Inductive generalization relies on category representations

Shelbie L. Sutherland[1] · Andrei Cimpian[1]

Abstract The ability to take information learned about one object (e.g., a cat) and extend it to other objects (e.g., a tiger, a lion) makes human learning efficient and powerful. How are these inductive generalizations performed? Fisher, Godwin, and Matlen (2015) proposed a developmental mechanism that operates exclusively over the perceptual and semantic features of the objects involved (e.g., furry, carnivorous); this proposed mechanism does not use information concerning these objects' category memberships. In the present commentary, we argue that Fisher and colleagues' experiments cannot differentiate between their feature-based mechanism and its category-based competitors. More broadly, we suggest that any proposal that does not take into account the central role of category representations in children's mental lives is likely to mischaracterize the development of inductive generalization. The key question is not whether, but *how*, categories are involved in children's generalizations.

Keywords Concepts · Categories · Induction · Cognitive development

Categories are prominent in young children's mental lives. The average preschooler wants to know what *kind* of thing everything is; wonders constantly which toys are for *boys* or which are for *girls*; has questions about where *bats* live,

✉ Shelbie L. Sutherland
ssuther2@illinois.edu

[1] University of Illinois, Urbana, IL, USA

whether *cats* can see in the dark, and why *kids* have to go to school. Young children even make up categories when there are none, such as when they claim—not without self-interest—that "*Adams* don't have to take naps" or that "*Naomis* are not ticklish" (Gelman & Brandone, 2010; Gelman, Goetz, Sarnecka, & Flukes, 2008). Thus, children express a budding curiosity about categories of things early and often. Notably, the roots of this tendency to carve up the world into categories are present even before children learn how to talk, in the first year of life (e.g., Dewar & Xu, 2009; Ferry, Hespos, & Waxman, 2013; Keates & Graham, 2008). These aspects of young children's cognition will serve as the backdrop for our commentary on a new proposal concerning the process by which children generalize information.

If you find out that a particular cat has a raspy tongue, what other animals can you reasonably assume have a raspy tongue as well? Another cat? A dog? A tiger? A robin? These sorts of inductive generalizations enable efficient learning about the world; without them, our learning would be impossibly narrow and slow. It is thus important to understand the cognitive mechanism that enables children to perform inductive generalizations. Fisher, Godwin, and Matlen (2015) proposed that this mechanism does not rely on category representations, but rather only on raw featural similarity (see also Sloman, 1993). In what follows, we summarize their perceptual and representational similarity (PaRS) account, contrasting it with its category-based (CB) competitors. We then argue that Fisher and colleagues' studies in fact provide little support for the PaRS account over CB accounts. Finally, we conclude with a few thoughts about the centrality of category representations in early cognitive development and the seeming implausibility of accounts that ignore this core feature of the human mind.

Featural similarity versus category-based accounts of inductive generalization

What's the difference?

The PaRS and CB accounts of inductive generalization are summarized schematically in Fig. 1. The core distinction between them is the following: According to PaRS, the mechanism that computes generalizations operates entirely on the features of the entities involved, both perceptual and semantic.[1] Specifically, this mechanism computes the overall featural similarity between the object that serves as the base for generalization and the object that serves as the potential target. If this similarity is higher than some threshold, then the child generalizes from the base to the target.

In contrast, CB accounts suppose that children conceptualize the base and the target objects as members of particular categories (e.g., CAT)[2] rather than simply as collections of disparate features (e.g., furry, four-legged, aloof). Moreover, children's generalizations are hypothesized to rely in one way or another on this categorical construal. Although Fisher and colleagues (2015) don't make this point clear, CB theories are in fact fairly inclusive concerning the *types* of computations performed over the relevant categories. For example, reasoners may determine whether the base and target objects belong to the same category (e.g., are the base and the target both CATS?; Gelman & Markman, 1986; Keates & Graham, 2008). Alternatively, they might perform set inclusion computations on these categories (e.g., how broad do I have to go to find a category that encompasses both CATS and ROBINS?; Osherson, Smith, Wilkie, López, & Shafir, 1990). Reasoners might also compute the semantic similarity of the two categories, which could involve a comparison of their features in semantic memory (e.g., López, Gelman, Gutheil, & Smith, 1992; Osherson et al., 1990). Or they may invoke the causal–explanatory relationships relevant to the two categories to determine whether to generalize from the base to the target (e.g., what function might the raspy tongue serve for CATS, and might this function benefit ROBINS as well?; e.g., Barrett, Abdi, Murphy, & Gallagher, 1993; Gelman & Markman, 1986; McCarrell & Callanan, 1995; Murphy & Allopena, 1994). These various ways of manipulating category information are thought to be employed flexibly, depending on the characteristics of the stimuli, the context, the reasoner, and so on (e.g., Gelman & Markman, 1986; Medin, Coley, Storms, & Hayes, 2003).

---

[1] Fisher and colleagues defined *semantic* features as those that are not available to the senses at the time of the judgment but that have been acquired through experience at some point in the past.
[2] For the rest of the article, we will follow the typical convention of denoting categories in small capital letters.

What's the evidence?

Fisher and colleagues (2015) tested two predictions meant to differentiate between the PaRS and CB accounts. However, neither actually does, as we go on to explain. Their first prediction boils down to the following:

> **Prediction 1:** Greater similarity between the base and the target objects should be accompanied by more generalizations from the base to the target.

For example, when given the option of generalizing a property of a lamb (the base object) to either a sheep or a "lure" (the target objects), children should make more generalizations to the lure when it is a cow than when it is a clock. Because the cow is more similar to the base object (the lamb) than the clock is, it should be more successful in competing for children's generalizations with the other target object (the sheep). However, this prediction does not follow uniquely from the PaRS account. Any CB account in which reasoners can compute the semantic similarity between two categories (e.g., LAMB–COW, LAMB–CLOCK) would make the same prediction. (As far as we know, all existing CB accounts allow for such computations.) Thus, Fisher and colleagues' first prediction cannot differentiate between PaRS and CB theories of inductive generalization.

The second prediction tested by Fisher and colleagues (2015) can be paraphrased as follows:

> **Prediction 2:** Individual differences in executive function (e.g., working memory, inhibitory control) and semantic knowledge should correlate with individual differences in inductive generalization. That is, children with greater executive function capacities and richer semantic knowledge should perform more normatively appropriate generalizations (e.g., generalizing from an alligator to a crocodile rather than to a butterfly, or from a sea to an ocean rather than to an apple).

It is unclear to us what theory would wish to deny that, as children's cognitive capacities and knowledge develop, their reasoning becomes more sophisticated. With respect to CB accounts more specifically, there are many ways in which such cognitive changes might matter for children's generalizations. For instance, the computations children are hypothesized to perform over the relevant categories (see Fig. 1) require working memory and processing capacity, and thus greater resources of this sort would allow children to generate more adequate judgments. (We should note that working memory and other executive functions are also involved in understanding the experimenter's instructions, remembering the information one is reasoning about, and so on. Improvements in these aspects of children's participation as
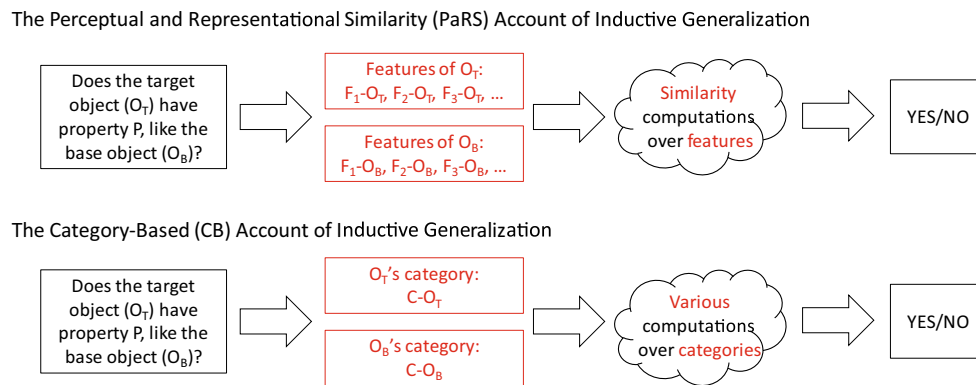
The Perceptual and Representational Similarity (PaRS) Account of Inductive Generalization

The Category-Based (CB) Account of Inductive Generalization

Fig. 1 Two accounts of inductive generalization. The major differences between them are depicted in red.

research subjects would also predict improved performance, even if there were no debates about the mechanisms of inductive generalization.) The PaRS model does not have exclusive claim to variables such as knowledge and cognitive resources. These factors are undoubtedly involved in any CB model of inductive generalization, as well.

In sum, Fisher and colleagues' (2015) studies provide no unique support for PaRS over its CB competitors. Neither semantic similarity computations (Prediction 1) nor individual differences in executive function and knowledge (Prediction 2) are the special province of the PaRS account. This is not to say that Fisher and colleagues' data are without value. These studies provide, to our knowledge, a first longitudinal investigation of inductive inferences. Moreover, the focus on individual differences in children's inductive generalizations, as well as on the factors that explain such differences (e.g., semantic knowledge, working memory), is innovative and likely to influence subsequent work on this topic. As such, Fisher and colleagues' studies mark an important advance in our descriptive, empirical knowledge of the development of inductive generalization, thereby also providing future theoretical accounts a more focused view of this phenomenon. However, these data do not in and of themselves provide unique support for the feature-based theory by which they were inspired, as we explained above.

In the final portion of this commentary, we step back and evaluate the very idea that categories—although omnipresent in children's cognitive lives more generally—would nevertheless be absent from their inductive generalizations.

How plausible is a category-less account of inductive generalization?

Children's proclivity for thinking in terms of categories is not just an anecdotal impression. Rather, the claim that category representations are available and influential from the earliest stages of life is supported by considerable empirical evidence. To begin with, even infants are able to represent categories per

se, not just collections of superficially similar objects, and use these categories in their reasoning (e.g., Booth, Waxman, & Huang, 2005; Dewar & Xu, 2009; Keates & Graham, 2008; Yin & Csibra, 2015). In addition, children begin to refer to categories in their speech as soon as they are able to string words together into multiword utterances (e.g., Gelman et al., 2008). Notably, children's category talk does not just mimic adults' category talk. The majority of children's utterances that express category content occur in self-initiated exchanges. (As a side note, the inference that certain linguistic structures [e.g., "Boys like trucks"] express thoughts about categories [e.g., BOYS, TRUCKS] is supported by decades of work in linguistics and philosophy [e.g., Carlson & Pelletier, 1995; Leslie, 2008], as well as in cognitive development [e.g., Brandone, Cimpian, Leslie, & Gelman, 2012]). Rather than being mere mimicry, children's category talk most likely reflects a deeper interest in gaining information at the category level—and, indeed, recent work has suggested that children are particularly eager to acquire knowledge about categories (Cimpian & Park, 2014; Cimpian & Petro, 2014; see also Martin & Ruble, 2004). Some researchers have even argued that the drive to learn about categories is a "built-in" feature of our cognitive systems, insofar as infants appear to be innately sensitive to certain communicative cues that signal the transmission of category-wide knowledge (for reviews, see Csibra & Gergely, 2009; Csibra & Shamsudheen, 2015). Once acquired, category-level information may be privileged in memory, as well. That is, children are better able to recall novel information about categories than to recall otherwise-identical information about noncategory sets (e.g., Cimpian & Erickson, 2012; Leslie & Gelman, 2012; Riggs, Kalish, & Alibali, 2014a, 2014b). This evidence, which is just a small portion of the evidence available in the literature, makes it clear that using categories is routine in young children's cognition.

In light of this fact, we need to examine the very idea that categories would not be involved in children's generalizations. How could it be that children (and people more generally) routinely think in terms of categories except when it comes

to generalizing information? Previous feature-based accounts explicitly acknowledged this tension and tried to resolve it. For instance, Sloman (1993) used a distinction between intuitive (i.e., low-effort) and logical (i.e., high-effort) ways of generalizing. His feature-based model (which is similar to PaRS) was intended to capture only the low-effort, intuitive pathway. We do not see this as a viable antidote, though, largely because reasoning with categories is not particularly effortful for children. For example, children find it easier to reason with categories (e.g., CATS) than with other broad sets (e.g., all cats), even though the latter are in fact much simpler from the viewpoint of formal semantics (e.g., Hollander, Gelman, & Star, 2002; Leslie & Gelman, 2012; Mannheim, Gelman, Escalante, Huayhua, & Puma, 2011; Tardif, Gelman, Fu, & Zhu, 2012; see also Hampton, 2012). That is, although the conditions under which claims about noncategory sets (e.g., that *all* cats have raspy tongues) are often trivial to spell out formally, children's ability to evaluate such claims lags behind their ability to evaluate comparable claims about categories (e.g., that CATS have raspy tongues), whose truth conditions are vastly more complex. This surprising contrast between the formal, in-principle complexity of category reasoning and children's early successes with it undermines claims that such reasoning is somehow effortful or unintuitive. In other words, category-based inductive generalization cannot be legitimately characterized as high-effort. Rather, categories are likely to be involved in the process by which everyday, intuitive generalizations are made, just as they are in so many other areas of child cognition.

Conclusion and questions for future work

Young children use categories in their everyday reasoning. Rather than trying to find ways of cordoning off category representations from the development of inductive generalization, it seems to us that a more productive way to move forward is to ask *how*—not whether—categories are involved in children's generalizations. Which of the category-based computations invoked by CB accounts (e.g., semantic similarity, set inclusion) are performed under which circumstances? Is there a hierarchy among these computations, whereby some are called on by default and others only as backup? Are there developmental differences in children's ability to perform these various computations? Which categorical information is retrieved from semantic memory under which circumstances? Are there developmental trends in the types of information retrieved (or retrieved most easily)? By answering questions such as these, we may be able to make meaningful progress toward understanding the development of inductive generalization.

References

Barrett, S. E., Abdi, H., Murphy, G. L., & Gallagher, J. M. (1993). Theory-based correlations and their role in children's concepts. *Child Development, 64,* 1595–1616.

Booth, A. E., Waxman, S. R., & Huang, Y. T. (2005). Conceptual information permeates word learning in infancy. *Developmental Psychology, 41,* 491–505. doi:10.1037/0012-1649.41.3.491

Brandone, A. C., Cimpian, A., Leslie, S. J., & Gelman, S. A. (2012). Do lions have manes? For children, generics are about kinds rather than quantities. *Child Development, 83,* 423–433.

Carlson, G. N., & Pelletier, F. J. (Eds.). (1995). *The generic book.* Chicago, IL: University of Chicago Press.

Cimpian, A., & Erickson, L. C. (2012). Remembering kinds: New evidence that categories are privileged in children's thinking. *Cognitive Psychology, 64,* 161–185. doi:10.1016/j.cogpsych.2011.11.002

Cimpian, A., & Park, J. J. (2014). Tell me about *pangolins*! Evidence that children are motivated to learn about kinds. *Journal of Experimental Psychology: General, 143,* 46–55.

Cimpian, A., & Petro, G. (2014). Building theory-based concepts: Four-year-olds preferentially seek explanations for features of kinds. *Cognition, 131,* 300–310.

Csibra, G., & Gergely, G. (2009). Natural pedagogy. *Trends in Cognitive Sciences, 13,* 148–153.

Csibra, G., & Shamsudheen, R. (2015). Nonverbal generics: Human infants interpret objects as symbols of object kinds. *Annual Review of Psychology, 66,* 689–710. doi:10.1146/annurev-psych-010814-015232

Dewar, K., & Xu, F. (2009). Do early nouns refer to kinds or distinct shapes? Evidence from 10-month-old infants. *Psychological Science, 20,* 252–257.

Ferry, A. L., Hespos, S. J., & Waxman, S. R. (2013). Nonhuman primate vocalizations support categorization in very young human infants. *Proceedings of the National Academy of Sciences, 110,* 15231–15235. doi:10.1073/pnas.1221166110

Fisher, A. V., Godwin, K. E., & Matlen, B. J. (2015). Development of inductive generalization with familiar categories. *Psychonomic Bulletin & Review.* doi:10.3758/s13423-015-0816-5. Advance online publication.

Gelman, S. A., & Brandone, A. C. (2010). Fast-mapping placeholders: Using words to talk about kinds. *Language Learning and Development, 6,* 223–240.

Gelman, S. A., Goetz, P. J., Sarnecka, B. W., & Flukes, J. (2008). Generic language in parent-child conversations. *Language Learning and Development, 4,* 1–31.

Gelman, S. A., & Markman, E. M. (1986). Categories and induction in young children. *Cognition, 23,* 183–209.

Hampton, J. A. (2012). Thinking intuitively: The rich (and at times illogical) world of concepts. *Current Directions in Psychological Science, 21,* 398–402. doi:10.1177/0963721412457364

Hollander, M. A., Gelman, S. A., & Star, J. (2002). Children's interpretation of generic noun phrases. *Developmental Psychology, 38,* 883–894.

Keates, J., & Graham, S. A. (2008). Category markers or attributes: Why do labels guide infants' inductive inferences? *Psychological Science, 19,* 1287–1293. doi:10.1111/j.1467-9280.2008.02237.x

Leslie, S. J. (2008). Generics: Cognition and acquisition. *Philosophical Review, 117,* 1–47.

Leslie, S. J., & Gelman, S. A. (2012). Quantified statements are recalled as generics: Evidence from preschool children and adults. *Cognitive Psychology, 64,* 186–214.

López, A., Gelman, S. A., Gutheil, G., & Smith, E. E. (1992). The development of category-based induction. *Child Development, 63,* 1070–1090.

Mannheim, B., Gelman, S. A., Escalante, C., Huayhua, M., & Puma, R. (2011). A developmental analysis of generic nouns in Southern Peruvian Quechua. *Language Learning and Development, 7,* 1–23.

Martin, C. L., & Ruble, D. (2004). Children's search for gender cues: Cognitive perspectives on gender development. *Current Directions in Psychological Science, 13,* 67–70.

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. K. (2003). A relevance theory of induction. *Psychonomic Bulletin & Review, 10,* 517–532. doi:10.3758/BF03196515

Murphy, G. L., & Allopena, P. D. (1994). The locus of knowledge effects in concept learning. *Journal of Experimental Psychology: Learning,* *Memory, and Cognition, 20,* 904–919. doi:10.1037/0278-7393.20.4.904

Osherson, D. N., Smith, E. E., Wilkie, O., López, A., & Shafir, E. (1990). Category-based induction. *Psychological Review, 97,* 185–200. doi:10.1037/0033-295X.97.2.185

Riggs, A. E., Kalish, C. W., & Alibali, M. W. (2014a). Property content guides children's memory for social learning episodes. *Cognition, 131,* 243–253. doi:10.1016/j.cognition.2014.01.004

Riggs, A. E., Kalish, C. W., & Alibali, M. W. (2014b). When you've seen one, have you seen them all? Children's memory for general and specific learning episodes. *Developmental Psychology, 50,* 1653–1659. doi:10.1037/a0036130

Sloman, S. A. (1993). Feature-based induction. *Cognitive Psychology, 25,* 231–280. doi:10.1006/cogp.1993.1006

Tardif, T., Gelman, S. A., Fu, X., & Zhu, L. (2012). Acquisition of generic noun phrases in Chinese: Learning about lions without an "-s.". *Journal of Child Language, 39,* 130–161.

Yin, J., & Csibra, G. (2015). Concept-based word learning in human infants. *Psychological Science, 26,* 1316–1324. doi:10.1177/0956797615588753