



Symbolic Versus Associative Learning

John E. Hummel

Department of Psychology, University of Illinois, Urbana-Champaign

Received 8 January 2010; received in revised form 11 January 2010; accepted 12 January 2010

Abstract

Ramscar and colleagues (2010, this volume) describe the “feature-label-order” (FLO) effect on category learning and characterize it as a constraint on symbolic learning. I argue that FLO is neither a constraint on symbolic learning in the sense of “learning elements of a symbol system” (instead, it is an effect on nonsymbolic, association learning) nor is it, more than any other constraint on category learning, a constraint on symbolic learning in the sense of “solving the symbol grounding problem.”

Keywords: Symbol systems; Relational representations; Relational learning; Symbolic learning; Symbol grounding

1. Introduction

Ramscar et al. (2010) present the “feature-label-order” (FLO) effect, in which subjects are better at learning to predict an object’s label from its features than the reverse. The authors’ primary concern is the FLO effect as a constraint on how people learn the referents of concrete nouns (object categories and colors). They conclude that FLO serves as a constraint on the symbol grounding problem and on symbolic learning more broadly.

The authors are either attempting to provide a nonsymbolic account of symbolic learning or they are raising a nonsymbolic phenomenon to the level of a symbolic one. In either case, their FLO effect fails to rise to the level of a constraint on symbolic learning. The authors’ arguments are symptomatic of a broader misunderstanding within cognitive science about the nature of symbol systems. I begin by reviewing what I see as the key differences between nonsymbolic and symbolic representations and processes, drawing on an analogy to first- and second-order logics, respectively. I follow this review with a discussion of why

Correspondence should be sent to John E. Hummel, Department of Psychology, University of Illinois, Urbana-Champaign, IL 61820. E-mail: jehummel@illinois.edu

FLO is not a constraint on symbolic learning, in general, and then with a critique of the authors' arguments about the role of FLO as a constraint on symbol grounding.

2. Symbolic and nonsymbolic systems

Simply being a representation is not by itself sufficient to render a representation symbolic, and simply being a process (e.g., learning) over one or more representations is not sufficient to render that process symbolic (Cheng, 1997; Chomsky, 1959; Deacon, 1997; Fodor & Pylyshyn, 1988; Holyoak & Thagard, 1995; Hummel & Holyoak, 1997, 2003; Penn, Holyoak, & Povinelli, 2008; Rumelhart, McClelland, and The PDP Research Group, 1986). As pointed out by Peirce (1879, 1903; see also Deacon, 1997), representations reside in a hierarchy consisting of *icons* (which denote things by virtue of resembling them), *indices* (such as words, which denote things without resembling them), and *symbols* (icons or indices that, by virtue of the system of which they are a part, can be combined and recombined to represent different kinds of relations among things). Key here is the notion of *different kinds of relations*.

2.1. Symbolic (explicitly relational) systems

Consider natural language, a symbolic (or, equivalently, *explicitly relational*; Hummel & Holyoak, 1997, 2003) representational system (Chomsky, 1959; Jackendoff, 1997). This system is explicitly relational in the sense that it represents relations as explicit entities (i.e., as nouns, verbs, adjectives, adverbs, prepositions, or as complete phrases or clauses) that take arguments. John can *love* Mary, or be *taller than* Mary, or be the *father of* Mary, or all of the above. The *vocabulary* of relations in a symbol system is open-ended (i.e., as elaborated below, *variabilized*), and relations can take other relations as arguments (e.g., *Mary knows John loves her*). More importantly, not only can John love Mary, but Sally can love Mary, too, and in both cases it is the *very same* "love" relation. The *love* relation is promiscuous, free to take different arguments on different occasions. Its arguments enjoy the same promiscuity: The Mary that John loves can be the very same Mary that is loved by Sally. This capacity for dynamic recombination is at the heart of a symbolic representation and is not enjoyed by nonsymbolic representations.

2.2. Representing relations versus explicitly relational representations

Relations are central to symbolic representations, but many things represent relations without being symbolic. The value of V_{ij} in Rescorla and Wagner's (1972) model of associative learning represents an associative relation between unconditioned stimulus (US) i and conditioned stimulus (CS) j . Similarly, connection weight w_{ij} in a connectionist network represents a coactivation relation over units i and j (roughly " i should be active when j is" when w_{ij} is positive, or " i should be inactive when j is active" when w_{ij} is negative), and the links in a causal Bayes network represent causal relations between the nodes they connect. Even the output of a ganglion cell in the retina represents a relation (a contrast)

between two adjacent locations on the retina. But none of these representations is explicitly relational (i.e., symbolic) in the manner described above.

The reasons are twofold. First is that the system (e.g., connectionist network, Rescorla–Wagner model, or Bayes net) has a very limited vocabulary of relations, usually consisting of only a single relation: V_{ij} always represents an associative relation; w_{ij} always represents a coactivation relation (in a connectionist model) or a causal relation (in a Bayes net); and the output of a ganglion cell always represents a particular local contrast on the retina. Second, and more fundamentally, these relations are not free to vary in their “arguments” (indeed, it is a misnomer to speak of their “arguments” at all): A given ganglion cell always represents the contrast at the same location on the retina (move to a new location and a different ganglion cell will represent the contrast there); for a given value of i and j , V_{ij} always represents the associative strength between the same US and the same CS (change the US [the value of i] or the CS [the value of j] and not only is the value of V_{ij} likely to change, the *variable itself* will have changed: $V_{1,2}$ is a different variable than $V_{1,3}$ or $V_{3,4}$ —that is, although you and I recognize that $V_{1,2}$ and $V_{3,4}$ represent the same relation between different arguments, the system itself has absolutely no appreciation of this fact); the same applies to the various w_{ij} in a connectionist or Bayes net. The “relations” in a nonsymbolic system lack the promiscuity of the relations (and arguments) in a symbol system.

2.3. *Why promiscuity matters*

It is the promiscuity of symbolic representations—the ability to bind the same relation to different arguments and the same arguments to different relations dynamically—that gives symbol systems the expressive (and therefore computational) power they enjoy over nonsymbolic (e.g., associative) systems (see, e.g., Deacon, 1997; Dumas, Hummel, & Sandhofer, 2008; Fodor & Pylyshyn, 1988; Gentner, 1983; Hummel & Holyoak, 1997, 2003). It is what makes humans capable of being engineers, scientists, and users of natural languages whereas our closest cousins in the animal kingdom are not (Penn et al., 2008). And it is what makes learning symbolic representations vastly more difficult than learning nonsymbolic representations (such as associations). A simple association, like those learned by the Rescorla–Wagner model, can be learned simply by tabulating the co-occurrence of various things (e.g., CSs and USs or objects and labels) in the world (which is precisely what the Rescorla–Wagner model does). A relation like *taller than*, in contrast, does not appear in the form of such simple co-occurrences: John can be taller than Mary, a beer bottle is taller than a beer can, and an apartment building is taller than a house. But in what sense, other than being taller than something, is John like a beer bottle or an apartment building? Making matters worse, Mary is taller than the beer bottle and the house is taller than John. Precisely because of their promiscuity, relational concepts defy learning in terms of simple associative co-occurrences (Dumas et al., 2008).

2.4. *Representational systems and logics*

The ability of symbol systems to variabilize relations whereas nonsymbolic systems cannot renders the difference between nonsymbolic and symbolic systems analogous to the

difference between first- and second-order logics.¹ This analogy can be illustrated by reference to the Peano Axioms (i.e., the axioms Italian mathematician Giuseppe Peano used to axiomatize integer arithmetic). A first-order logic makes it possible to variablize (i.e., quantify over) the arguments of functions or relations, but the functions or relations themselves must be constants (much like relations must be constants in nonsymbolic systems). For example, one of the Peano axioms states that for every natural number, x , $x = x$. The variablization in this statement is the “for every,” the relation is “=” and the argument is “ x .” Note that the variablization refers to the argument, “ x ,” not the relation, “=.” For this reason, this axiom is expressible in first-order logic, as are the majority of the Peano axioms. However, it is well known that first-order logic is not powerful enough to axiomatize integer arithmetic. For this, we need second-order logic. The difference between first- and second-order logic is that the latter makes it possible to variablize not only arguments but also functions and relations (much like a symbolic system can variablize relations). The ability to variablize relations makes it possible to state the *induction axiom*: If ϕ is a unary predicate such that (a) $\phi(0)$ is true and (b) for every natural number, x , the truth of $\phi(x)$ implies the truth of $\phi(S(x))$ (where $S(x)$ is simply the successor function, $x + 1$), then $\phi(x)$ is true for all x . Without the induction axiom, it is not possible to prove theorems about infinite sets of numbers, that is, about arithmetic itself.² Note that what gives the induction axiom the power to support such generalization is precisely the same thing that renders it an axiom of second-order logic, namely the fact that it variablizes not only arguments, x , but also functions, ϕ , over those arguments. It is this same capacity to variablize relations that gives symbol systems their expressive power.

3. Why FLO is not a constraint on symbolic learning

Just as second-order logics permit qualitatively different kinds of inferences than do first-order logics (a difference so profound that it is possible to axiomatize integer arithmetic using one but not the other), symbol systems permit qualitatively different kinds of processing (such as learning and inference) than do nonsymbolic systems (a difference so profound that our symbolic species dominates the planet, whereas our nonsymbolic cousins do not). So the difference between symbolic and nonsymbolic systems is real and important. But why does this fact imply that the FLO effect is not an effect on symbolic learning?

The answer begins with the fact that the FLO effect was predicted by the decidedly nonsymbolic Rescorla–Wagner model (see Cheng, 1997). FLO is a constraint on (nonsymbolic) association learning. Inasmuch as association learning is a component of symbolic learning (and there is a reason to believe it is; Dumas et al., 2008), it may also be a constraint on symbolic learning. But in this case, calling it a “constraint on symbolic learning” is a bit like calling a constraint on addition a “constraint on the integral calculus.” The “larger sum law” (which I just made up and which states that the sum of two positive nonzero addends is larger than either addend) is a “law” that belongs in the purview of addition; trying to publish said law as “a constraint on the integral calculus” would, while not exactly wrong in the very strongest sense, nonetheless be misleading and possibly harmful,

especially to the extent that anyone in the article's target audience struggles to understand the difference between addition and integration: The "larger sum law" really is a law of addition, not integration. In the same way, the FLO effect (if true) is a constraint on associative, not symbolic learning.

But it gets worse. Even though first-order logic is, in some sense, a component of second-order logic, not all constraints that apply to first-order logic apply to second-order logic. Most obviously, the constraint "thou shall not quantify over predicates" applies to first but not second-order logic. Less obviously, Kurt Gödel proved the completeness of first-order logic. But it would be a mistake to assume that this law also applies to second-order logic: The same Kurt Gödel also proved that no logic powerful enough to axiomatize integer arithmetic (i.e., second- and higher-order logics) can be both complete and consistent. The point is that even if x is a component of y (as first-order logic is a component of second-order logic, and as association learning may be a component of symbolic learning) not everything that is true of x will necessarily be true of y . Thus, even if Ramsar et al. (2010) demonstrated that FLO is a constraint on a kind of association learning, and even if association learning can be plausibly argued to be a component of symbolic learning, it does not follow that FLO is also a constraint on symbolic learning.

Where does this leave us? We know that FLO may serve as a constraint on a kind of associative learning. We know that at least some kind of associative learning is likely to be a component of symbolic learning. Given this, either (a) FLO is a constraint on both associative and symbolic learning, in which case it is better described as a "constraint on learning" than as a "constraint on symbolic learning," or else (b) it is a constraint on associative learning but not symbolic learning, in which case the statement "FLO is a constraint on symbolic learning" is simply wrong. Possibilities (a) and (b) are exhaustive and neither is consistent with the claim that FLO is a constraint on symbolic learning.

3.1. Symbol grounding

What about the more modest claim that FLO is a constraint on symbol grounding (the authors' primary claim)? Even this more restricted interpretation of the authors' argument runs into difficulties. The authors demonstrated the FLO effect using concrete concepts: feature-based object categories in Experiment 1 and colors in Experiment 2. And their data suggest that for these concrete, nonrelational concepts, FLO may serve as a constraint on learning to associate labels with them.

However, many concepts for which we have names are much more relational in nature (see, e.g., Gentner & Kurtz, 2005; Jackendoff, 1997). These include verbs (Gentner, 1978, 1981), adjectives (such as *taller than*), adverbs (such as *very*), prepositions, and even nouns, in the case of *role-governed* categories (Markman & Stilwell, 2001)—categories such as *friend*, *barrier*, and *container*, which are defined by the relational role in which an object participates, rather than the literal features of the object itself. A nonsymbolic system such as the Rescorla–Wagner model cannot even represent such concepts, much less learn anything about them.

Ironically, the best hope for the authors' argument may be that the "symbol grounding problem," as described in their article, does not require symbolic representations at all. The

FLO effect is at best a constraint on how we learn to associate labels—*indices*, in the parlance of Peirce (1879, 1903)—with category members. But recall that not all indices are symbols (as such the “symbol grounding problem” is perhaps better called the “index grounding problem”). An index is only a symbol if it is part of a system that allows it to combine promiscuously with other symbols in order to express different kinds of relations. There is little doubt that many nonhuman animals have something like indices—as evidenced, for example, by my dog’s ability to retrieve her toys by name (“Go get your squirrel!”). But this fact does not render their mental representations symbolic, and it is not what’s hard about “symbolic learning”: The hard part of symbolic learning is dissociating indices from their specific features and contexts to render them promiscuous. My dog has learned, perhaps via mechanisms very much like those the authors propose, to go get her squirrel; but from this, it does not follow that I will ever be able to count on her to go find me a container, a barrier, or a friend (or even to find me a toy squirrel in a different context), much less to engineer me a bridge—activities that depend on genuinely symbolic representations.

The authors’ data suggest that FLO may serve as a constraint on learning to associate indices with their (concrete) referents, and it is in this context that it may be called a constraint on the index grounding problem. Or to put it another way, it is a constraint on how we learn categories (under speeded conditions, when the categories have a complex, nonfamily-resemblance structure; see the Methods section of Ramscar et al.’s Experiment 1). But this fact does not raise it to the level of a “constraint on symbolic learning” any more than it does any other constraint on category learning. And even if FLO were a major constraint on index grounding, it is by no means clear that this fact would contribute much to our understanding of how the mind solves the difficult problem of symbolic learning: Although “symbol grounding” is a hard *philosophical* problem, that it is not what’s hard about symbolic learning as a *psychological/biological* problem as evidenced by the fact that many animals appear to have something like categories (i.e., grounded indices), whereas of the species investigated to date, only humans appear to have genuinely symbolic representational systems (see Penn et al., 2008; see Doumas et al., 2008, on the difficulty of symbolic learning).

4. Conclusion

Feature-label-order is a comparatively minor effect on category learning, predicted by a 40-year-old model of associative learning in rats, demonstrated under cognitively limiting circumstances, which the authors have gone through extensive verbal acrobatics to imbue with meaning it does not have.

Notes

1. In fact, nonsymbolic systems might be more accurately compared with propositional logics (which are weaker than first-order logics), but for the purposes of this example, it is sufficient to “promote” them to the level of a first-order logic. That is, even if we

assume that nonsymbolic systems are more powerful than they actually are, their weakness relative to symbolic systems can be illustrated by analogy to the difference between first- and second-order logics.

2. It is difficult to overstate the importance of the induction axiom. Without it, it would be impossible to prove such a simple statement as “for all x , $S(x) > x$.” Instead, you would have to prove it for all natural numbers *individually*: Having proven it for 1, 2, 3, ... up to 1,000,000, you would not know whether it is true for 1,000,001. The induction axiom, which requires the power of a second-order logic, is thus absolutely essential for axiomatizing integer arithmetic.

Acknowledgment

Preparation of this article was supported by AFOSR grant FA9550-07-1-0147.

References

- Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, *104*, 367–405.
- Chomsky, N. (1959). Review of verbal behavior, by B.F. Skinner. *Language*, *35*, 26–57.
- Deacon, T. W. (1997). *The symbolic species: The co-evolution of language and the brain*. New York: W. Norton and Company.
- 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.
- Fodor, J. A., & Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: A critical analysis. In S. Pinker & J. Mehler (Eds.), *Connections and symbols* (pp. 3–71). Cambridge, MA: MIT Press.
- Gentner, D. (1978). On relational meaning: The acquisition of verb meaning. *Child Development*, *49*, 988–998.
- Gentner, D. (1981). Some interesting differences between nouns and verbs. *Cognition and Brain Theory*, *4*, 161–178.
- Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, *7*, 155–170.
- Gentner, D., & Kurtz, K. J. (2005). Relational categories. In W. K. Ahn, R. L. Goldstone, B. C. Love, A. B. Markman, & P. W. Wolff (Eds.), *Categorization inside and outside the lab* (pp. 151–175). Washington, DC: American Psychological Association.
- Holyoak, K. J., & Thagard, P. (1995). *Mental leaps: Analogy in creative thought*. Cambridge, MA: MIT Press.
- Hummel, J. E., & Holyoak, K. J. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review*, *104*, 427–466.
- Hummel, J. E., & Holyoak, K. J. (2003). A symbolic-connectionist theory of relational inference and generalization. *Psychological Review*, *110*, 220–264.
- Jackendoff, R. (1997). *The architecture of the language faculty*. Cambridge, MA: MIT Press.
- Markman, A. B., & Stilwell, C. H. (2001). Role-governed categories. *Journal of Experimental and Theoretical Artificial Intelligence*, *13* (4), 329–358.
- Peirce, C. S. (1879, 1903). Logic as semiotic: The theory of signs. In J. Buchler (Ed.), *The philosophical writings of Peirce (1955)* (pp. 98–119). New York: Dover Books.
- Penn, D. C., Holyoak, K. J., & Povinelli, D. J. (2008). Darwin’s mistake: Explaining the discontinuity between human and nonhuman minds. *Brain and Behavioral Sciences*, *31*, 109–178.

- Ramscar, M., Yarlett, D., Dye, M., Denny, K., & Thorpe, K. (2010). The effects of feature-label-order and their implications for symbolic learning. *Cognitive Science*. doi: 10.1111/j.1551-6709.2009.01092.x
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64–99). New York: Appleton-Century-Crofts.
- Rumelhart, D. E., & McClelland, J. L., & The PDP Research Group (1986). *Parallel distributed processing: Explorations in the microstructure of cognition*, Vol. 1. Cambridge, MA: MIT Press.