frontal gyri (IFG) showing sensitivity to each component process of analogical reasoning. Separate regions that showed exclusive sensitivity to each component process were also identified within IFG. In addition, the degree of activation increase in the right ventral IFG during trials in which participants had to integrate three relations (compared to one) was greater for individuals whose performance accuracy was higher.

Although the above studies do not directly deal with the development of analogy during childhood, they do clearly demonstrate several component processes involved in analogical reasoning that are dependent on prefrontal cortex, an area of the brain that actively develops throughout childhood (Diamond 2002). In an effort to explore these processes directly in children, Richland et al. (2006) developed a scene-analogy task manipulating both relational complexity and featural distraction. Even 3-year-olds could solve simple (one-relation, no-distraction) problems, but they had difficulty if the problem required integration of multiple relations or ignoring a featurally similar object. Similarly, Wright et al. (2007) performed an fMRI study with children using another semantically rich visual analogy task, and found that brain activation in areas associated with relational integration was the best predictor of analogy performance. Wright et al. also found that these areas, which are not associated with semantic retrieval (Bunge et al. 2005), become more and more engaged over the same time period in which children dramatically improve in their ability to solve more relationally complex problems (Richland et al. 2006).

We are highly sympathetic with the target article's efforts to computationally model the development of analogy, and we certainly don't dispute the importance of relational knowledge in development. However, we believe that a successful model of development must (1) explain how knowledge representation and process constraints interact to produce the changes in analogy observed in children, including increases in ability to perform relational integration and resist featural distraction; and (2) explain how an architecture consistent with the demands of adult analogical reasoning develops. Unfortunately, the connectionist model described in the target article does not meet these requirements. In contrast, Morrison and collaborators have used LISA (Learning and Inference with Schemas and Analogies; Hummel & Holyoak 1997; 2003), a neurally plausible model of analogical reasoning, to successfully simulate many of the developmental and neuropsychological results discussed in this commentary (e.g., Morrison et al. 2004; 2006; Viskontas et al. 2004).

We believe that the development of analogical reasoning is best conceptualized as an equilibrium between children's relational knowledge and their current processing ability. As children mature, their prefrontal cortices more efficiently implement WM and thereby can process more complex analogies. However, more efficient relational representations can impose fewer processing demands at any given age, which is why a child who becomes an expert in a given domain can show rapid progress even though the child's WM system has not improved (Morrison et al. 2007). This framework can account for the observed changes in children's analogical reasoning, as well as subsequent changes in analogy during normal and abnormal human aging. It can also be simulated in symbolic-connectionist models of relational learning and reasoning (e.g., Doumas et al. 2008; Hummel & Holyoak 1997; 2003).

ACKNOWLEDGMENTS

Generous support for the authors was provided by the Northwestern University Mechanisms of Aging and Dementia Training Grant funded by the National Institute of Aging (2T32AG020506; RGM), the Office of Naval Research (SBIR OSD03-DH07; RGM), and the Kwanjeong Educational Foundation (SC).

Relational priming plays a supporting but not leading role in adult analogy-making

doi:10.1017/S0140525X08004627

Alexander A. Petrov

Department of Psychology, Ohio State University, Columbus, OH 43210.

apetrov@alexpetrov.com

http://alexpetrov.com

Abstract: Leech et al.'s analysis adds to an emerging consensus of the role of priming in analogy-making. However, their model cannot scale up to adult-level performance because not all relations can be cast as functions. One-size-fits-all accounts cannot capture the richness of analogy. Proportional analogies and transitive inferences can be made by nonstructural mechanisms. Therefore, these tasks do not generalize to tasks that require structure mapping.

Leech et al. argue forcefully that adult-level models of analogy-making must make contact with the developmental constraint. This argument cuts both ways: Developmental models must also make contact with adult-level capability. We argue that although relational priming does play a role in adult analogical reasoning, it does not play the leading role that Leech et al. suggest.

Relational priming. The role of priming in analogical reasoning is well documented empirically (e.g., Kokinov 1990; Schunn & Dunbar 1996). It also features prominently in several models, including Associative Memory-Based Reasoning (AMBR) (Kokinov 1994; Kokinov & Petrov 2001) and Copycat (French 1995; Hofstadter 1984; Mitchell 1993). All of these models implement priming as residual activation. The present proposal thus adds to an emerging consensus of the importance of priming and of its underlying mechanism.

Not all relations can be cast as functions. Leech et al. claim that "for the purposes of analogy it may be sufficient to conceptualize relations as transformations between items" (sect. 2.2, para. 2). The main idea is to cast each binary relation $R(a,b)$ as an equivalent univariate function[1] $b = F_R(a)$. The model uses hand-coded representations, *rep*, such that $rep(F_R(a)) = rep(a) + F_R(a)) = rep(a) + rep(R)$. The authors argue this is beneficial because "relations do not have to be represented explicitly, avoiding the difficulties of learning explicit structured representations" (sect. 5.1.1, para. 1). However, this benefit comes at the cost of rendering the model incapable of scaling up to adult-level performance.

The problem is that a relation can be cast as a function only if it is deterministic: that is, if for each a there is precisely one b that satisfies $R(a,b)$ (Halford et al. 1998). Many important relations violate this condition. Consider the transitive inference task: $taller(Ann,Beth)$, $taller(Beth,Chris) \rightarrow taller(Ann,Chris)$. Now, if the relation $taller(a,b)$ is cast as a function $b = shrink(a)$, the query $shrink(Ann) = ?$ becomes ambiguous. There are techniques for supporting nondeterministic functions in connectionist networks (e.g., Hinton & Sejnowski 1986) that can be incorporated into the model. However, the priming account faces a deeper challenge: Why should *Chris* be produced as the answer to the above query after the system has been primed with $Beth = shrink(Ann)$?

Many relationships in the world are indeed near-deterministic transformations such as *bread* → *cut bread*. It is an important developmental constraint that young children find such regular, familiar relations easier to deal with (e.g., Goswami & Brown 1989). These strong environmental regularities shape coarse-coded distributed representations that can support generalization and inference (Cer & O'Reilly 2006; Hinton 1990; Rogers & McClelland 2004; St. John & McClelland 1990). The target article demonstrates the utility of relational priming in these cases. However, there are also relationships such as *left of* that are quite accidental and changeable. To process them, the brain relies on sparse conjunctive representations (McClelland et al. 1995) that do not support priming well. Finally, adult-level analogies involve higher-order relations and nested propositions (Gentner 1983). Their brain realization is an active research topic (e.g., Smolensky & Legendre 2006). One promising approach relies on dynamic gating in the basal ganglia and prefrontal cortex (O'Reilly 2006; Rougier et al. 2005). Priming does play a role in these gated networks, but the critical functionality rests on other mechanisms.

The role of mapping. Proportional analogies are often presented in a multiple-choice format (e.g., Goswami & Brown 1989; 1990). An important limitation of the priming model is that its activation dynamics is not influenced by the available responses. The network simply produces an output pattern and stops. Then some unspecified control mechanism compares this pattern to the response representations. The limitations of this approach can be demonstrated by analogies with identical premises but different response sets, as illustrated in Figure 1. As Leech et al. argue in Figure 11 of the target article, the model should select response R2 when the choices are R1 and R2. Arguably, it should select response R3 when the choices are R1 and R3. To do this, the model must produce a pattern that is less similar to *rep*(*R1*) than it is to both *rep*(*R2*) and *rep*(*R3*). This seems to contradict the reasonable assumption that *rep*(*R1*) lies between *rep*(*R2*) and *rep*(*R3*) because of the intermediate size of R1.

Examples such as this highlight the role of mapping in analogy-making. The most important contribution of the target article, in our opinion, is to lay bare that a model (or a child or an ape) lacking any mapping capabilities can still perform proportional analogies quite well. The bold claim that "explicit mapping is no longer *necessary* for analogy to occur, but instead describes a subset of analogies" (sect. 5.4, para. 6) is a terminological matter. The take-home lesson for us is that proportional analogies can be solved by nonstructural means and thus cannot represent the class of analogies for which mapping *is* necessary.

The "psychologist's fallacy." This alerts us to a variant of the *psychologist's fallacy* wherein experimenters confuse their own understanding of a phenomenon with that of the subject (Oden et al. 2001). Proportional analogies can be solved by structure mapping; they are also solved at above-chance levels by many 4-year-olds. Still, it does not follow that "the ability to reason by analogy is present by at least age four" (Goswami 2001, p. 443), not if this ability is understood to imply structure mapping.

The transitive inference task is another case in point. It has been argued that this task is more complex than proportional analogies (Halford et al. 1998; Maybery et al. 1986). And yet even pigeons and rats can make transitive inferences (Davis 1992; Van Elzakker et al. 2003; von Fersen et al. 1991). Does that mean that the ability to reason by analogy is present in pigeons and rats? No, it means that transitive inferences can be made by nonstructural mechanisms (Frank et al. 2003). Human adults can make such inferences by verbal and nonverbal strategies that can be dissociated (Frank et al. 2005; 2006).

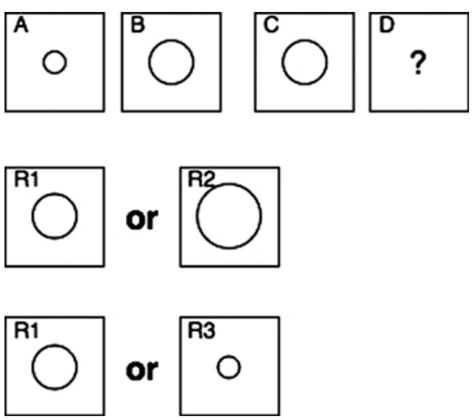

Figure 1 (Petrov). Demonstration of the importance of the available responses in a proportional analogy. Different response sets (R1–R2 vs. R1–R3) produce different analogies when paired with the same premises (A:B::C:?). Compare with Figure 11 in the target article.

Conclusion. The field can no longer treat analogy-making as a uniform skill. We need to identify the computational demands of analogies of different kinds, explicate the various strategies available for solving them, and design appropriate controls to discriminate among the strategies. Only then would developmental comparisons be meaningful. Relational priming is indeed a point of developmental continuity. However, it hardly constitutes a foundation strong enough for the formidable weight of adult analogical reasoning. After all, "it is probably safe to say that any program capable of doing analogy-making in a manner truly comparable to human beings would stand a very good chance of passing the Turing Test" (French 2002, p. 204).

ACKNOWLEDGMENTS

I thank Dave Noelle, Jeremy Reynolds, Jonathan Cohen, Mike Frank, Randy O'Reilly, and Todd Braver for their insightful discussions of these ideas and their feedback on the text.

NOTE

1. We use the standard predicate-calculus term *function* instead of *transformation*.

Abstract analogies not primed by relations learned as object transformations

doi:10.1017/S0140525X08004639

Steven Phillips

Neuroscience Research Institute, National Institute of Advanced Industrial Science and Technology (AIST), Tsukuba, 305-8568, Japan.

steve@ni.aist.go.jp

http://staff.aist.go.jp/steven.phillips

Abstract: Analogy by priming learned transformations of (causally) related objects fails to explain an important class of inference involving abstract source-target relations. This class of analogical inference extends to ad hoc relationships, precluding the possibility of having learned them as object transformations. Rather, objects may be placed into momentarily corresponding, symbolic, source-target relationships just to complete an analogy.

A glaring concern with Leech et al.'s "relations as transformations" account of analogy is the amount of training needed to attain a capacity for analogical inference. Adults reach a stage in development where analogical inference extends to ad hoc relationships outside the sphere of prior experience. Modeling this capacity is a problem for common feed-forward and simple recurrent networks, which rely on stimulus-driven response-error correction (Phillips 1999; 2000); and for similar reasons, this level of development is unreachable with the sort of connectionist model proposed in the target article. The analogizer cannot prepare in advance all possible transformations that could be primed. Moreover, any degree of generalization afforded to the model via similarity-based transformation is thwarted by analogies demanding transformations inconsistent with previous tasks.

Learning set transfer (Kendler 1995) or relational schema induction (Halford et al. 1998) involves testing participants on a series of stimulus-response tasks having a common structure (e.g., *transverse patterning*), where each task instance consists