R_{obs} . We then consider a third problem—identification—that arises when the available data specify information about object and relation tokens, and the reasoner may be unsure whether two tokens correspond to the same relation or object. The three problems we consider are captured by the hierarchical framework in Figure 2. The ultimate goal of the reasoner is to recover the repository R at the top of the hierarchy, and this repository must be inferred given a partially observed repository R_{obs} or given token data T that form an incomplete specification of a partially observed repository.

Figure 2: A hierarchical framework for specifying inductive problems. Semantic repository R contains information about objects (o_i) and relations (r_j) and can be formalized as a list of statements. R_{obs} is a partially observed version of R . The token data T are an incomplete specification of R_{obs} that include information about object tokens (o_i^*) and relation tokens (r_j^*) and identity relations over these tokens. If the identity relations I_o and I_r are fully observed, then T and R_{obs} will be equivalent. Note, however, that the identity relations are typically partially observed, reflecting uncertainty about which tokens correspond to the same relation or object.

Generalization

Suppose first that a reasoner is given a partially observed repository R_{obs} —for example, a version of Figure 1 where one or more of the entries are replaced by question marks. *Generalization* is the problem of making inferences about these unobserved entries. A partially observed repository can be represented as a list of statements like the examples in Figures 2 and 3a. Given these statements as input, generalization can be formalized as the problem of deciding whether or not a new statement is true. In Figure 3a, for example, the reasoner must decide whether $r_2(o_3) = 1$ or $r_2(o_3) = 0$.

At least two special cases of the problem of generalization can be distinguished. *Object generalization* can be defined as the problem of making inferences about a sparsely observed object. Suppose, for instance, that you meet a new person and observe only some of his properties and relationships with others. Making inferences about this new acquaintance is an example of object generalization. *Relation generalization* can be defined as the problem of making inferences about a sparsely observed relation. Suppose, for instance, that you learn about a new property (e.g. carries the T4 gene) or a new category (e.g. mesomorph), and observe a single instance of a person with the property or belong to the category. Deciding which other people have the property or belong to the category is an example of relation generalization.

As characterized here, the problem of generalization includes several problems that go by different names in the psychological literature. Stimulus generalization and property induction are two prominent examples—the first is similar to

object generalization, and the second to relation generalization. Categorization or classification is a third problem that falls under the heading of relation generalization. Since we formalize a category as a unary relation, reasoning about the extension of a novel category reduces to the problem of relation generalization.

Discovery

One family of generalization problems includes cases where the partially-observed repository R_{obs} includes all objects and relations of interest and the inductive challenge is to infer statements about these relations and objects that are true but unobserved. A second family includes cases where some objects or relations are not mentioned at all in R_{obs} . Note, for example, that R_{obs} in Figure 2 does not mention relation $r_2(\cdot)$ or objects o_3 and o_4 . *Discovery* is the problem of inferring the existence of an object or relation that has not been observed.

The problem of *relation discovery* has received some attention in the psychological literature. In one version of the problem, the relations to be discovered are unary relations that specify the category assignments of a set of objects. The problem of discovering these categories is sometimes known as *unsupervised categorization*. We previously suggested, for example, that a learner might group Sikhs into a category ($r_2(\cdot)$ in Figure 2) without being taught about the existence of this category. To give a second example, the first European explorers to visit Australia were able to organize the animals they saw into categories without needing a teacher to provide category labels. Further examples of relation discovery can be found in the literature on theory learning [9] and inductive logic programming [16]. One common idea is that learners should search for a short description of the observations they have made, and this shortest description will sometimes rely on relations that have not been observed but that help to explain the available data. For example, if Alice and Bob both simultaneously come down with a rare illness, we may infer that the two have recently come into contact.

The problem of *object discovery* involves inferences about the existence of objects that have not been observed. Given the partially observed repository R_{obs} in Figure 2, for example, a reasoner may infer that Arjun had parents (one of whom is o_3), and that properties of these individuals can help to explain some of the properties of Arjun. Scientists have been responsible for some of the most striking examples of object discovery. Before the planet Neptune was directly observed, the existence of this object was inferred based on the way that it interacted with known objects. Many microorganisms have also been discovered without the benefit of direct observation, and Koch’s postulates characterize one case in which the existence of an unobserved organism can be confidently inferred.

Identification

So far we have assumed that the information available to a reasoner takes the form of a partially observed repository

a) $r_1(o_1)=1$ $r_1(o_2)=1$ $r_2(o_2)=1$ $r_3(o_1, o_3)=1$ $\text{name}(o_1, \text{"Arjun"})=1$	b) $r_1^*(o_1^*)=0$ $r_2^*(o_2^*)=1$ $r_3^*(o_3^*)=1$ $r_4^*(o_4^*, o_5^*)=1$ $\text{name}(o_6^*, \text{"Arjun"})=1$	I_o	I_r
		$o_1^* \ o_2^* \ o_3^* \ o_4^* \ o_5^* \ o_6^*$	$r_1^* \ r_2^* \ r_3^* \ r_4^*$
		o_1^*	r_1^*
		o_2^*	r_2^*
		o_3^*	r_3^*
		o_4^*	r_4^*
		o_5^*	
		o_6^*	

Figure 3: Inductive inferences are based on information about objects, relations and labels. (a) In some cases, the available information corresponds to a partially observed semantic repository and can be formalized as a list of statements. (b) In other cases the available information is better described as a list of statements about object and relation tokens together with identity relations I_o and I_r that specify which tokens the same object or relation. The identity relations here are fully observed, and the information in (b) is equivalent to the information in (a). In general, however, the identity relations will be partially observed as they are in Figure 2.

R_{obs} involving objects, relations, and labels. In many cases, however, this partially observed repository cannot be directly experienced and must instead be constructed from more primitive kinds of observations. Here we assume that the primitive observations take the form of the token data T in Figure 2. These token data specify information about object tokens, relation tokens, and labels, where each token is an instance of an underlying object or relation. Suppose, for example, that you have met Arjun on two occasions. The information you gather will refer to two distinct object tokens, and it is likely you will understand that these tokens are instances of the same individual. Note, however, that it is entirely possible to meet the same person twice and to think that the two object tokens are instances of different individuals.

The notion of a relation token can be introduced using a common word learning scenario. Suppose that a young child has noticed that some objects are round—in other words, she has acquired the unary relation $\text{round}(\cdot)$, although she does not yet know the name of this relation. Suppose now that the child’s father points to an orange and tells her that this object is “round.” The child now knows that there is some unary relation with this label, but may still be unsure whether the label refers to the $\text{round}(\cdot)$ relation or to some other relation. Relation tokens provide a natural way to handle this uncertainty. A relation token is created when the child thinks of the relation $\text{round}(\cdot)$, and another relation token is created and associated with the label “round.” Given these two tokens, the child may well be unsure whether they are instances of the same underlying relation.

It is possible to introduce a distinction between a label, or a sequence of phonemes, and a label token, or an utterance of a label that is spoken in a particular accent and that may include speech errors. For simplicity, we avoid making this distinction here, and assume instead that labels and naming relations $\text{name}(\cdot, \cdot)$ can both be directly observed.

The problem of identification can now be formalized. Suppose that a reasoner is given a list of statements that specify information about object tokens and relation tokens (Figure 3b). Suppose also that there are two identity relations,

one for object tokens (I_o) and one for relation tokens (I_r). $I_o(o_i^*, o_j^*) = 1$ if tokens o_i^* and o_j^* are instances of the same object, and $I_r(r_i^*, r_j^*) = 1$ if tokens r_i^* and r_j^* are instances of the same relation. Each identity relation may be partially observed—in Figure 2, for example, the reasoner is not sure whether o_1^* and o_2^* correspond to the same object, and whether r_1^* and r_2^* correspond to the same relation. The problem of identification is to infer which tokens correspond to the same underlying entity—in other words, to infer the relations I_o and I_r .

A closely related problem is known as *recognition* [13]. Suppose, for example, that Bill, Bob and Ben are triplets and that you have met all three. One day you see a boy in the store and *recognize* him as a person you have seen before, even though you cannot *identify* him as Bill, Bob or Ben. As this example suggests, recognition is the problem of deciding whether a token is an instance of a previously-observed entity without necessarily identifying the entity involved. This problem can again be formulated as an inference about the identity relations I_o and I_r in Figure 2. In the recognition setting, a reasoner may be uncertain about the contents of I_o , but may infer enough about this relation to know whether a given object has previously been observed.

Although this section has focused on identification, the two inductive problems previously described (generalization and discovery) can be formulated given raw data in the form of object and relation tokens. Once the relations I_o and I_r have been inferred, the observations in Figure 3b uniquely specify a partially observed repository, and we are back to the inductive setting considered in previous sections. The problems considered in previous sections can be posed even if a reasoner is uncertain about the identity relations I_o and I_r . Even if a reasoner is uncertain about R_{obs} she can still make inferences about unobserved properties of object tokens, and inferences about the existence of objects and relations that have not been observed.

A taxonomy of inductive problems

We have now described a framework for characterizing inductive reasoning (Figure 2). The input data T include statements about object and relation tokens along with partially observed identity relations I_o and I_r . The goal of the reasoner is to complete the identity relations, thereby specifying a partially observed repository R_{obs} , and to infer the true repository R that is partially captured by R_{obs} . In one sense we have described a single inductive problem that is very general. In another sense our framework captures many inductive problems, and this section attempts to organize these problems into a taxonomy.

The previous sections described five basic problems: generalization, object discovery, relation discovery, object identification, and relation identification. Additional problems can be created by combining two or more of these basic problems. Since there are five basic problems, there are 31 combinations that include at least one basic problem. We do not suggest that

the 31 problems specified by these combinations are equally likely to be encountered in the real world, and expect that some will turn out to be more fundamental than others. We propose, however, that many of the 31 combinations specify problems that are worthy of psychological investigation.

The rest of this section focuses on the nine combinations shown in Table 2. Each row represents a combination, and the first five columns correspond to the five basic problems. The combinations above the double line have been discussed by previous researchers, but the combinations below the line appear to be novel. Some of these combinations have previously been discussed in this paper, but for completeness we briefly review them here.

Note that our taxonomy of 31 problems is only one way of organizing the inductive problems that emerge from the framework in Figure 2. For example, Table 2 treats property induction and supervised categorization as instances of the same basic problem, but we could separate the two by distinguishing between categories and unary relations and increasing the number of problems in the taxonomy. The taxonomy could also be expanded by including separate columns for object generalization and relation generalization instead of grouping these problems. Our taxonomy provides a useful way to think about the space of inductive problems, but is by no means the only taxonomy that could be constructed.

Familiar problems

Generalization. The first row of Table 2 specifies a combination that includes only the problem of generalization. This problem has been extensively discussed, and the relevant literature includes work on stimulus generalization, property induction, and supervised categorization.

Relation discovery. The second row specifies a combination that includes both relation discovery and relation generalization. Relation discovery is needed to infer the existence of unobserved relations, and generalization is needed to infer the extensions of these relations. This combination has been previously addressed by research on unsupervised categorization and predicate invention.

Object identification. The third row specifies a combination that includes only the problem of object identification. This problem has previously been discussed by researchers including [12] and [15], and is closely related to the problem of recognition [13].

Object identification and generalization. The fourth row specifies a problem where reasoners are required to make inferences about unobserved properties of object tokens and may be uncertain whether two tokens correspond to the same object. Few researchers have set out to study this problem, but some have explored it inadvertently by designing generalization experiments where the stimuli are highly confusable. There is some debate about whether generalization gradients are exponential or Gaussian in character. One proposed resolution is that pure generalization curves are exponential, but that inferences about confusable stimuli include an identification component that produces near-Gaussian generalization

Generalization	Object discovery	Relation discovery	Object identification	Relation identification	Problem
✓					Stimulus generalization [15] Property induction [7] Supervised categorization [12]
✓		✓			Unsupervised categorization [2] Predicate invention [16]
			✓		Object identification [12] Object recognition [13]
✓			✓		Object identification and generalization [6]
✓		✓	✓		Object identification and categorization
				✓	Property identification [8] Category identification [10]
✓				✓	Inferences about ‘Property P’ [14] Word learning [17]
			✓	✓	Object and property identification [8]
✓	✓				Object discovery [5]

Table 2: A taxonomy showing 9 of the 31 problems specified by our framework. The framework includes three basic problems—generalization, discovery, and identification—and additional problems can be specified that include these basic problems as components. Each problem above the double line has been previously discussed in some detail. The problems below the double line are discussed less often, although some are connected with previous work.

curves [6]. 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 problems posed by a given task.

Novel problems

Object identification and categorization.

We have just seen that object identification and generalization can be combined, and the fifth row specifies a problem where identification is combined with unsupervised categorization. The raw data in this case are observations of object tokens, and the reasoner must decide how many distinct objects have been observed and organize these objects into categories. Infants may solve a version of this problem early in development when they are simultaneously discovering which objects their world contains and organizing these objects into categories. This problem, however, is rarely discussed in the psychological literature.

Relation identification. There are many studies of object identification, but the problem of relation identification (row 6 of Table 2) has received very little attention. In one version of this problem, the relation to be identified is a property, or a unary relation. Suppose for example that a reasoner learns that a polar bear and a dove both have ‘Property P.’ The reasoner may be able to infer that ‘Property P’ is the property of being white. Motivated in part by the taxonomy outlined here, Kemp et al. [8] have recently addressed this problem.

Real-world instances of property identification often arise when learning new words. We previously mentioned the case of a child who is told that an orange is “round”:

$$\begin{aligned} r_j^*(o_i^*) &= 1 \\ \text{name}(r_j^*, \text{“round”}) &= 1 \end{aligned} \quad (1)$$

where o_i^* is a token of a given orange. Token r_j^* will either correspond to one of the properties that is familiar to the child, or to a property that has never previously been encountered. Deciding which possibility is true is a problem of property identification.

Relation identification and generalization. Row 7 of Table 2 shows a closely related problem where the learner must not only decide whether a relation token corresponds to a familiar property, but must also decide whether the relation token can be applied to additional objects. Given, for example, that a polar bear and a dove both have property P, a reasoner can be asked to decide whether a swan has property P.

Generalization problems of this kind should be distinguished from property induction problems that do not include an identification component. Some studies ask participants to make inferences about “property P,” but others use properties like “has biotinic acid in its blood.” Property P may be interpreted as a token of a familiar property, but the “biotinic acid” example uses a novel property that does not raise the problem of identification.

The literature on property induction often blurs the distinction between these two kinds of properties, but this distinction helps to explain some findings that seem puzzling at first. Consider, for example, two studies of inductive reasoning that compared German speakers with Mandarin speakers [14]. Unlike German and English, Mandarin is a language with numeral classifiers, or linguistic categories that organize objects into groups on the basis of properties like shape (e.g. whether an object is long and thin) and function (e.g. whether an object has a handle). Mandarin speakers are therefore likely to think of these properties given problems that include a property identification component, but less likely when given pure generalization problems. The results of Saalbach and Imai

[14] are consistent with this prediction. In their first study, participants judged whether two objects were likely to “carry the same bacteria”, and the second study was identical except that the bacteria property was replaced by “property X.” No difference was found between the two groups when the bacteria property was used, but Mandarin speakers were more likely to give responses that matched their numeral classifiers in the “property X” experiment.

Real-world examples of property identification and generalization often arise in the context of word learning. Given the information in Equation 1, for example, a reasoner can be asked to extend the label “round” to other objects. Generalization tests of this kind often explore cases where a learner acquires a label (e.g. “round”) for a pre-existing concept (e.g. the property $\text{round}(\cdot)$). Many approaches to word-learning (including the modeling work of Xu and Tenenbaum [17]) are consistent with the idea that word-learning includes an identification component. Analyzing word-learning in this way helps to explain cases of ‘one-shot learning’ or ‘fast mapping’ [4] where children appear to learn a new concept in a single exposure. In reality, children are often learning a label for a pre-existing concept, which is a much simpler challenge.

Object and property identification. Object and property identification can be studied individually, but the task in the eighth row of Table 2 requires both problems to be solved. Suppose, for example, that Tweety the canary has property P , that Animal A also has property P , and that Animal A has a mane [8]. Humans can combine all of this information in order to identify the property (P is more likely to be $\text{yellow}(\cdot)$ than $\text{feathered}(\cdot)$) and the object (A is more likely to be Leo the lion than Hans the horse).

Object discovery. The final row in Table 2 represents the problem of object discovery. This problem is occasionally mentioned in the psychological literature [3] but is rarely addressed by modeling or experimental work. Note that the problem includes a generalization component, since postulating a new object is of little use unless its properties or its relationships to other objects can be inferred.

Conclusion

Psychologists dream of developing unified theories of cognition [11], and our long-term goal is only slightly more modest: we aim for a unified theory of inductive inference. In order to reach this goal it will be necessary to understand the space of inductive problems that people are able to solve. We took a step in this direction by providing a systematic description of more than thirty inductive problems involving objects, relations, and labels. All of these problems are closely related and it is surprising that many of them have received little previous attention. Future work can explore all of these problems in detail.

This paper has not focused on computational models, but future work can aim to develop a single formal framework that addresses all of the problems in our taxonomy. The problems we described can potentially be addressed by several dif-

ferent approaches, including connectionist, logical, and probabilistic approaches. For example, a probabilistic approach can be developed by defining a prior distribution over semantic repositories and a procedure by which object tokens, relation tokens, and labels are sampled from a true but unobserved repository. Given these components, a learner who observes a collection of tokens can compute a posterior distribution over identity relations and repositories, and can use this distribution to address all the problems in our taxonomy. Previous researchers [2, 8, 15, 17] have described probabilistic models that address some of the individual problems in our taxonomy, and the approach just sketched may help to unify many of these models.

Acknowledgements We thank Kai-min Chang and Michael Lee for valuable suggestions, and Momme von Sydow for comments on the manuscript. This work was supported in part by NSF grant IIS-0835797.

References

- [1] Adelson, E. H. and Bergen, J. R. (1991). The plenoptic function and the elements of early vision. In Landy, M. and Movshon, J. A., editors, *Computational models of visual processing*, pages 3–20. MIT Press, Cambridge, MA.
- [2] Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review*, 98(3):409–429.
- [3] Bruner, J. A., Goodnow, J. S., and Austin, G. J. (1956). *A study of thinking*. Wiley, New York.
- [4] Carey, S. and Bartlett, E. (1978). Acquiring a single new word. *Papers and reports on child language development*, 15:17–29.
- [5] Csibra, G. and Volein, A. (2008). Infants can infer the presence of hidden objects from referential gaze information. *British Journal of Developmental Psychology*, 26:1–11.
- [6] Ennis, D. M. (1988). Confusable and discriminable stimuli: Comment on Nosofsky (1986) and Shepard (1986). *Journal of Experimental Psychology: General*, 117(4):408–411.
- [7] Heit, E. and Hayes, B. K. (2005). Relations among categorization, induction, recognition, and similarity: Comment on Sloutsky and Fisher (2004). *Journal of Experimental Psychology: General*, 134(4):596–605.
- [8] Kemp, C., Chang, K., and Lombardi, L. (Submitted). Category and feature identification.
- [9] Kemp, C., Goodman, N. D., and Tenenbaum, J. B. (2008). Learning and using relational theories. In *NIPS 20*.
- [10] Maratsos, M. (2001). How fast does a child learn a word? *Behavioral and Brain Sciences*, 24(6):1111–1112.
- [11] Newell, A. (1989). *Unified theories of cognition*. Harvard University Press, Cambridge, MA.
- [12] Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, 115:39–57.
- [13] Nosofsky, R. M. (1992). Exemplar-based approach to relating categorization, identification and recognition. In Ashby, F. G., editor, *Multidimensional models of perception and cognition*, pages 363–393. Erlbaum, Hillsdale, NJ.
- [14] Saalbach, H. and Imai, M. (2007). Scope of linguistic influence: does a classifier system alter object concepts? *Journal of Experimental Psychology: General*, 136(3):485–501.
- [15] Shepard, R. N. (1987). Towards a universal law of generalization for psychological science. *Science*, 237:1317–1323.
- [16] Stahl, I. (1993). Predicate invention in ILP—an overview. In *ECML 93: Proceedings of the European Conference on Machine Learning*, London, UK. Springer-Verlag.
- [17] Xu, F. and Tenenbaum, J. B. (2007). Word learning as Bayesian inference. *Psychological Review*, 114(2):245–272.