THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Developing structured representations

Citation for published version:

Doumas, LAA & Richland, LE 2008, 'Developing structured representations', *Behavioral and Brain Sciences*, vol. 31, no. 04, pp. 384-385. <https://doi.org/10.1017/S0140525X08004536>

Digital Object Identifier (DOI):

[10.1017/S0140525X08004536](https://doi.org/10.1017/S0140525X08004536)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

Behavioral and Brain Sciences

Publisher Rights Statement:

© Doumas, L. A. A., & Richland, L. E. (2008). Developing structured representations. *Behavioral and Brain Sciences*, 31(04), 384-385. [10.1017/S0140525X08004536](https://doi.org/10.1017/S0140525X08004536)

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Analogy as relational priming: A developmental and computational perspective on the origins of a complex cognitive skill

Robert Leech

*School of Psychology, Birkbeck, University of London, London, WC1E 7HX,
United Kingdom*
r.leech@bbk.ac.uk
<http://www.alphalab.psyc.bbk.ac.uk/RobertLeech.html>

Denis Mareschal

*School of Psychology, Birkbeck, University of London, London, WC1E 7HX,
United Kingdom*
d.mareschal@bbk.ac.uk
<http://www.bbk.ac.uk/psyc/staff/academic/dmareschal>

Richard P. Cooper

*School of Psychology, Birkbeck, University of London, London, WC1E 7HX,
United Kingdom*
r.cooper@bbk.ac.uk
<http://www.bbk.ac.uk/psyc/staff/academic/rcoper>

Abstract: The development of analogical reasoning has traditionally been understood in terms of theories of adult competence. This approach emphasizes structured representations and structure mapping. In contrast, we argue that by taking a developmental perspective, analogical reasoning can be viewed as the product of a substantially different cognitive ability – relational priming. To illustrate this, we present a computational (here connectionist) account where analogy arises gradually as a by-product of pattern completion in a recurrent network. Initial exposure to a situation primes a relation that can then be applied to a novel situation to make an analogy. Relations are represented as transformations between states. The network exhibits behaviors consistent with a broad range of key phenomena from the developmental literature, lending support to the appropriateness of this approach (using low-level cognitive mechanisms) for investigating a domain that has normally been the preserve of high-level models. Furthermore, we present an additional simulation that integrates the relational priming mechanism with deliberative controlled use of inhibition to demonstrate how the framework can be extended to complex analogical reasoning, such as the data from explicit mapping studies in the literature on adults. This account highlights how taking a developmental perspective constrains the theory construction and cognitive modeling processes in a way that differs substantially from that based purely on adult studies, and illustrates how a putative complex cognitive skill can emerge out of a simple mechanism.

Keywords: Analogical reasoning; cognitive development; connectionism; relational priming; transformation similarity

1. Introduction

“Analogy lies at the core of human cognition,” as Holyoak et al. (2001, p. 2) point out. Analogies underlie creative thought and problem solving, and as such are implicated in virtually all aspects of human life. Analogies are found in science (e.g., comparing an atom with the solar system), in politics (e.g., the first President Bush comparing Saddam Hussein with Adolf Hitler), and in everyday living. It is not altogether surprising, then, that analogy is an equally powerful force in cognitive development. Children use analogy to extend their knowledge about the biological, physical, and psychological world and to solve problems (e.g., Brown & Kane 1988; Holyoak et al. 1984; Inagaki & Hatano 1987; Pauen & Wilkening

1997). Spontaneous analogies have been observed in very young children (Pauen & Wilkening 1997; Tunteler & Resing 2002), and there is even some evidence that infants are able to reason by analogy from around their first birthday (Chen et al. 1997).

Given its importance to cognition, it is equally unsurprising that analogy has been the focus of a number of detailed theoretical accounts, many of which have been implemented as working computational models (for a recent review, see French 2002). The majority of this work has focused on adult reasoning, and existing developmental accounts are largely adaptations of adult models (e.g., Gentner et al. 1995). From a developmental perspective, these accounts are wanting in that they posit specific mechanisms (e.g., structure-mapping) with no plausible

explanation as to how such mechanisms arise during cognitive development.

The work presented here is an attempt to fill this gap by providing a theory of the emergence of analogical reasoning abilities. The theory suggests that basic analogical abilities may arise from the normal functioning of a memory system as its domain knowledge increases. Importantly, no special additional mechanisms are required to deal with simple analogies. Rather, the ability to generate and use such analogies is argued to arise from the priming of relations that hold between terms in the analogy. The account thus illustrates how a putative complex skill could emerge out of relatively simple mechanisms. Our approach is comparable to emergentist theories of other cognitive skills such as language, which is sometimes envisaged as a “new machine constructed entirely out of old parts” (Bates & MacWhinney 1989). A further consequence of our approach is that analogy may best be understood not as a uniform cognitive skill but, instead, as an umbrella term that describes different task-specific constellations of basic memory and control processes.

The rest of this target article unfolds as follows. First, we present a brief overview of the current state of research into analogical reasoning and its development. In the second section we consider key aspects of current

accounts of analogical reasoning and consider why these are difficult to reconcile with some aspects of cognitive development. We then present our suggestions for a more developmentally constrained theory based on priming within a semantic memory system. Subsequently, this verbal theory is implemented in a model that, although simple, illustrates how our account of analogical completion functions. We then demonstrate how this account is able to tie together a wide range of developmental findings into a single explanatory framework. Finally, we consider the theoretical implications of the model for the development of analogical reasoning.

1.1. Key features in the development of analogical reasoning

Because we are primarily concerned with development, we start by considering general accounts (and evidence) of how analogical reasoning develops before going on to consider more detailed and specific (largely adult) accounts of analogy. One approach (e.g., Halford et al. 1998; Hummel & Holyoak 1997; Piaget et al. 1977; Sternberg & Rifkin 1979) sets analogy alongside other cognitive skills and attempts to provide a domain-general explanation for all such skills. The second approach focuses more on whether analogical reasoning arises from increasing knowledge (e.g., Brown 1989; Gentner 1989; Gentner et al. 1995; Goswami 1992, 1996). While reviewing the experimental results from both approaches in sections 1.1.1–1.1.4, we enumerate the key phenomena that a developmental account of analogy must capture.

1.1.1. Analogy as a domain-general cognitive skill. Early research into the development of analogical reasoning (e.g., Piaget et al. 1977; Sternberg & Rifkin 1979) found scant direct evidence of analogy use by young children. This was generally taken as evidence of domain-general structural changes in children’s reasoning abilities. For example, Piaget tested 5- to 12-year-old children on picture-based “*a is to b as c is to what?*” analogies and found only occasional and uncertain evidence of analogical reasoning (Piaget et al. 1977). Piaget interpreted this as suggesting that analogical development should be understood in terms of his more general account of the development of logical reasoning. In a similar vein, Sternberg and colleagues (Sternberg & Nigro 1980; Sternberg & Rifkin 1979) used children’s reaction time data to argue that there was an age-modulated shift from solving analogies using largely associative strategies to using more genuine analogical reasoning strategies. However, both Piaget’s and Sternberg’s experimental results, and consequent theoretical positions, have been criticized on the grounds that they failed to take into account the children’s knowledge of the relations underlying the analogies. Consequently, they greatly underestimated children’s analogical reasoning abilities (Goswami 1991).

Some more recent theorists have also argued that domain-general changes have a particularly important role in young children’s emerging analogical reasoning abilities. These accounts focus on the development of capacity limits in active memory instead of structural changes in underlying reasoning mechanisms. Halford argues (see Andrews & Halford 2002; Andrews et al. 2003; Halford 1993; Halford et al. 1998) that one of the

ROBERT LEECH is a postdoctoral researcher in the School of Psychology at Birkbeck, University of London. He studied at Cambridge and Birmingham Universities before receiving his Ph.D. in Developmental Psychology at Birkbeck in 2004. His Ph.D. research focused on the development of analogical reasoning, primarily using computational models. In addition, his current research also concerns the development of linguistic and auditory processing using behavioral, neuro-imaging and computational methodologies.

DENIS MARESCHAL obtained his first degree in Physics and Theoretical Physics from Cambridge University, UK. He then completed a Masters in Psychology from McGill University, Canada, before moving on to complete a Ph.D. at Oxford University, UK. He has received the Marr prize from the Cognitive Science Society (USA), the Young Investigator Award from the International Society on Infant Studies (USA), and the Margaret Donaldson Prize from the British Psychological Society. His research centers on developing mechanistic models of perceptual and cognitive development in infancy and childhood. He is currently Professor at Birkbeck, University of London.

RICHARD P. COOPER obtained his first degree in Mathematics with Computer Science from the University of Newcastle, Australia. He then completed a Ph.D. in Cognitive Science, focusing on computational linguistics, at the Centre for Cognitive Science, University of Edinburgh, UK, before effectively retraining in cognitive psychology at the Department of Psychology, University College London. His research is primarily concerned with computational modelling of neurological disorders, though he has also published extensively on the methodology of cognitive modelling. He is currently Reader in Cognitive Science at Birkbeck, University of London.

most fundamental constraints acting on cognitive development is the maximum relational complexity that can be processed in parallel in working memory (see also Hummel & Holyoak [1997] for a similar account concerning the LISA model, discussed in section 1.2.1 of the target article). Halford and colleagues define complexity as “the number of related dimensions or sources of variation” (Halford et al. 1998, p. 803). Tasks involving one source of variation start to be processed around the first birthday. Binary relations (i.e., with two sources of variation), including rudimentary analogical reasoning, can be understood by about 2 years of age. By age 5 children are able to process ternary relations and so are able to demonstrate skills such as transitivity. For example, Richland et al. (2006) tested 3- to 14-year-old children on analogical mappings between pictures that varied in the number of relations that had to be integrated and the presence or absence of perceptual distractors. Although the youngest children performed well on the simpler analogies, relational complexity severely disrupted their mapping success, a result that diminished with age. The authors interpret these results as evidence that the general maturation of working memory and inhibitory mechanisms is at the heart of increased performance on analogical reasoning tasks. Finally, it is worth pointing out that these conclusions are not uncontested. Indeed, Goswami (1998) and Gentner and Rattermann (1998) provide some evidence that children younger than 5 years of age are able to form analogies involving ternary relations, while Halford et al. (1998) argue that these latter findings can be explained in terms of decomposing ternary relations into binary relations.

1.1.2. The development of analogical completion as knowledge accretion. In contrast to the early null findings, recent researchers have demonstrated analogical reasoning very early in development. Most strikingly, Chen et al. (1997) demonstrated 10- and 13-month-old infants’ ability to use analogy to solve a simple task. Here, the infant’s parent modeled a task where the infant had to combine two sub-goals to reach a toy – removing a barrier, pulling a cloth, and then pulling a string to reach the toy. The infants subsequently had to disregard superficial similarities to transfer the parent-modeled solution to a novel task involving the same underlying structure. This result suggests that at least the precursors of analogical reasoning are present from before the first birthday. Similar studies have demonstrated that 17- to 36-month-olds (see Brown 1989) and 2- to 4-year-olds (Crisafi & Brown 1986) benefit from analogical transfer in simple problem solving paradigms. Furthermore, children from 3 to 4 years of age can, given sufficiently familiar domains, solve analogies for more complex “*a is to b as c is to what?*” type analogies (Goswami & Brown 1989; 1990; Rattermann & Gentner 1998a).

Results such as these have been taken as evidence that the crucial constraint on analogical development is the knowledge that the child has, not some kind of general structural change (e.g., Goswami 1992). As children’s knowledge about the world becomes richer, they can better use this knowledge to form and understand analogies. It is worth noting that there is no inherent contradiction between domain-general changes in processing relational complexity and knowledge accretion. Indeed,

domain-general accounts also acknowledge a strong role for knowledge accretion as a driving force in analogical development. However, a substantial difference between the positions of Halford and colleagues and that of Goswami is that the latter places a far greater importance on the development of relational representations and downplays the importance of maturational change in working memory capacity.

Observation 1: There is a strong relationship between accretion of relational knowledge and successful analogical reasoning (Goswami & Brown 1989; Rattermann & Gentner 1998a).

To investigate this knowledge accretion account, Goswami and Brown (1989; 1990) developed a picture-based test similar to Piaget’s *a:b::c:d* analogies, for use with 3- to 9-year-olds. However, unlike Piaget’s experiments, Goswami and Brown used relations that were familiar to the young children, such as cutting (e.g., *Playdoh is to cut Playdoh as apple is to cut apple*). They found evidence of analogical completion in even the youngest children and found a correlation between analogical completion and an independent test of relational knowledge. They concluded that analogical reasoning is domain specific. Children start using analogies at different points in development as a result of the development of the appropriate knowledge representations. Further evidence for knowledge constraining analogical reasoning comes from studies investigating 4-year-olds’ analogical transfer solving biological problems (Brown & Kane 1988) and physical problems (Pauen & Wilkening 1997), and young children’s analogical inferences (Inagaki & Hatano 1987; Vosniadou 1989).

Observation 2: Given that children’s analogical reasoning depends on their underlying domain-specific knowledge, there is domain-specific as opposed to a domain-general change over development in children’s ability to form analogies (Goswami & Brown 1989).

An additional developmental phenomenon is the extent to which children use analogy spontaneously. Inagaki and Hatano (1987), Goswami and Brown (1989), and Pauen and Wilkening (1997) all reported some degree of analogical transfer in the absence of any explicit guidance. Spontaneous transfer was demonstrated more systematically in 4-year-olds by Tunteler and Resing (2002), whose results not only suggest that spontaneous analogical transfer occurs even in young children, but also that it becomes more likely with increasing experience of the problem domain – again consistent with the knowledge accretion account. Evidence for spontaneous analogical reasoning is particularly interesting in that it further suggests that analogy is an emergent phenomenon (with important developmental implications considered later).

Observation 3: Analogical ability occurs spontaneously within a domain (Goswami & Brown 1989; Inagaki & Hatano 1987; Pauen & Wilkening 1997; Tunteler & Resing 2002).

1.1.3. The relational shift and knowledge accretion. Gentner and colleagues (Gentner 1988; 1989; Gentner & Toupin 1986; Gentner et al. 1995) have suggested an

account of the development of analogy based on Structure-Mapping Theory (SMT), which is discussed in more detail in section 1.2.1. A key component of this account is that children undergo a relational shift whereby their analogical reasoning changes over time from being initially based on the similarity of object attributes to gradually including relational information between objects and subsequently incorporating systems of relations. Gentner et al. (1995) proposed that this change results from children progressively re-representing relations. Children move from using first order predicate-argument representations, for example, *darker* (a, b) (i.e., a is darker than b) where the dimension, *darkness*, and the comparison, *greater than*, are conflated in the same representation, to using increasingly abstract relational representations that support more complex mappings (e.g., *greater[darkness(a), darkness(b)]* with the comparison clearly separated from the dimension).

Observation 4: Over development, children show a “relational shift,” changing their preference for judging similarity from surface similarity to relational similarity (Gentner 1988; Gentner & Toupin 1986; Rattermann & Gentner 1998a).

Evidence for the relational shift comes from a number of areas, including a replication of the Goswami and Brown (1989) study taking into account object similarity (Rattermann & Gentner 1998a), studies of children’s metaphor comprehension (Gentner 1988), object similarity as a constraint on analogical problem solving (Holyoak et al. 1984), and children’s performance on cross-mapping tasks. We focus specifically on the last of these. In a cross-mapping task (see Fig. 1), children must make an appropriate analogical transfer in the presence of a conflict between object similarity and perceptual similarity. A number of experiments (see Gentner et al. 1995) suggest that younger children find cross-mapping problems harder than older children, although

the delay in cross-mapping performance may be driven by specific aspects of the task design and very rich, highly detailed stimuli. Goswami (1995) demonstrated that even very young children (3 years old) can solve the kind of simple cross-mapping tasks involving relative size represented in Figure 1. Although changes in cross-mapping ability could arise from development in domain-general abilities such as working memory capacity, they are also consistent with the hypothesis that children become better at processing relational information with age.

Observation 5: Children can solve analogies even when there is a conflict between object and relational similarity (cross-mapping), although these analogies may be harder (Gentner et al. 1995; Kotovsky & Gentner 1996; Rattermann & Gentner 1998b; although see Goswami 1995).

One interesting experimental finding is that the age at which children can solve cross-mapped analogies can be manipulated by teaching children relevant relational labels (Kotovsky & Gentner 1996; Rattermann & Gentner 1998b). With appropriate teaching of relational labels, a 3-year-old child can solve cross-mapped analogies normally only solved at 5 years or older. Similarly, Loewenstein and Gentner (2005, Experiment 3) found that preschool children who heard words for spatial relations performed better on a spatial mapping task than those who did not, and Goswami (1995) also demonstrated that 3- and 4-year-olds can use the familiar relational structure from “Goldilocks and the Three Bears” (i.e., Daddy Bear > Mummy Bear > Baby Bear) to solve transitive mapping problems. Such evidence is strongly consistent with the knowledge accretion account, although the importance of relational labels is also consistent with Halford’s (e.g., Halford et al. 1998) relational complexity account of analogical development.

Observation 6: Provision of labels affects the formation of representations and subsequent performance in analogical reasoning tasks (Kotovsky & Gentner 1996; Loewenstein & Gentner 2005; Rattermann et al. 1990).

Figure 1. An illustration of cross-mapping with the relation larger-than. The circle in the center of the top row can be mapped onto two of the circles in the lower row as indicated by the arrows. The dashed arrow shows a literal similarity mapping where identical objects are mapped onto one another. The full arrow shows the relational cross-mapping where object identity is ignored and the mapping is based on the relative size of the objects.

1.1.4. Indicators of discontinuous change. A very different approach to the development of analogical reasoning comes from dynamical systems theory (van der Mass & Molenaar 1992). Hosenfeld et al. (1997) investigated whether the development of analogical reasoning could be understood as a phase transition (specifically a bifurcation) in a dynamical system. Their longitudinal study with eighty 6- to 8-year-olds looked for indicators of such a discontinuous change in children’s performance on twenty geometric analogies. The authors found evidence for three such indicators, all occurring at approximately the same point in development. First, a rapid improvement in children’s test scores was observed, consistent with a sudden jump in children’s performance. Second, Hosenfeld et al. noted an increase in the inconsistency of children’s responses, which they interpreted as an increase in anomalous variation typical of transitional dynamical systems. Third, they reported evidence of a critical slowing down in solution times around the time of the sudden jump.

Observation 7: During the course of analogical development, children display several indicators of a discontinuous change (Hosenfeld et al. 1997): (i) sudden jump in correct responses, (ii) anomalous variation, and (iii) critical slowing down.

We have summarized the developmental literature into seven key phenomena that any successful theory of the development of early analogical reasoning must take into account. The primary focus of the following work is on simple analogies of the kind first solved by very young children (i.e., involving comparisons between pairs of relations) demonstrating these developmental markers. As such we, initially, do not consider some aspects of more complex analogical reasoning, (like the importance of relational complexity; Andrews & Halford 2002; Andrews et al. 2003). However, in sections 4 and 5 of the target article we consider more complex analogies and analogical mapping, and we will briefly return to the issue of the interrelationship between relational complexity and analogical reasoning at that point.

1.2. Models of analogical reasoning

Having reviewed behavioral data relevant to the development of analogical reasoning, we now turn to existing accounts and computer models of analogical reasoning in general and the extent to which such accounts can capture the developmental phenomena listed in the previous section. The following review is necessarily brief, but more extensive surveys of computational models of analogy can be found in Holyoak and Barnden (1994), Barnden and Holyoak (1994), and French (2002). We start with Structure-Mapping Theory (SMT) (Gentner 1983; 1989), which has probably been the most influential theoretical position in analogical reasoning for the past two decades (for other models strongly influenced by SMT, see Forbus et al. 1994; Keane & Brayshaw 1988; Keane et al. 1994; Larkey & Love 2003). SMT is important because it illustrates several key aspects of many current theories of analogy. It has also led to a substantial body of empirical research (e.g., Bowdle & Gentner 1997; Clement & Gentner 1991; Markman & Gentner 1997).

1.2.1. Structure-Mapping Theory and related models.

Structure-Mapping Theory is an account of how people use analogies to draw inferences (e.g., inferring that an electron rotates around a nucleus based on the more familiar knowledge that the planets rotate around the sun). For SMT, the actual attributes of objects such as color, size, and so on, are normally irrelevant; what is important are the relations between objects (such as Object A revolves around Object B). In SMT, mental representations are highly structured, being composed of predicates with arguments. Given this assumption, SMT distinguishes between object attributes, and relations between objects, at a purely syntactic level, with no regard for semantic content.

Analogical reasoning, under this account, involves first selecting a base domain from memory based on surface similarity, followed by creating a structural mapping between base and target. This process involves trying to put objects in the base in a one-to-one correspondence with objects in the target. Predicates between target objects are then matched with identical predicates in the

base domain. Which relations are mapped from base to target are governed by a preference for systematicity among the relations, that is, a preference for higher order relations between relations. This preference determines what inferences will result from an analogy.

Structure-Mapping Theory has been implemented in the two-stage MAC/FAC (“many are called but few are chosen”) computational model (Forbus et al. 1994). This model is a compromise between the desire to have structural analogies and the computational requirements of searching through long-term memory. The MAC component economically selects a few candidates from a vast number of options in long-term memory using a nonstructural matcher based on the number of occurrences of a predicate in a candidate. The FAC stage then uses the much more computationally expensive Structure-Mapping Engine (SME) (Falkenhainer et al. 1989) to choose the best structural match from the candidates. This two-stage process was designed to be consistent with the body of psychological research which suggests that nonstructural similarity constrains the retrieval of a base domain (e.g., Gick & Holyoak 1983; Novick 1988; but see Blanchette & Dunbar [2001; 2002] for a discussion of the role of other factors such as audience characteristics and goals on analogical reasoning, as well as the important distinction between analogical retrieval and generation).

Importantly, SMT has been used to simulate the development of analogical reasoning (Gentner et al. 1995). The same process model was tested on an analogical task under different representational assumptions, mirroring Gentner et al.’s assumptions about the knowledge base of 3- or 5-year-old children. By changing how much higher-order relational information was represented, Gentner et al. simulated key differences in performance seen in experiments with 3- and 5-year-olds. SMT demonstrated successful cross-mapping for the older children – and demonstrated a relational shift. Furthermore, by varying the abstractness of representations, Gentner et al. were able to use SME to simulate the age-related change observed in children from making only within-dimension mappings (e.g., within the dimension *size*) to cross-dimension mapping (e.g., mapping greater *size* onto greater *brightness*).

A consideration of SMT is important because two central features of the model – explicit, structured representations and some form of structure mapping – appear in some form or another in many other models. This can be seen even in accounts that use seemingly fundamentally different architectures. For example, the Analogical Constraint Mapping Engine (ACME; Holyoak & Thagard 1989) and Learning and Inference with Schemas and Analogies (LISA; Hummel & Holyoak 1997) are both hybrid connectionist/symbolic systems. Yet, they both focus largely on the process of mapping between base and target domains, incorporating multiple constraint satisfaction of semantic information as well as structural information. ACME (as SMT) starts with predicate-argument representations. Three constraints then govern the mapping process: structural, semantic, and pragmatic. These constraints are intended to act together as pressures to make local decisions about correspondences between elements in order to produce a psychologically plausible global mapping. LISA is an attempt to build upon ACME to incorporate more aspects of traditional connectionist models. Key to the working of

LISA is that relevant objects and roles are bound via temporal synchrony (Shastri & Ajanagadde 1993). For example, the proposition loves(Jim, Mary) is represented in LISA's semantic units by the patterns for Jim and Lover being active simultaneously while the patterns for Mary and Beloved are independently synchronously active. Hummel and Holyoak (1997) hypothesize that maturation in the number of concepts that can be simultaneously processed could explain the development of analogy. Halford et al. (1994; 1998) make similar developmental claims from a different Structured Tensor Analogical Reasoning (STAR) symbolic connectionist approach (see also Eliasmith & Thagard 2001).

The Copycat models (Hofstadter 1984; Mitchell 1993) are another example of hybrid accounts of analogical reasoning. These models involve a long-term memory, interconnected nodes forming a semantic network, and a working memory where structures representing the analogical problem are built. These models also consist of "codelets" or "agents" that cooperate and compete to construct descriptions of relationships between objects. Copycat differs from other models by actively constructing new representations. The representations are influenced by both top-down and bottom-up processes, motivating Hofstadter to describe Copycat as "high level perception" (for similar accounts, see also Barnden 1994; French 1995; Kokinov & Petrov 2001).

1.2.2. Distinguishing features of models. A useful characterization of models of analogical reasoning might split them according to how information is represented. Thus, models such as SME, ACME, LISA and STAR, which all use structured predicate-argument type representations as input to the model, may be contrasted with models such as Copycat, that start with less structured representations. In the former type, representation formation is an integral component of the modeling process, whereas in the latter, analogy is viewed as more closely related to perceptual processes.

2. Analogy as relational priming

Although explicit mapping appears to be a near ubiquitous feature of the theories and models surveyed above (e.g., Gentner 1983; Holyoak & Hummel 2001; Kokinov & French 2003; Novick 1988), there is evidence to suggest that it may not be necessary for successful analogical reasoning. For instance, Ripoll et al. (2003) demonstrated a disassociation between analogical mapping and analogical transfer. Explicit mapping accounts (e.g., SMT) predict that analogical transfer occurs only after a mapping is established between the base and the target. Consequently, if analogical transfer requires prior mapping, then the reaction time of analogical transfer should increase with reaction time for mapping. Ripoll et al. (2003) found that while reaction time increased in a cross-mapping condition when participants were asked to perform a mapping task, reaction time did not increase when participants were tested for analogical transfer. Therefore, there is reason to consider alternatives to explicit mapping as the fundamental mechanism behind the development of analogy-making abilities.

In the following sections we describe the two mechanisms that form the backbone of our developmental account of analogy-making. First we suggest that priming is centrally implicated in analogical reasoning. Second, we propose that relations are best represented as transformations between states, rather than as explicit symbols.

2.1. Analogical reasoning as priming

Many cases of analogy involve seeing the similarity between the prominent relation in one domain and the prominent relation in a second domain. Within traditional approaches to analogical reasoning, the cognitive mechanism for an *a is to b as c is to ...?* analogy involves mapping the *a* term onto the *c* term and the *b* term onto the unknown *d* term and then transferring across the relation between the *a* and *b* terms to the *c* and *d* terms. We suggest that this is accomplished via a simpler mechanism based on relational priming. Put most simply, we propose that exposure to the *a* (e.g., *puppy*) and *b* (e.g., *dog*) terms of an analogy primes a semantic relation (e.g., *offspring*), which then biases the *c* term (*kitten*) to produce the appropriate *d* term (*cat*).

What evidence is there for this proposal? First, medium- to long-term semantic priming is a ubiquitous phenomenon that is well established in the adult literature (e.g., Becker et al. 1997; see also Chapman et al. [1994] for a review of priming in children). Second, several studies now suggest a general role for priming in analogical reasoning. For example, Schunn and Dunbar (1996) presented participants with a biochemistry problem on one day and with a genetics problem on the following day. Although the two problems were different, their solutions both involved inhibition. Participants were significantly more likely than controls to propose an inhibition solution to the genetics problem following the biochemistry problem. However, the participants did not mention the prior biochemistry problem either during the experiment or in a post-task questionnaire. Consequently, the authors explained the results as implying a form of priming. Kokinov (1990) also demonstrated priming in analogical problem solving. Prior to being given a difficult target problem (e.g., heating water in a forest), participants were primed with a different analogical problem whose commonsense solution was well known to them (e.g., heating a cup of tea in a mug). The performance of the participants rose substantially immediately after priming, returning to control levels after 24 hours.

Third, recent studies have repeatedly demonstrated the presence of relational priming under a variety of different conditions (Estes 2003; Estes & Jones 2006; Gagne 2001; 2002; Gagne & Shoben 1997; Gagne et al. 2005; Gerrig & Murphy 1992; McKoon & Ratcliff 1995; Spellman et al. 2001; Wisniewsky & Love 1998), including relational priming resulting directly from analogical reasoning (Green et al. 2007). In these studies, prior exposure to a relation (e.g., by presentation of two nouns joined by a relation, such as *apple* and *cake*) facilitates subsequent judgments involving that relation (e.g., *made of*).

In summary, semantic priming effects are commonly reported across many areas of cognitive psychology (see Tulving & Schacter 1990). More importantly, relational priming is a robust psychological phenomenon that does not require explicit strategic control. Consequently,

relational priming is a choice candidate mechanism for a developmental account of analogy, emerging from simple memory processes. The ubiquity of priming effects in development suggests that it is a plausible building block for a theory that posits the emergence of analogical completion from simple cognitive mechanisms (for a similar view of the relation between inhibition – negative priming – and cognitive development, see Houdé 2000). Indeed, a strength of a relational priming account of analogical reasoning is that it does not posit analogy-specific mechanisms.

2.2. Representation of relations

Some models of analogical reasoning (e.g., Copycat) attempt to integrate analogical reasoning with lower-level perceptual representations. Other models (SME, LISA) use explicitly structured predicate and argument type representations. The issue of how object attributes and relations are represented is important for modeling analogy because it constrains what else can take place; for example, mapping is one obvious way to compare systems of explicit structured representations (e.g., predicates). However, there are serious concerns about how these predicate representations can be acquired (see Elman et al. 1996; Shultz 2003). Hence, predicate and argument representations may not be appropriate for a developmental account of analogical reasoning.

In fact, within the context of analogy, relations need not be represented in predicate/argument terms at all. For the purposes of analogy it may be sufficient to conceptualize relations as transformations between items. This has the important advantage that relations may be learned via well-understood (connectionist) procedures. Moreover, there are notable precursors in the literature for viewing relations as transformations. For instance, Thomas and Mareschal (1997) and Hahn et al. (2003) argue that transformations underlie similarity judgments and, by extension, analogical reasoning. One such account, the metaphor as pattern completion model (MPC) (Thomas & Mareschal 2001; Thomas et al. 2001), is particularly relevant because of its strong focus on development and because it simulates the emergence of metaphor, closely related to analogy. Similarly, Rogers and McClelland (2004) present an account of the development of semantic cognition that proposes that relations are transformations. Finally, viewing relations as transformations in a semantic space suggests that “relational similarity” might be a performance factor in analogical completion; this is indeed the case, at least in adult participants (Leech et al. 2007).

3. A model of the Goswami and Brown paradigm

The seven key developmental phenomena we detailed in earlier sections can all be seen in variants of the “*a is to b as c is to what?*” analogies, such as the Goswami and Brown (1989; 1990) paradigm. According to Sternberg (1977a), “*a is to b as c is to what?*” analogies incorporate the core information processing components required for analogical reasoning. Therefore, to facilitate comparison with the developmental evidence, we have embodied our two central theoretical tenets (relational priming and

relations as transformations) in a connectionist model of the Goswami and Brown paradigm. We will first describe the task and model and then present a number of specific simulations exploring this paradigm, before finally presenting a simulation which takes a step back and tentatively suggests how the underlying principles of the specific model could form part of a much more general explanatory relational priming framework.

3.1. The Goswami and Brown paradigm

The Goswami and Brown forced choice task involves children selecting a picture to complete an analogical sequence involving simple causal relations (see Fig. 2). After seeing three pictures (e.g., *bread*, *cut bread*, and *lemon*) the child is given four response options: (a) the analogically appropriate response (e.g., *cut lemon*); (b) the correct transformation applied to the wrong object (e.g., *cut cake*); (c) the wrong transformation applied to the correct object (e.g., *squeezed lemon*); and (d) an object-similarity match (e.g., *yellow balloon*).

Consistent with the Goswami and Brown paradigm, the model focuses on simple causal domains (e.g., *cutting*, *melting*, *turning on*, *burning*) such as those used to test young children (Goswami & Brown 1989; 1990; Rattermann & Gentner 1998a). In these tasks, common objects (e.g., *apples*, *bread*) are transformed by a causal agent (e.g., *knife*), as when an apple is cut by a knife. The event sequence experienced by the network in the example would be: first, to be presented with representations of an apple and a knife, and then, with representations of a cut apple and a knife. The task of the network is to learn the transformation from apple to cut apple in the context of knife, which, consistent with our theoretical assumption about transformations and relations, is equivalent to learning the relation *cutting*. Once the network has learned such relations, analogical completion may be modeled by first exposing the network to the *a* and *b* terms of the analogy (e.g., *apple* and *cut apple*), thus priming a relation (in this case *cutting*); and then presenting the network with the *c* term of the analogy (e.g., *bread*). The network should then settle into a state consistent with the product of the *c* term and the primed relation (in this case, *cut bread*).

3.2. The model

3.2.1. Network architecture. Figure 3 shows the architecture of the connectionist network used to model both the acquisition of relational information and the completion of analogies within the Goswami and Brown paradigm. All network weights are bidirectional and symmetrical,

Figure 2. An illustration of the Goswami and Brown paradigm (adapted from Rattermann & Gentner 1998a).

Figure 3. Top: schema of the model architecture. Below: an example of the different activation phases during training. Darkly shaded layers are clamped onto a given target activation pattern (e.g., *apple*, etc.); unshaded layers are unclamped and free to change.

thereby enabling the flow of activation in all directions throughout the network. The bottom layer (roughly corresponding to the input layer) is split into two banks of units, representing the presentation of two different objects in a “before,” or pre-transformation, state. The two different banks of units correspond to an object [i.e., *Object(t1)*] that can be transformed (e.g., *apple*) and a causal agent object [i.e., *CA(t1)*] that causes the transformation (e.g., *knife*). Similarly, the upper layer (corresponding to the output layer) is split into two banks, representing the same two objects [i.e., *Object(t2)* and *CA(t2)*] in their “after,” or post-transformation, state. In the current simulations, we assume that objects are encoded in terms of perceptual features only (e.g., shape, size, color) at both input and output.

One can think of the “before” and “after” representations as two temporally contiguous states of the world. Because both the “before” and “after” states can be obtained by direct observation of the world, learning of relational information does not require an external teacher, and constitutes a form of self-supervised learning (Japkowicz 2001). Both the “before” and “after” states of the object representations [i.e., *Object(t1)* and *Object(t2)*] and the hidden layer have 40 units each. The “before” and “after” states of the causal agent [*CA(t1)* and *CA(t2)*] have 4 units each. The activation of any unit varies according to a sigmoidal activation function from 0 to 1 (this is sometimes referred to as an “asigmoid activation” unit; Shultz 2003). The initial weights are randomized uniformly between ± 0.5 .

Because of the bidirectional connections, input activation can cycle throughout the network before settling into a stable attractor state. During training, contrastive Hebbian learning is used to change the connection weights such that attractors on the output units coincide

with target output states of the network (for an explanation of contrastive Hebbian learning and a presentation of the benefits – in terms of biological plausibility – of this algorithm with regard to other supervised training algorithms, see O’Reilly 1996; O’Reilly & Munakata 2000). As with the better-known backpropagation training algorithm, contrastive Hebbian learning creates internal representations across the hidden units to solve complex problems. As the name suggests, within the learning algorithm weight changes are calculated locally as the difference between a Hebbian and an anti-Hebbian term. These terms correspond to different states of activation of a unit. Contrastive Hebbian learning requires two phases of activation during training. The first phase, the minus phase, involves clamping some of the units [e.g., the *Object(t1)* units] to a desired pattern and letting the remaining units’ activation spread through the network (we used five activation cycles between each weight update). For example, in Figure 3, the *Object(t1)* and *CA(t1)* units are clamped, and the hidden units and *Object(t2)* and *CA(t2)* are free to change and settle into a stable state. The resulting activation state of *Object(t2)* is taken as the response the network arrives at for a given input. In the second phase, the plus phase, all the external units (inputs and outputs) are clamped on. Only the hidden units’ activation settles into a stable state, constrained by all the external units [i.e., the *Object(t1)* and *Object(t2)* and *CA(t1)* and *CA(t2)* units]. The state of activation of the plus phase corresponds to the desired activation of the network given a certain input.

Contrastive Hebbian learning uses the difference between the plus and minus phases to update the connection weights as follows:

$$\Delta w = \alpha [x^+ y^+ - x^- y^-]$$

where α is a learning-rate parameter (set to 0.1 in all simulations reported here), x and y are the activations of two interconnected units, and the superscripts distinguish between the values of the plus and minus activation phases. As learning proceeds, the difference between the weights in the plus and minus phases reduces as the activation in the minus phase comes to replicate that in the plus phase.

3.2.2. Training: The learning of causal relations. In the current model, networks were trained on input patterns produced on-the-fly by adding Gaussian noise ($\mu = 0.0$, $\sigma^2 = 0.1$) to prototypes selected at random from a predefined pool of 20 different possible *Object(t1)* prototype patterns, and 4 different *CA(t1)* prototype patterns. The prototypes consisted of randomly generated input vectors with slot values within the range [0, 1], and where each vector slot value was set to 0 with a probability of $p = 0.5$. Slot values were set to 0 in the prototypes to increase the sparsity of external representations, whereas the addition of noise to the prototypes was intended to capture the fact that although two instances of, for example, cutting an apple with a knife are similar, they are not identical.

Four transformation vectors were also randomly generated but were set to have a Euclidean distance from every other transformation of less than 10. The transformation vectors encode the relation between the pre- and post-transformation states of the object. In fact, the

transformed state of the object, Object(t2), is obtained by adding a transformation vector to Object(t1). For example, Object(t1) (e.g., *apple*): [0.5 0.0 0.2 0.0 0.8 0.2 0.0 0.4] might be transformed by the vector (e.g., *cut*): [-0.4 0.0 0.0 0.0 0.0 0.7 0.0 0.0], resulting in Object(t2) (*cut apple*): [0.1 0.0 0.2 0.0 0.8 0.9 0.0 0.4]. Note that although the transformation vector is used to generate the target pattern corresponding to any particular input, the transformation vector itself is never presented to the network. Different objects (e.g., *bread* or *apple*) transformed by the same relation (e.g., *cut*) are transformed by the same vector. Thus, the network can learn about a particular transformation by generalizing across sets of Object(t1)/Object(t2) pairings that are affected by that transformation.

In the model, CA(t1) represents a causal agent (e.g., *knife*) which when presented concurrently with certain (but not all) objects at Object(t1) (e.g., *apple*), leads to a transformed Object(t2) representation (e.g., *cut apple*). Consequently, the target pattern for the Object(t2) depends on CA patterns. In the simulations presented here CA(t1) always remains the same at CA(t2) (i.e., CA is never transformed). Training consists in randomly selecting an object and a causal agent, computing the transformed state, Object(t2), and updating the weights such that the actual Object(t2) state produced by the network approaches the target Object(t2). The partitioning of banks of units into object and causal agent layers is actually a property of the training regime, not of the network architecture. More complex training environments could also lead to a change in the state of CA at t2 (e.g., *knife* at t1 to *wet knife* at t2).

Each of the 20 Object(t1) representations can be affected by 2 of the 4 causal agents (and thus 2 of the 4 transformations). When an object is presented in conjunction with one of the remaining 2 causal agents, the target Object(t2) pattern is the same as the untransformed Object(t1) pattern. Thus, whereas the causal agent *knife* transforms *apple* to *cut apple*, the causal agent *water* (for example) has no effect on *apple*. Equally, whereas the causal agent *knife* transforms *apple* to *cut apple*, the causal agent *knife* has no effect on *rock*. Hence, the presence of the causal agent alone is not a predictor of whether a transformation will occur. Given this organization, there are 360 potential analogies (20 objects $\times$ 2 causal agents $\times$ 9 other objects that can be affected by the same causal agent) on which the network may be tested.

3.2.3. Testing: Analogical completion. The testing of analogy completion proceeds in a different way from the learning of relation information. As we have stressed, priming is fundamental to our account of analogical completion. It occurs in the network because the bidirectional connections allow the hidden and “after” layers to maintain activity resulting from an initial event. The activity that is maintained in the network biases how new external input is then subsequently processed.

To illustrate how activation-based priming and pattern completion combine to complete analogies, we consider the archetypal case of *a:b::c:d* analogies. First, the units are clamped with the representation of *apple* at Object(t1) and *cut apple* at Object(t2), while CA(t1) and CA(t2) are initially set to 0.5, the resting value. This corresponds to

being presented with the information *apple:cut apple* (i.e., the first half of an *a:b::c:d* analogy). The causal agent is not presented to the network at any point during testing. The network settles into the attractor by filling in CA(t1) and CA(t2) and arriving at hidden unit activations consistent with the transformation cutting. Following this, the Object(t1) and Object(t2) units are unclamped, and a second pattern, corresponding to *bread*, is presented to Object(t1) and nothing presented to Object(t2). CA(t1) and CA(t2) are initially presented with resting activation patterns and then unclamped. This corresponds to being presented with the information *bread:?* (i.e., the second probe-half of the *a:b::c:d* analogy). By unclamping the original object and causal agent units and by presenting a different Object(t1) pattern, the network is no longer in equilibrium and settles into a new attractor state. During training, the network has encoded, in the connections to and from the hidden layer, the similarities in the transformations corresponding to relations such as *cutting*. Consequently, the prior priming of the *apple* and *cut apple* transformation biases the network to settle into the attractor state consistent with the transformation *cutting*, which gives the *cut bread* pattern at Object(t2). The network has now produced the appropriate response at Object(t2) to complete the analogy (i.e., *apple:cut apple::bread:cut bread*).

3.2.4. An example of developing analogical ability. In the Goswami and Brown paradigm, children are presented with the *a:b::c* terms of the analogy and four response options. Figure 4 shows, at three different stages of learning, the sum of squared distance (SSD) between the actual output of the network when tested on the *bread:cut bread:apple::?* analogy and four possible trained Object(t2) patterns as activation propagates throughout the network over five cycles. The lower the *y*-axis value, the closer the actual

Figure 4. The *y*-axis shows the sum-of-squared distance calculated between the actual output of the network and four target patterns (the four different lines). The lower the *y*-axis value, the closer the actual activation is to that target pattern. Figures (a–c) show the network’s response to an analogy (e.g., *bread is to cut bread as apple is to what?*), and (d) shows an example of the network’s response to a non-analogy (e.g., *bread is to bread as apple is to what?*).

activation is to that possible output pattern. Consequently, Figure 4 shows which of the four objects that the network has been trained on is closest to the network's actual response after different amounts of training. The four Object(t2) target patterns (taken from the Goswami and Brown paradigm) that the network's response is compared with are: (1) the analogically appropriate transformed object (i.e., *cut apple*); (2) a possible Object(t2) which is perceptually identical to the Object(t1) representation (e.g., *apple*); (3) the Object(t1) changed by an inappropriate transformation (e.g., *bruised apple*); and (4) a different Object(t1) pattern transformed by the correct transformation (e.g., *cut banana*).

After 100 epochs of training (Fig. 4a), the network is unable to complete the analogy appropriately. Instead its output is closest to *apple* (i.e., the object similarity response). After 2,000 epochs of training (Fig. 4b), the output is ambiguous, equally close to both *apple* and *cut apple*. After 5,000 epochs of training (Fig. 4c), the network settles into the appropriate state (i.e., *cut apple*).

3.2.5. An example of non-analogy. To infer correct analogical completion, it is not enough to demonstrate that it is occurring in the appropriate context. It is also important that the network does not produce an analogical response when it is *not* appropriate. Consistent with the results in Figures 4a–c, it could be the case that the network has developed the attractor corresponding to *cut apple* with a basin so wide that the activation settles into it whenever the network is presented with *apple*. However, consideration of the performance of the network after 5,000 epochs of training (Fig. 4d) demonstrates that this is not the case. Here, the network was presented with *bread* at Object(t1) and the untransformed *bread* pattern at Object(t2). Subsequently, Object(t1) was clamped to *apple* and the network allowed to settle. The resulting activation state was consistent with the Object(t2) for the untransformed *apple* pattern. Thus, when primed with a non-transformation example, the network appropriately produces the non-analogical response.

3.3. Simulating the developmental markers of analogical completion

Having established that the model captures the broad-brushed developmental profile of children's analogical completion, we now consider how the model captures the seven key more detailed developmental phenomena highlighted above.¹

3.3.1. The relationship between knowledge accretion and successful analogical reasoning. The thick line in Figure 5 shows the network's performance when tested on all 360 possible analogies across training. An analogy is assumed to have been successfully completed if the sum of squared difference between the actual activation and the analogically appropriate target is lower than the sum of squared difference between the actual activation and each other possible response. After 100 epochs of training, less than 20% of analogies are completed successfully. However, by 5,000 epochs of training, the network produces the analogically appropriate response for almost 100% of possible analogies. The thin line in the same figure shows the mean sum of squared error at Object(t2).

Figure 5. Percentage of analogically appropriate responses at Object(t2) (thick line) and the normalized mean sum of squared error at output (thin line) over training.

This is a measure of how well the network has mastered the causal domain on which it is trained. The proportion of analogies correct and sum of squared error are strongly negatively correlated (Spearman's $\rho = 0.99$; $p < 0.001$).² Thus, consistent with developmental evidence, the network's performance on analogical completion is highly correlated with its domain knowledge of causal transformations.

The relation between domain knowledge and analogical completion may also be shown by extracting the causal agent responsible for a given transformation. Goswami and Brown (1989) tested relational knowledge by asking children to choose the causal agent responsible for transforming different objects. This may be tested in the network by prompting it with *apple* and *cut apple* at Object(t1) and Object(t2) and resting patterns at CA(t1) and CA(t2). The network should then produce the appropriate causal agent (i.e., *knife*) at CA(t1) and CA(t2). This ability is important because it demonstrates that the analogical completion observed in the network is not simply a matter of forming a simple input-output (or stimulus-response) link.

Figure 6 presents the performance of the network at extracting the causal agent over training. The proportion of correctly produced causal agents closely correlates with performance on analogical completion (Spearman's $\rho = 0.508$; $p < 0.001$). This strong correlation mirrors the results obtained by Goswami and Brown (1989) with young children.

3.3.2. Domain-specific change in children's ability to reason analogically. Domain specificity is a natural corollary of the strong and drawn out relationship between relational knowledge and analogical ability – as the network gradually becomes able to master relational knowledge in different domains it will acquire the ability to solve analogies in that domain. Figure 7 illustrates this phenomenon by showing the network's performance over training when a single object is tested on two analogies involving distinct causal transformations. For this example, the network solves one analogy over 2,500 epochs earlier than the other. Again, this is consistent with the developmental

Figure 6. The thick line shows the percentage of correctly produced causal agents over training when the network is prompted with a transformation. For comparison purposes, the percentage of analogies completed appropriately (the thin line) is also included (averaged across 50 replications).

literature: the profile parallels the analogical performance observed with children, suggesting that the ability to solve analogies in different domains arises at different points in development.

3.3.3. The spontaneous production of analogical completion. As noted above several developmental studies have shown that children use analogies without explicit teaching (e.g., Goswami and Brown 1989; Ingaki & Hatano 1987; Pauen & Wilkening 1997; Tunteler & Resing 2002). The network mirrors this performance. At no point is the network trained on an example of an analogy; instead the network is only trained on transformations. Nor does the

Figure 7. Domain specificity of analogical completion. The thick and thin lines show performance on a single example object involving different transformations (e.g., analogy 1 = *apple:cut apple::bread:cut bread*; analogy 2 = *apple:bruised apple::banana:bruised banana*). (The figure shows raw unfitted data averaged over 10 replications.)

network have any dedicated architecture for performing analogy. Analogical completion is an emergent phenomenon resulting from the way relational information is represented and the way analogies are tested.

3.3.4. A shift in analogical judgment from surface similarity to relational similarity. Children demonstrate a relational shift over development. They appear to move from judging similarity in terms of object features to judging similarity on the basis of relational similarity in analogy tasks (Rattermann & Gentner 1998a). To investigate whether the network undergoes a relational shift over training we compared the types of errors produced by the network with those made by children. Two of these types of errors are: *object-similarity errors* (where the network responds at Object(t2) with the same pattern as at Object(t1)) and *wrong transformation errors* (where the network produces the appropriate object at Object(t2) but transformed by a non-primed causal agent).

Object-similarity errors and wrong transformation errors are the kinds of errors predominantly made by 4- to 5-year-old children (Rattermann & Gentner 1998a). Table 1 presents a comparison of children's analogical completion over development and the network's analogical completion over training. The network provides a reasonable approximation of the children's response profiles. Importantly, it shows the same shift over training with a considerable decrease in the proportion of appearance responses (from 22.4% to 2.5%). This is matched with an increase in correct responses (i.e., correct transformation responses) from 39.3% to 63.5%. Such behavior is consistent with the relational shift phenomenon in which children produce more transformation-based analogies as they get older.

Bootstrap re-sampling tests (Efron & Tibshirani 1998) were used to compare the mean performance for children in each cell to the distribution of means found by repeatedly sampling subsets of the individual networks' responses at 2,300 and 2,800 epochs. Children and the model do not differ on either correct response ($p > 0.1$) or object similarity ($p > 0.1$). The networks produce significantly fewer wrong transformation errors than the children ($p < 0.01$), however, the difference in means between 4- and 5-year-old children on wrong transformation responses does not differ significantly from the differences

Table 1. *The profile of responses made by children and the network, averaged over 50 replications (in percentages)*

Response Type	Children		Networks	
	4-year-olds	5-year-olds	2,300 epochs	2,800 epochs
Correct analogy	35	67	39.3	62.5
Incorrect	35	28	16.1	8.9
Object-similarity	22	3	22.4	2.5

Note. The children's data are taken from Rattermann and Gentner (1998a).

in network performance at 2,300 and 2,800 epochs ($p > 0.1$).

3.3.5. The effect of relational labels on analogical performance. Gentner and her colleagues (Kotovsky & Gentner 1996; Loewenstein & Gentner 2005; Rattermann et al. 1990) have repeatedly found that prior training or exposure with appropriate relational labels facilitates children's analogy-making. A variant of the network shows a parallel effect on analogical completion when trained with an analogue of relational labels.

To model the effects of labeling on analogical completion the network architecture was modified to include two additional layers of four relational units: RL(t1) and RL(t2) (see Fig. 8). These were bidirectionally connected to the hidden layer (as are the CA(t1) and CA(t2) layers). During training, each of the relational label units uniquely coded for a transformation, thereby labeling that relation. For example, when the network was presented with *apple* and *knife* it was also presented with the relational label *cutting* (e.g., $RL(t1) = [1\ 0\ 0\ 0]$, whereas the relational label *bruising* might be $RL(t1) = [0\ 1\ 0\ 0]$). Therefore, the RL(t1) and RL(t2) served to uniquely identify a transformation irrespective of the Object(t1) and Object(t2), and as opposed to the CA(t1) and CA(t2), which as in the original simulations were ambiguous. Although the addition of the RL units simplifies the learning task confronting the network, it provides no additional information that is not already conjointly present in the object and causal agent layers.

There was a marked difference in analogical completion over training with the addition of relational labels (Fig. 9). On average, relational labels resulted in the network completing analogies earlier. To assess whether this difference is statistically meaningful, we found the maximum value of the first derivative of the fitted curves for each individual network – this value corresponds to the sudden jump in performance for each network (See Sect. 3.3.6). The median number of epochs before the maximum slope was 1,830 when labels were supplied and 2,665 when labels were not supplied. A Wilcoxon signed-rank test indicated that this difference was highly significant, $p < 0.0001$. Despite the substantial difference in performance across relational label conditions, both early (prior to 1,500 epochs) and later (after 4,000 epochs) analogical completion were similar to completion without relational labels. Thus, the presence of relational labels does not change the overall developmental picture. What it does is move forward developmentally the point at which

Figure 8. Schema of the model architecture incorporating relational labels.

Figure 9. The percentage of analogies solved appropriately both with relational labels (the thin line) and without (the thick line). Data are averages of 20 networks in each condition.

there is a sudden shift in the ability to complete analogies. Again, this behavior mimics the performance observed experimentally with children. The role that the label plays here is to provide clearer, more consistent, task relevant constraints. The network's performance is accelerated because it can make use of this more consistent cue.

3.3.6. Indicators of a discontinuous change. The network also exhibits all three indicators of discontinuous change in children's analogical reasoning abilities identified by Hosenfeld et al. (1997). First, Hosenfeld et al. reported a sudden jump or rapid acceleration in the proportion of children's correct responses between test sessions. The network shows a similar rapid acceleration in performance at an average of 2,230 epochs of training (Fig. 10a). This figure also shows the first derivative (rate of change) of analogy performance. The first derivative reveals not only the sharp discontinuity of the sudden jump, but also the complex developmental trajectory. In particular we see reductions in the rate of change of analogical performance surrounding the sudden jump, and also that the model exhibits secondary, less extreme, sudden jumps in analogical performance.

Hosenfeld et al. (1997) also demonstrated an increase in the inconsistency of children's responses occurring alongside the sudden jump in correct performance. One way to assess the inconsistency of the network's response is to present the same network twice with noisy versions of the same analogy. The percentage of analogies that are not completed in the same way is a measure of the network's inconsistency at a given training epoch. Figure 10b shows the percentage of inconsistent responses calculated on presentation of each analogy twice to 50 networks. Figure 10b reveals that the inconsistency of the network's performance undergoes a rapid increase and peaks at around 2,500 epochs, shortly after the onset of the sudden jump. A nonparametric correlation comparing the epoch of the maximum value for the first derivative of accuracy (the sudden jump) with the maximum first derivative for the inconsistency scores across individual networks

Figure 10. (a) Evidence for a “sudden jump.” The thick line is the percentage of analogies completed correctly. The thin line is first derivative (rate of change) of the percentage correct across time, demonstrating a “sudden jump” centered around 2,200 epochs of training (over 50 replications). (b) Inconsistency across training with standard error curves below and above the mean (over 50 replications). (c) Number of activation cycles before a unique response, across training with standard error curves below and above the mean (over 50 replications).

reveals a strongly significant relationship (Spearman’s $\rho = 0.54$; $N = 50$; $p < 0.001$).

Finally, Hosenfeld et al. (1997) also found a critical slowing down in children’s solution times accompanying the onset of the sudden jump in correct responses. If we take the number of cycles necessary before the network settles into its final response as a measure of solution time, then the same pattern can be observed with the behavior of the network. Figure 10c shows the mean number of activation cycles required to settle over training. The number of cycles peaks around 2,300 epochs, in the neighborhood of the “sudden jump,” and across networks this correlates significantly with the occurrence of the sudden jump (Spearman’s $\rho = 0.285$; $N = 50$; $p < 0.05$).

3.3.7. The relative difficulty of cross-mapped analogies.

Children (and adults) can complete analogies appropriately even when there is a strong conflict between object similarity and relational similarity (e.g., Gentner et al. 1995). This is most clearly shown in cross-mapping situations, where the same object appears in both the base and the target but with a different role. In repeated studies (see Gentner et al. 1995), children have solved cross-mapped analogies. However, these analogies are hard, and older children perform considerably better than younger children.

Because the current model is designed to capture performance in the Goswami and Brown (1989; 1990) and Rattermann and Gentner (1998a) studies, with analogies consisting of two objects, the exact cross-mapping experiments presented in Rattermann et al. (1990) and Gentner and Toupin (1986) cannot be directly simulated. However, Figure 11 illustrates how the network can be tested on an analogue of the three-object cross-mapping experiments. There are two important aspects of the analogy presented in Figure 11: first, identical circles (B and R1) have different roles (B is larger than A, whereas R1 has the same role as C); and second, there is response competition between the literal similarity response (R1) and the relational response (R2). Thus, analogies of the type presented in the figure constitute a genuine test of cross-mapping.

We trained a modified network on a different set of stimuli in order to simulate a cross-mapping situation. The training environment consisted of three objects composed of nine units and three causal agents and associated

Figure 11. An illustration of cross-mapping with the relation *larger than* (the network was not tested on this analogy). The same object appears in both the base and the target domains, but with a different role. The network has to choose between the analogically inappropriate identical object (R1 on the left) and the correct response (R2 on the right).

transformations. Each object could be transformed when presented in conjunction with two out of the three causal agents. Training and testing were conducted in the same way as in the other simulations with the same network parameters and architecture. Crucially, one object vector when transformed by one causal agent was identical to a different untransformed object, leading to the potential for response competition. Given its environment, the network could be tested on six possible analogies, with one cross-mapped analogy.

Figure 12 shows the results for cross-mapped and non-cross-mapped analogies, averaged across 20 replications. After 15,000 epochs of training the cross-mapped analogy was completed almost with complete accuracy. However, the networks failed to complete any cross-mapped analogy in any replication before approximately 9,900 epochs. This is in contrast to the non-cross-mapped analogies, which were solved substantially earlier in training, reaching close to 100% performance after approximately 6,400 epochs of training. The maximum first derivative occurred 950 epochs earlier for the non-cross-mapped analogies than for the cross-mapped analogies (Wilcoxon signed-rank test: sum of positive ranks = 105; $N = 20$; $p < 0.001$).

In summary, the network is able to disregard all object similarity to solve analogies even when the same object representation occurs in different roles in the base and the target. The results also suggest that cross-mapped analogies are harder for the network to learn (although how difficult a cross-mapped analogy is to solve still depends on how easy it is for the network to learn the appropriate transformations). Both of these aspects of the model's behavior are consistent with the developmental evidence: that children can solve cross-mapped analogies (consistent with Goswami 1995) but that they are more difficult than other similar analogies.

3.4. The development of representations within the model

Many of the parallels with developmental phenomena observed in the network's performance arise from the

Figure 12. The network's performance on normal and cross-mapped analogies (20 replications). The thin line is the average percentage correct of the non-cross-mapped analogies. The thick line is the average percentage correct for the cross-mapped analogy.

interaction between the contrastive Hebbian learning algorithm and the regularities in the training set. It is worth considering how this interaction affects the network's internal representations over training.

The learning task facing the network can be seen in terms of two distinct processes: auto-association [replicating the pattern of activation at Object(t1) at Object(t2)] and transformation [producing a transformed version of the Object(t1) activation at Object(t2)]. In fact, auto-association can be viewed as a transformation encoded by a vector filled with 0s. Because this auto-association (or null transformation) occurs more frequently than the other transformations, the network initially learns auto-associations, replicating the activation at Object(t1) at Object(t2) and ignoring the transformations. It is this aspect of the learning algorithm that results in object similarity errors early in training instead of the less frequent transformation responses. Later in training, the network has learned to perform both auto-association and transformations. Consequently, the learning algorithm encodes the transformations given the appropriate conjunctively active CA(t1) activation pattern. This allows the network to build up more complex internal representations which support the production of the desired activation pattern at Object(t2) and which are a prerequisite for the appropriate completion of analogies.

The shift in strategy is also observed in the changes that take place in the network's internal representations at different points in training. Although the network is made up of bidirectional connections, information about a state at a given time [i.e., Object(t1) and CA(t1)] must pass through the hidden layer before it has any impact on the Object(t2) and CA(t2) response units. Hence, the hidden units provide an area in which information about objects and causal agents can be combined. The network's output [i.e., Object(t2) and CA(t2)] is driven by how the network organizes (represents) the object and causal agent information at the hidden layer.

Figures 13 and 14 show the location, in the first two principal components space, of the hidden unit activations for each possible Object(t1) and CA(t1). The grey ellipses illustrate the clustering of hidden unit activation following presentation of every pattern at Object(t1) alongside each pattern at CA(t1). After 100 epochs of training (Fig. 13), the hidden unit representations group the inputs according to the pattern at Object(t1). For example, the 1s (i.e., 1a, 1b, 1c, 1d) are grouped together, as are the 2s, the 3s, and so on.

Figure 14 shows the same analysis after 5,000 epochs of training (i.e., after the behavioral spurt observed at around 2,400 epochs). This suggests that the hidden units no longer group the inputs on the basis of Object(t1) features. Instead, the network has developed more complex internal representations. The different shaded ellipses illustrate the clusterings of hidden unit activations for the different causal agents (a–d). There is considerable overlap for the clusterings corresponding to causal agents a–d. The overlap reflects the fact that the 2-dimensional display is no longer sufficient to illustrate the complexity of the separation that is embedded in a 40-dimensional space.

We used a *k*-means cluster analysis technique to explore the representations that exist in this higher dimensional space. This technique is appropriate because we have strong prior theoretical reasons for investigating whether

Figure 13. A two-dimensional principal components analysis of the hidden units after 100 epochs of training. Each dot represents the response of the hidden units following presentation of a specific object and causal agent pair. Each object is represented by a number (1–20), and each causal agent is represented by a letter (a–d). Thus, the point marked by 5d is the location in the first two principal component space of the hidden unit activation resulting from presenting the fifth object and fourth causal agent to the network.

the network's hidden units cluster into four different groups according to causal agent. Table 2a shows the *idealized* outcome of a *k*-means cluster analysis showing perfect grouping by causal agent (when four groups are specified). Table 2b shows the actual results of the *k*-means analysis after 100 epochs of training. The fit of the model at 100 epochs was assessed against the null hypothesis that each cell has a count of 5. This analysis revealed that the model clustering did not differ from that expected by chance ($\chi^2(9) = 11.2$; $p > 0.25$). Thus, at 100 epochs, the hidden units do not cluster by causal agents. This is consistent with the principal components analysis showing that at 100 epochs the network groups events by object similarity and not by causal agent. However, Table 2c indicates that after 5,000 epochs of training the hidden unit activations group largely by causal agents. For each causal agent there is a separate dominant cluster similar to the idealized results in Table 2a. Frequency counts for the model cluster data correlated highly with the idealized data (Spearman's $\rho = 0.712$; $N = 16$). Thus, later on in training the hidden units are grouping the input according to causal agent, and consequently, are representing the transformations and not just the object attributes.

One way of describing how the internal representations change over training is to say that the learning algorithm “pays attention” to different aspects of the environment and develops different representations over time. Initially, the learning algorithm considers only the coarse object patterns irrespective of transformations (i.e., relations), whereas later in training the learning algorithm pays increasing attention to the transformations. This characterization of the network's behavior parallels Gentner's (1989) proposed explanation for the developmental changes observed with the relational shift. She suggested

Figure 14. A two-dimensional principal components analysis after 5,000 epochs of training. The grey ellipses illustrate the clusters of hidden unit responses to different causal agents. The darker regions represent the overlap between clusters.

that children's changing performance in analogical reasoning and metaphor comprehension tasks results from changes in what children pay attention to (from object attributes to relations) and how the objects and relations are represented. Note, however, that one of the strongest implications of the analogy as relational priming account is that the relational shift does not arise simply from a maturing system shifting from generally representing the world in terms of objects to representing the world in terms of relational systems. Instead, the apparent relational shift is a consequence of acquiring greater and richer relational knowledge – as suggested by Goswami's notion of “relational primacy” (Goswami 1991).

In summary, the work presented so far accounts for the development of early analogical reasoning in young children in terms of simple priming mechanisms and increasing world knowledge. The current framework stresses the importance of the interplay between the learning mechanism and the environment in determining not only the final representations, but also the developmental trajectory of how the network represents objects and transformations.

Of course, explaining the simple analogies used to test young children in terms of relational priming begs the question of how such a mechanism might explain the more difficult and complex analogies used to test adults – analogies on which young children typically fail. Thus, in the next section, we ask whether a priming-based account can be extended to account for complex analogies with multiple objects and multiple relations typically used to assess adult analogical reasoning. Our aim here is *not* to develop a full theory of all aspects of analogical reasoning in adults, but rather, to demonstrate the potential of our relational priming framework for modeling adult performance.

4. Analogies with multiple objects and multiple relations

In our view, analogical reasoning is actually something of an umbrella term referring to several different cognitive

Table 2. (a) Idealized results of a *k*-means cluster analysis showing perfect clustering by causal agent. The clusters are represented as numbers at the top; the different causal agents are the letters A–D. Each cluster contains all the hidden unit activations corresponding to a causal agent. (b) Actual clustering after 100 epochs. (c) Actual clustering after 5,000 epochs (where the results closely resemble the idealized results)

	1	2	3	4
A	0	0	0	20
B	0	0	20	0
C	0	20	0	0
D	20	0	0	0

(a)

	1	2	3	4
A	3	6	4	7
B	3	6	3	8
C	3	6	3	8
D	3	6	4	7

(b)

	1	2	3	4
A	1	0	2	17
B	3	1	16	0
C	1	12	7	0
D	15	3	2	0

(c)

Note. WW2 = World War II.

processes working in concert and heavily dependent on specific task demands. Our position echoes earlier approaches. For instance, Goswami (1991) suggests that analogical problem solving and *a:b::c:d* analogical completion may, to some degree, tap distinct cognitive processes. We maintain that complex analogies involving systems of relations and simple analogies involving relational priming may use similar underlying memory processes (e.g., pattern completion and relational priming) in considerably different ways as elaborated in section 4.1.3.

The following simulation exemplifies our approach with an account of how relational priming could build up an analogy between the Gulf War in 1991 and World War II involving multiple objects and networks of relations. This simulation is not intended to be a definitive account but instead is intended to illustrate that there is no essential conflict between the relational priming account (with additional cognitive control) and complex analogical reasoning. However, we acknowledge that there is considerable future work ahead of us before the model presented here could capture all the subtleties of fully fledged adult analogy.

In this simulation we address how the relational priming account could explain: (1) networks of relations (through iterative unfolding); and (2) one-to-many mappings (through additional controlled inhibitory mechanisms). The core theoretical mechanisms of the model (i.e., relational priming and relations implemented as transformations, pattern completion, gradual adaptation of connection weights) remain from earlier simulations. The main difference between the simulations is in additional controlled inhibitory processes present only during testing of analogical reasoning.

4.1. The model

4.1.1. The model architecture. There are two architectural differences from that of the previous simulations (see Fig. 15):

(i) *A single context layer instead of two causal agent layers.* In earlier simulations it was useful to talk about a causal agent in order to make the relationship between the model and the Goswami and Brown paradigm as transparent

as possible, the idea being that children learn about one object which is transformed (the apple) and another object which is instrumental for the transformation (the knife) but which does not undergo a transformation itself. It was therefore logical to represent the causal agent at both time step 1 (i.e., the “before” state) and time step 2 (i.e., the “after” state). However, the “before” and “after” causal agent layers can be collapsed into a single layer (in terms of the network architecture, the single [context] and duplicate [causal agent] representations are functionally identical). This single layer can, more generally, be thought of as representing the context in which the auto-association or transformation of an object (e.g., an apple) occurs. The principal benefit of using a single context layer is that it allows greater flexibility of the types of situations that can be represented and therefore simulated. For instance, instead of the Object(t1) and Object(t2) layers representing different temporal states, these layers could equally be interpreted as representing different objects (e.g., *robin* and *bird*) at a single time point. In this case the context layer would represent a more abstract relational label such as *ISA* (i.e., robin *ISA* bird); see Rogers and McClelland (2004) for further discussion.

(ii) *Additional inhibitory connections.* In addition to receiving the normal input via connections from other layers, in the adapted network all external units (i.e.,

Figure 15. Schema of the adapted model architecture. Arrows with round ends represent inhibitory connections. Object 1 and Object 2 layers = 12 units each; context layer = 6 units; hidden layer = 15 units.

Object 1 and Object 2 and the context layers) can also be selectively inhibited (i.e., switched off). This inhibition only occurs during analogical completion (testing), not during knowledge acquisition (training), and can be understood in terms of inhibitory connections from an additional control system (see Davelaar et al. 2005). Relational priming combined with selective inhibition of either context or object layers underlies how the network can demonstrate complex analogical completion representative of a wider variety of behaviors.

4.1.2. Training. In the current simulations, the training environment implements a version of an analogy comparing the Persian Gulf War of 1991 with World War II. As such, the 11 object prototypes in the training set group into two categories: (i) five from the Persian Gulf: these are *Saddam Hussein*, *George Bush Senior*, *Iraq*, *Kuwait*, and *the United States*; and (ii) six from World War II: *Adolf Hitler*, *Winston Churchill*, *Germany*, *Poland*, *Austria*, and *the Allies*. We chose these two domains (the Gulf War and World War II) because they have previously been the focus of analogy research (e.g., Spellman & Holyoak 1992), and can support analogies involving networks of relations.

In contrast with our earlier simulations, object prototype representations are binary vectors (i.e., each unit is either on or off). Each object representation consists of two components, an orthogonal component uniquely identifying the object (this can be thought of as the object's label or name) and an orthogonal component representing which category the object belongs to (i.e., either Gulf War or World War II). Although, these object representations are highly simplified, they do permit a straightforward illustration of how relational priming is compatible with complex analogical reasoning.

In addition, there are six distinct orthogonal context representations, each consistent with a different relation. The network is trained (for 2,000 epochs) on only a subset of contexts and objects designed to reflect some of the relational structure underlying the analogical comparison of the Gulf War and World War II (see Table 3). As in previous simulations, each prototype pattern presented

to the network has Gaussian noise added ($\mu = 0.0$; $\sigma^2 = 0.01$).

4.1.3. Testing. In testing, there are three processes involving the proposed inhibition mechanism that, working in concert, illustrate the iterative unfolding of a complex analogy. At each time step some combination of the object or context layers is inhibited (this could be thought of as shifting conscious attention between different contexts and objects), and the network falls into a new attractor state with either the new context or a new object. We consider each of the three proposed processes in turn.

(1) *Selecting a new object with the same role (i.e., relation) in the parallel (or target) domain.* This process involves finding a different pair of Object 1 and Object 2 representations that correspond to the same context representation. This process can be thought of as giving the network two objects and then asking what other two different objects have the same relation, that is, $a:b:c \dots ?$ with the response $c:d$, in contrast to previous $a:b:c \dots ? d$ analogies.

After presentation of Object 1 (e.g., *Iraq*) and Object 2 (e.g., *Kuwait*) and settling into an attractor state (e.g., with the context representation *occupies*), the Object 1 and Object 2 representations are inhibited. This inhibition reflects the active, volitional search that is involved in completing some complex analogies. The pattern completion properties of this type of recurrent network ensure that the network subsequently settles into an attractor state consistent with the prior context representation (*occupies*) but with different Object 1 and Object 2 representations (e.g., *Germany* and *Austria*, respectively). This allows the network to form analogies from limited information.

In order to traverse a network of relations, building up a large multipart analogy, the network has to find new relations for consideration in a controlled way. The following two processes address this issue:

(2) *Selecting a new role in the same domain for the same object.* One possibility for finding a new relation is to simply inhibit the context layer and allow the network to settle into a new attractor consistent with the Object 1 representation and a new context representation. This can be understood as: Given an Object 1 activation (e.g., *Iraq*), what other relation (other than *Occupies*) goes with this object? In response, the network would produce the context representation *Defies* and the Object 2 representation *United States*.

(3) *Selecting a new role in the same domain for a different object.* It is also possible to find a pair of different objects connected by a different relation. This involves inhibiting the previous Object 1 and Object 2. This process results in a pair of new objects and a new relation (e.g., Bush *Motivates* United States), which are from the same domain as the prior objects. Both processes (2) and (3) can subsequently be used to form a new mapping with the other domain, using process (1) above.

4.2. Results

The three processes detailed above can be used sequentially to explore similarities between two domains involving multiple relations and objects to build up a large, complex analogy. Figures 16 and 17 demonstrate the network's performance on one such analogy using the Gulf War and World War II domains (the network's exploration

Table 3. *The objects and relations that the network was exposed to over training*

	Object 1	Relation	Object 2
Gulf War	Bush	Motivates	United States
	Saddam	Threatens	Kuwait
	Bush	Orders_attack_of	Iraq
	Saddam	Dictator_of	Iraq
	Iraq	Defies	United States
	Iraq	Occupies	Kuwait
WW2	Churchill	Motivates	Allies
	Hitler	Threatens	Poland
	Churchill	Orders_attack_of	Germany
	Hitler	Dictator_of	Germany
	Germany	Defies	Allies
	Germany	Occupies	Poland
	Germany	Occupies	Austria

Figure 16. An example of the network's trajectory through "conceptual" space. The two grey boxes represent the different domains (i.e., Gulf War, bottom left; World War II, top right). The x,y -coordinates of each point correspond to the maximally active unit in the Object 1 and Object 2 layers, respectively, at different time steps. Initially, Gaussian noise is presented to the network, and the network randomly falls into the activation state corresponding to Saddam, Kuwait in the Gulf War domain. On each subsequent step, the inhibitory processes described in (1), (2), and (3) drive the network to different points in the conceptual space. First, process (1) finds a match in the World War II domain (Hitler, Poland). Then process (2) finds a different transformation in the same World War II domain (Hitler, Germany) with a different relation [dictator_of], with process (3) subsequently finding the appropriate mapping in the Gulf War domain (Saddam, Iraq). This interplay of relational priming with processes (1), (2), and (3) continues until the full set of mappings between the two domains has been explored.

of the analogy presented in these figures is one example out of many possible trajectories). Figure 16 shows the network's journey through the space of multiple concepts as it discovers the multirelation analogy, whereas Figure 17 shows the successive activation states of the network as analogical completion unfolds through time. First, the network is given a domain as the base by presenting the network with Object 1 and Object 2 layers consisting of the component representing the domain (i.e., *Gulf War*) and resting values for all other Object 1 and Object 2 units. The network then completes the Object 1, Object 2, and context layers consistent with this domain (i.e., *Saddam threatens Kuwait*). This activation state constitutes the starting point of the analogy. Subsequently, process (1) is used to find a mapping in the alternative domain (i.e., *Hitler threatens Poland*). Following this, either process (2) or process (3) is used to find a different relation involving the same Object 1 representation (i.e., *Hitler dictator of Germany*). Process (1) is then repeated to find a mapping in the Gulf domain (i.e., *Saddam dictator of Iraq*). Subsequently, process (1) and either process (2) or (3) are interleaved iteratively to traverse the network of relations, thereby building up a complex analogy (see Figs. 16 and 17). Importantly, this model reveals how relational priming may serve as a fundamental subprocess when building analogies involving

Figure 17. The activation states of the units of the network in the Object 1, Object 2, hidden, and context layers across different time steps 1–12 – tracing out the unfolding of the complex analogy set out in Figure 16. The enlarged dark grey circles indicate the maximally active unit for either the Object 1 or Object 2 layer. The network was tested following 2,000 epochs of training.

networks of relations and many-to-many comparisons across domains.

4.3. Discussion

In the current implementation, iterative unfolding is guided by an essentially blind process of relation selection (i.e., processes (2) and (3) above involve finding a different random relation or object). This, however, is not to say that iterative unfolding is necessarily an unguided bottom-up process. One way in which top-down information can play a strong role is in terms of systems of relations and the way that these are gradually encoded in the network's internal representations. In the network's training environment, systems of relations will co-occur with greater frequency and across different examples over individual relations that do not co-vary consistently. Connectionist networks are good at picking up these statistical regularities and representing them in the hidden units, such that systems of relations are represented as closer in hidden unit space. Any consequent analogical mappings (i.e., given a certain relation) made with these hidden units will thus contain a bias for selecting a new relation from within the same coherent relational system. Rogers and McClelland (2004) discuss the importance of this type of consistent covariance for the related but distinct development of semantic cognition. Here, we see how this consequence of the statistics of the input could explain biases for systems of relations which have been repeatedly reported in the analogy literature, and, more importantly, could explain how and why these biases develop.

Importantly, a process of iterative unfolding also resonates with evidence as to how children learn about relational structures. Children's analogy abilities become more systematic and sophisticated as they gradually absorb and internalize more of the richness of the structural relations (particularly causal systems of relations) in the biological, physical, and psychological domains in which they find themselves (for a review, see Goswami 2001).

We do not consider in detail how the postulated additional control processes are implemented in a full connectionist account, or how these processes develop – this is beyond the scope of the current paper (although certainly not beyond the scope of connectionist modeling, e.g., Davelaar et al. 2005; O'Reilly 2006; Rougier et al. 2005). Tentatively we suggest that the work of O'Reilly and colleagues provides a useful template for understanding how flexible cognitive control systems could develop in response to task demands (see Rougier et al. 2005). This work considers the interaction of a network that can rapidly update or maintain activation according to a dynamic gating mechanism trained with reinforcement learning. Task-relevant representations self-organize over training. These representations are capable of capturing data from Stroop and Wisconsin-card sorting tasks – both tasks that require the kind of active maintenance and inhibition necessary for the proposed iterative unfolding of analogies. Furthermore, one of the interesting aspects of the Rougier et al. work is that the network is able to generalize to novel tasks. Therefore, such an approach is a useful *starting point* for investigating the origins of the controlled iterative unfolding that we propose for complex analogical mapping and how such skills generalize to analogies in novel domains, although there are likely to be considerable challenges ahead in satisfactorily marrying the simple analogical processes detailed earlier with the controlled processes in a single developmental model.

5. General discussion

The account of analogical completion presented here is an attempt at a novel developmental theory of analogical reasoning. Contrasting our model with models inspired by adult reasoning (e.g., SME, LISA) indicates that taking a developmental perspective constrains the modeling process in a substantially different way. Unlike most accounts of adult analogical reasoning, which are largely concerned with analogies involving systems of relations, our developmental account has been forced to emphasize different qualities such as knowledge acquisition and the importance of reconciling analogy competence with other lower-level cognitive processes. This difference in emphasis results in an explanatory framework that brings together a wide range of seemingly disparate developmental experimental findings in a single computational model. In contrast to other models of analogy that have been applied to developmental phenomena, the model presented here is able to account for all of the key developmental phenomena listed in the introduction. As such, it is worth reflecting on some of the theoretical implications of the model's two underlying core assumptions: the central role of priming and the treatment of relations as transformations.

5.1. Relations as transformations

5.1.1. The relational shift revisited. The key benefit of viewing relations as transformations between states is that relations do not have to be represented explicitly, avoiding the difficulties of learning explicit structured representations. By attempting to simulate the acquisition

of relational knowledge, the current account is consistent with and provides a possible developmental mechanism for knowledge accretion theories of analogical development. This is especially clear for simple causal relations (e.g., *cutting*) such as those considered with the model, where there is a direct physical change in an object. However, as emphasized in the final simulation, potentially all relations could be viewed as transformations.

Our account of the relational shift depends crucially on how well the network learns object transformations versus auto-associations (i.e., null transformations). At the heart of the relational shift is a change from ignoring transformation information (auto-association) to incorporating transformation information into the network's internal representation of the object. The rate at which this occurs reflects the relative mix of transformation versus non-transformation experiences that the network encounters, thereby grounding the relational shift squarely within the experiences that the child encounters. This position contrasts markedly with Gentner's Structure-Mapping account, which posits that the relational shift derives from a process of knowledge becoming increasingly abstract. Instead, our position resembles far more the relational primacy hypothesis of Goswami.

One of the more unexpected consequences of the current account as a developmental theory is that it implies that there is no necessary relational shift for any given relation in a child's similarity judgments. The relational shift is modulated by the frequency with which children encounter transformations versus auto-associations (i.e., null transformations). This even implies that for high-frequency relations children might show an initial bias for transformational similarity over object-based similarity. For the network exposed to high-frequency transformations, this early preference for transformation over auto-association would correspond to it being able to solve analogies but not non-analogies (i.e., the network would make errors like *apple:apple::bread:cut bread*). For such high-frequency relations the network would, accordingly, undergo something akin to a relational shift in reverse, gradually becoming able to produce the object-similarity response and producing the transformation response only when appropriate rather than all of the time. Thus, the theory and model strongly predict that with a subset of highly familiar relations children will, in contrast to the expectations of the relational-shift hypothesis of Gentner (1988), choose relational responses over object-similarity responses even when an object-similarity judgment would be more appropriate. A useful test of the validity of the model would investigate the relational judgments of children to similar stimuli that differ in the appropriateness of a relational or object-based interpretation. Following the analogy as relational priming account, we would expect that over development there should be a shift away from relational responses for the object-based stimuli – an inverse relational shift.

5.1.2. Transformation size in analogical reasoning. One factor that affects how well the network learns a transformation is the degree of confusability or overlap between different input-output training patterns (e.g., how similar are the representations between *apple* and *cut apple*). The lesser the overlap between the auto-associative (e.g., *apple*) and the transformed (e.g., *cut apple*) responses

(i.e., the greater the size of the transformation) the easier it is for the network to learn the pattern sets that define the transformation. As a consequence, the network will also be better at completing analogies involving states that are very different as compared to those that do not differ very much. The model therefore predicts that analogies involving larger transformations will be solved more readily than analogies involving smaller transformations. This prediction is a consequence of our two key theoretical assumptions (analogy as relational priming and relations as transformations), and their implementation within the network model. No other model of analogical reasoning makes a similar prediction.

In order to explore this prediction, Leech et al. (2007) tested adolescents and adults on analogies where the transformations involved varied along some dimension. Consistent with our model's predictions, it was found that for both adults and adolescents, analogies involving larger transformations were solved more accurately than analogies involving smaller transformations.

5.1.3. Transformation and predicate structure. Some readers may argue that relations are not really implemented as relations, but, instead, we have built predicate structure into the network architecture, with the Object(t1) layer being the patient and the CA(t1) (or context) layer being the agent or instrument. However, this is not the case. Any layer (or slot) of the network can be trained as a patient or an agent in a transformation. Indeed, architecturally the CA(t1) layer and Object(t1) layer are not functionally different. For example, we could train the model on *knife* at CA(t1) and *broken knife* at CA(t2) (in addition to training the network that *knife* cues a physical transformation in *apple*). Then a single network could solve analogies involving *knife* as agent and *knife* as patient.

5.2. Analogy as priming

Central to the question of whether priming can play an important part in analogical reasoning is whether analogy (and related underlying mechanisms) need be explicit. Priming is normally considered as an automatic, implicit mechanism (e.g., Tulving & Schacter 1990), whereas analogy is generally characterized as an essentially explicit ability. Therefore, analogy as relational priming seems at first glance to be a somewhat paradoxical theoretical position. However, we believe that this argument has only superficial validity, primarily because it is important to distinguish between cognitive processes and the behavioral and cognitive results of those processes. For instance, although priming – an implicit process – may be a core mechanism of analogy, the result of that process can still be explicit and accessible to the system (i.e., the resulting analogy can be verbalized or used as input to some other cognitive process). Language comprehension provides an illustrative example of a domain that involves some implicit mechanisms (including semantic priming) and verbalizable explicit outcome (e.g., see Kutas & Federmeier 2000). Secondly, our account of analogy as relational priming is compatible with both implicit analogical reasoning mechanisms (as suggested by the earlier simulations of analogical completion with uncontrolled relational

priming) and with much more deliberative mechanisms (i.e., the controlled use of inhibition and relational priming to iteratively unfold a large and complex analogical mapping).

Priming effects are ubiquitous in perception and cognition and have been observed across development. Viewing analogy as related to a form of priming demonstrates how a high-level cognitive process can be rooted in lower-level processes. This provides a possible explanation of the early and natural use of analogical reasoning by young children. The proposed relationship between analogical reasoning and more basic priming mechanisms can be understood as analogous to how some researchers have related cognitive skills (including reasoning and categorization) to more general, lower level processes of inhibition (e.g., Dempster & Brainerd 1995; Houdé 2000).

Consistent with this perspective, we envisage analogy as something of a heterogeneous phenomenon: Different types of analogy will utilize a variety of memory and control processes in considerably different organizations and to considerably different effect. While our aim has not been to provide an exhaustive account of analogical reasoning, we have suggested two possible configurations of cognitive processes that may underlie distinct types of analogical reasoning. These configurations illustrate the explanatory power of analogy as relational priming, and demonstrate how the framework is compatible with complex adult levels of performance.

Although there is strong evidence that relational priming may be involved in analogy, further work needs to be done to establish how well the two processes are interrelated. In fact, the analogy as relational priming account stands or falls on the predicted intimate relationship between analogy and priming, especially through development. In this vein, a strong test of our account concerns whether relational priming can be demonstrated in children and if so whether this correlates with performance on more standard analogical reasoning measures.

5.3. The role of world knowledge

Both children and adults can readily form analogies involving novel stimuli. They do this rapidly and following very little exposure. This may appear to be inconsistent with the lengthy training regime and slow connection weight adjustments that form the basis of analogical reasoning in our simulations. This is not the case. Even stimuli that appear novel often share a great deal of underlying information with prior experience. For example, novel square drawings moving on a computer monitor will tap into considerable existing world knowledge about the possible relations between moving items. More generally, when the network is presented with a “novel” problem, the task will be made easier if the network can co-opt existing representations into learning the new problem (e.g., Altmann 2002; Shultz & Rivest 2001; Shultz et al. 2006). A similar explanation of why children can rapidly draw causal inferences about novel objects that they have never encountered before is given by McClelland and Thompson (2007), who demonstrate that providing a network with

substantial previous experience in a domain can lead to rapid “one-shot” causal learning.

5.4. Explicit mapping

In our final simulation we demonstrated how relational priming could be used deliberately to build up a complex explicit analogy. The final simulation illustrates how a situation that is normally assumed to be an example of explicit structure mapping is consistent with a simpler conception of analogy combined with meta-cognitive processes used to elaborate an initial mapping. We acknowledge that the current adult model is underdeveloped and in particular fails to explain the etiology of the control and memory processes necessary for handling complex adult analogies; and that a fuller account of adult analogy would, therefore, be necessary to convince many researchers in the structure-mapping community of the explanatory power of the relational priming account of analogical reasoning. However, we believe that the iterative mechanism for building up mappings based on relational priming is useful for illustrating some important distinctions between our account and existing models.

In our view the elaboration of an initial mapping is both deliberative (i.e., non-automatic) and task-directed. To demonstrate, reconsider the World War II/Gulf War analogy. Holyoak and colleagues (e.g., Holyoak & Hummel 2001; Spellman & Holyoak 1992) have shown that people can come up with complex mappings between 1990 and 1939 (e.g., Saudi Arabia is comparable to France; Kuwait is comparable to Poland; George Bush could be Roosevelt or Churchill). However, most of these explicit mappings are irrelevant for the analogy to fulfill its intended purpose of stressing the similarity between invasions from Hitler's Germany and Saddam Hussein's Iraq, and consequently emphasizing that if Saddam Hussein were not stopped something terrible – like World War II – would be visited on the world. The analogy still holds if only Saddam Hussein and Hitler (both conceived as dangerous, aggressive dictators) are compared and nothing else is mapped. Thus, in this situation at least, there does not need to be an explicit mapping of relational structure to form an analogy (though making such structure explicit may well serve to increase the intensity of any argument based on the analogy).

A mechanism of iterative unfolding, such as the illustrative one presented in the final simulation, also enables a fuller comparison of the current model's scope with respect to other phenomena such as the systematicity effects shown in conjunction with Structure-Mapping Theory. As we noted in the previous section, relations that constitute a system will vary coherently in a naturalistic training environment. This means that a connectionist network learning about that domain will develop internal representations that reflect this relational structure (something that developing children also do – see Goswami 2001). Consequent inferences or analogical reasoning based on these internal representations would, therefore, contain a bias towards systems of relations (see Rogers & McClelland 2004; 2005; Thomas & Mareschal 2001).

How world knowledge is acquired is also of particular importance in determining the kinds of relations that prime others. In the adult relational priming literature, priming effects have typically been quite small. These

small effects (although semantic priming effects are typically larger in children than adults; Chapman et al. 1994) could in part result from the fact that experimenters have focused on very general relations defined in linguistic taxonomies (e.g., the relation *have*; for a discussion, see Estes & Jones 2006), rather than on the more specific types of relations that are salient and useful for explaining the world. This parallels the evidence from Goswami and colleagues that causal explanations are particularly relevant to children in making sense of their environments, because many real-world events feature cause-effect patterns (see Goswami 2001). It follows that in order to further develop our account of iterative unfolding with systems of relations it will be necessary to provide training regimes that better reflect a child's environment, and, in particular, training regimes that emphasize the relations and relational structures that are most salient within the child's environment, and as such are most likely to serve as primes.

We have so far not discussed the important issue of relational complexity (Halford et al. 1998) and how the iterative unfolding account could marry simple proportional analogies such as $a:b::c:d$ with the developmental effects of relational complexity on children's reasoning. However, to address this issue, one open question that would first need to be broached is how n -ary relations would be implemented (e.g., relations with three arguments such as *gives* in *John gives Mary the book*). One simple way, following event semantics (Davidson 1967), to generalize our account to analogies involving n -ary relations would be to decompose an n -ary relation into multiple binary relations around an event: for example, GIVER(event, John), GIVEE(event, Mary), and GIVEN(event, book). Higher arity relations, according to the iterative unfolding account, would then require more temporary storage of partial results (e.g., dealing with the ternary relation *give* would require temporary storage of the giver, the givee, and the given). Hence, this approach would also predict that relational complexity would constrain analogical reasoning and that, across development, ability with higher arity relations will correlate with some measure of working memory efficacy.

The important distinction underlying our approach to more complex analogies is between the explicit mapping framework where analogies are worked out completely by some cognitive mechanism (e.g., LISA), and an alternative view where simple analogies are first made before subsequently being checked and expanded if necessary using explicit iterative unfolding. To reiterate, according to our account, explicit structure mapping is a meta-cognitive skill: a relational priming mechanism reveals a relational similarity between two domains, but the reasoner can iteratively unfold this by repeatedly applying the simpler mechanism over and over again to components of a domain in order to extend the analogy or to discover where the analogy breaks down. This second account has an important role for explicit mapping; the key difference is that explicit mapping is no longer *necessary* for analogy to occur, but instead describes a subset of analogies.

Future work involving, for example, patients following neurological insult, or possibly transcranial magnetic stimulation with healthy participants, could provide strong evidence for the disentangling of mapping from analogy. In particular, given that explicit mapping is likely to employ more frontal regions, whereas relational priming is likely to be more temporal, we predict that

frontal damage should have relatively little impact on analogical reasoning when explicit mapping is not central to performance, and that more temporal damage should severely reduce relational priming and consequent analogical reasoning.

One interesting prediction from the model that also suggests a dissociation of analogy and mapping concerns a developmental asymmetry in analogy completion when base and target are reversed. For each analogy, it is predicted that there will be a period when the model can complete an analogy one way but fails to complete its reverse. For example, *apple:cut apple::bread:cut bread* at a given point in development may be easier than the reverse analogy: *bread:cut bread::apple:cut apple*. This phenomenon arises in the model because pattern completion is differently constrained in the base domain and the target domain. The base domain involves greater external constraint (i.e., both the *a* and *b* terms) than the target domain (just the *c* term). Consequently, the model is more likely to appropriately complete an analogy if the less well learnt relation is in the more highly constrained base domain than if it is in the less constrained target domain. This prediction is hard to reconcile with structure-mapping accounts and so constitutes a further strong test of the validity of analogy as relational priming model.

5.5. Development revisited

One of the principal lessons from this work is that it is vital to place development squarely at the heart of any account of cognition. This is not a new proposal (e.g., see Karmiloff-Smith 1998; Mareschal et al. 2007; Piaget 1970; Thelen & Smith 1994) but one that is often overlooked by investigators of adult cognition. Many models of adult cognition have become very complex, often positing a myriad of specialist mechanisms, but are also very powerful at explaining many different aspects of adult performance on a range of complex tasks. However, in many cases, these models make no attempt to explain how the complex structures assumed to be part of adult cognition emerged. In contrast to this, we have emphasized the need to explain how cognitive mechanisms emerge over time with experience of the world. The result is that a very different kind of model is arrived at. As discussed in sections 4 and 5.4, our current account still has a substantial way to go to capture the complexity and richness of adult analogical reasoning. Indeed, in section 4 we sketch one possible way forward. That objection notwithstanding, it still remains for adult-level models to make contact with the developmental constraint: namely, that all proposed mechanisms must have a developmental origin in order to be plausible. Thus, while the developmental model does not reach adult levels of competence, the adult model does not make sufficient contact with its developmental origins. A complete explanation of analogical reasoning must breach this gap.

In summary, relational priming has been presented as a developmentally viable account of early analogical completion. We have shown that the account, implemented in a connectionist model, captures a broad range of developmental phenomena, including seven detailed developmental markers of analogical ability. Our final simulation demonstrates how the simple relational priming mechanism can be applied iteratively to traverse complex analogies. This approach promises to provide a fuller developmental

picture of the mechanisms underlying the gradual transition from simple to more complex reasoning.

ACKNOWLEDGMENTS

This work was supported by the Economic and Social Research Council (UK), MRC NIA grant G0400341 and European Commission NEST grant 020988 (ANALOGY).

NOTES

1. Unless otherwise stated, all figures shown in this section portray the raw network data after fitting by 15 b-spline functions with a lambda smoothing coefficient of 0.5 using the Functional Data Analysis MatLab toolbox (Ramsay & Silverman 2002). This results in much clearer trajectories that can be interpreted in terms of higher derivatives (see Shultz [2003] for a full discussion of the benefits of using this method of presenting network data for estimating developmental trajectories). This smoothing process does not change the essence of the results we present nor the arguments that we make based on these network simulations. We thank Tom Shultz for suggesting this additional analysis.

2. It is important to note that the sum of squared error at output is calculated for the network appropriately producing either a transformational or auto-associative response (which it is being trained on) – not analogical completion (which it is never trained on). As such, sum of squared error and percentage analogies correctly completed are logically independent, and non-parametric correlation is an appropriate statistical test to assess their relationship.

Open Peer Commentary

Relational priming: Obligational nitpicking

doi:10.1017/S0140525X08004470

Varol Akman

Department of Philosophy, Bilkent University, Bilkent, Ankara, 06800, Turkey.
akman@bilkent.edu.tr
<http://www.phil.bilkent.edu.tr/akman.htm>

Abstract: According to the target article authors, initial experience with a circumstance primes a relation that can subsequently be applied to a different circumstance to draw an analogy. While I broadly agree with their claim about the role of relational priming in early analogical reasoning, I put forward a few concerns that may be worthy of further reflection.

As someone interested in naturalized epistemology, I have benefited from reading Leech et al.'s target article for it argues that there is a need to study early analogical reasoning in a developmental framework. Although complex analogical reasoning – as observed in adults with proper mental makeup and education – may be explicated in the abstract (Falkenhainer et al. 1989), a developmental approach requires, more than anything else, psychological plausibility. In one of the better works in the latter regard, Thagard et al. (1990, p. 290) affirm that the computational model they describe is “intended to be a model of human cognition, so it is essential that it be consistent with the results of psychological experiments.” It is uplifting to watch Leech et al. pay careful attention to present-day cognitive psychology and never let the results of their work drift away from correlation with its accomplishments.

Leaving aside the connectionist implementation for others to comment upon, I think it is fair to say that the core of Leech et al.'s argument is the cognitive apparatus for an “A is to B as

C is to ...?" (denoted as "A:B::C:?" in what follows) analogy, which is customarily handled thus:

map A onto C

map B onto D (this is the unknown term to occupy the ... position)

transfer across the relation between A and B to C and D (so as to correctly determine D)

Their conjecture is that *relational priming* accomplishes this task neatly. To wit, exposure to A and B primes a relation R that then biases C to show the way to the appropriate D. To give an example from the simple causal domains of childhood, after seeing the triple "cake, slice of cake (i.e., cut cake), pizza," the analogically appropriate response would be "slice of pizza" (i.e., cut pizza). Here, the terms "cake" and "slice of cake" prime a relation ("cutting") that then biases "pizza" to produce the appropriate response "slice of pizza."

Relational priming is a commonsense notion, and its relatives are widespread. Estes and Jones (2006, p. 89) define *relation* [*sic*] *priming* as a phenomenon in which grasp of a word pair ("plastic spoon") is facilitated by the past presentation of another ("wooden plate") that instantiates the same conceptual relation ("made of"). Recanati (2004, pp. 148–51), inspired by Wittgenstein, elucidates how someone learns a predicate P – a central concern in cognitive linguistics. The learner observes the application of P in a particular context (situation) K and associates P and K. In another context K', the learner will assume that P applies in case she discerns that K' is adequately "similar" to K (Akman 2007). If her judgment regarding the resemblance of K and K' is flawed (in other words, K and K' are similar in a way not pertinent for the application of P), then communal help would come and correct the learner. The learning phase of the learner boils down to noting enough contexts in which P is reasonably applicable. In some sense, we can say that P and K cause a "legitimate context of application" to be primed.

Unfortunately, the notion of relational priming is fairly problematic. It is true that an analogy problem such as "puppy:dog::kitten:?" can be solved rather easily with relational priming. Basically, exposure to "puppy" and "dog" primes the relation "offspring," which then biases "kitten" to produce "cat" as the D term. But the question arises: Why not the relation "younger than"? Let me borrow the notation of situation theory for a second (Barwise & Perry 1983) and say that if R is a binary relation and A, B are objects appropriate for the respective argument places of R, then we shall write $\langle\langle R, A, B, 1 \rangle\rangle$ to denote the informational item that A, B stand in the relation R. My point is that exposure to "puppy," "dog" could potentially prime several candidate relations, and it is no easy feat to decide which one of these is to be preferred. (It is a different matter to inquire whether this would affect the correct solution of the analogy problem.) In a nutshell, $\langle\langle\text{offspring, puppy, dog, 1}\rangle\rangle$ and $\langle\langle\text{younger than, puppy, dog, 1}\rangle\rangle$, $\langle\langle\text{cutter than, puppy, dog, 1}\rangle\rangle$, and so on, are all possible informational items, not to mention the more theoretical $\langle\langle\text{isa, puppy, dog, 1}\rangle\rangle$. Presumably, this problem does not arise in the aforementioned cake example because of the implicit presence of a causal agent (e.g., a knife), but even that is suspect. Surely, the primed relation could have been "larger than" in that example. (Or "more expensive than," if you truly want to make life difficult for a fan of relational priming.)

It is also crucial to observe that R has two argument roles, yet these "slots" into which appropriate objects can be placed are *not* ordered in any way – and "certainly not linearly ordered as a finite sequence" (Devlin 1991, p. 116). In the case of the relation "offspring," the argument roles can be termed as "ascendant" and "descendant." Accordingly, one apparent objection has to do with the "direction" of the relationship. Maybe "ancestry" (rather than "offspring") should be the primed relation as it does the job equally well; after all, one of the informational items $\langle\langle\text{ancestry, ascendant: dog, descendant: puppy, 1}\rangle\rangle$ and $\langle\langle\text{offspring, descendant: puppy, ascendant: dog, 1}\rangle\rangle$ is

redundant given the other. In all likelihood, young children would never use these abstract relations anyway; they would prefer a tangible relation such as "mother of." At least this is what my 5-year-old consistently does. Whenever she sees a lonely cat in the backyard, she says that the animal is seeking "the mom" (and never "the dad" or "the parents").

To make things more complicated, it is not clear whether other examples superficially similar to "puppy:dog::kitten:?" can be solved effortlessly. Take "rake:leaves::magnet:?" In this case, exposure to the first pair should presumably prime the relation "conveniently collects in one place (location)." Then the answer "paper clips" would be apparent. However, I trust that this must be a tough one for a child of five.

Finally, take the celebrated line attributed to Groucho Marx: "Military justice is to justice as military music is to music." When this is formulated as "military justice:justice::military music:?", it is straightforward to say that the first two terms prime the theoretical relation "ako" – representing one class being a subset of another – and hence the D term can be trivially worked out as "music." However, and I say this somewhat hesitantly, the correct understanding here starts from the reverse direction. In other words, what makes the analogy a striking one (one that works) is that we are supposed to know that military music is invariably dreadful as a harmonious experience and that military justice should fail in the same vein. Maybe another Groucho quip is in order: "A child of five would understand this. Send someone to fetch a child of five."

A neural-symbolic perspective on analogy

doi:10.1017/S0140525X08004482

Rafael V. Borges,^{a,b} Artur S. d'Ávila Garcez,^a and Luis C. Lamb^b

^aDepartment of Computing, City University London, Northampton Square, London, EC1V 0HB, United Kingdom; ^bInstitute of Informatics, Federal University of Rio Grande do Sul, Brazil, Porto Alegre, RS, 91501-970, Brazil.

Rafael.Borges.1@soi.city.ac.uk

<http://www.soi.city.ac.uk/~Rafael.Borges.1>

aag@soi.city.ac.uk

<http://www.soi.city.ac.uk/~aag>

LuisLamb@acm.org

<http://www.inf.ufrgs.br/~lamb>

Abstract: The target article criticises neural-symbolic systems as inadequate for analogical reasoning and proposes a model of analogy as transformation (i.e., learning). We accept the importance of learning, but we argue that, instead of conflicting, integrated reasoning and learning would model analogy much more adequately. In this new perspective, modern neural-symbolic systems become the natural candidates for modelling analogy.

The target article identifies two different stages as important for analogical reasoning: the learning of transformations (or relations) between objects, and the application of the acquired knowledge. The importance of learning, building a knowledge base, is highlighted in the article, as the analogy performance improves when more expertise is acquired. In what regards the application of knowledge for analogy, the process can be divided into two steps: (1) the recognition of the context, exemplified by the search for the most appropriate relation to be considered, and (2) the further reasoning over a different object, which might be seen either as a search for the most relevant target, or as the application of a transformation (as advocated in the article).

Among other works, the authors mention Shastri and Ajjana-gadde (1993), an important reference for the research on the integration of neural and symbolic artificial intelligence approaches. Since then, the research on neural-symbolic systems has evolved considerably with a strong focus on the integration

of expressive reasoning and robust learning into effective computational systems. Neural-symbolic systems might be considered nowadays as an alternative to traditional intelligent systems, enhancing the main features of trainable neural networks with symbolic representation and reasoning. Among the different neural-symbolic systems, we will consider those that translate symbolic knowledge into the initial architecture of a neural network, such as Connectionist Inductive Learning and Logic Programming (CILP) (d'Avila Garcez et al. 2002). In CILP, knowledge is represented by a propositional logic program, and the translation consists in setting up a three-layer feed-forward neural network in which input and output neurons represent propositional variables, and the hidden neurons represent clauses (rules) involving such variables, with the overall network presenting a semantic equivalence with the original logic program.

In light of these developments, we can propose a neural-symbolic approach to model the procedures involved in the experiments shown in the target article. Instead of focusing on the dichotomy between relations and transformations, we prefer to analyse each stage of the described analogical procedure as subsuming aspects of learning and reasoning. Regarding the stage of building the knowledge base, neural-symbolic systems would cater for empirical learning with the possibility of integrating background symbolic knowledge about the domain, which can improve performance as shown in Towell and Shavlik (1994) and d'Avila Garcez et al. (2002).

As for the application of analogical reasoning, in an "A:B::C:D"-like example, we can consider two steps. First, the recognition of the relation (or transformation) between A and B. Considering, again, a neural-symbolic representation of possible functions, the system should be able to receive two objects as input, recognise the relation R between them, and keep this information in a short-term memory for immediate use. This kind of mapping with the use of short-term memory can be found in neural-symbolic systems such as Sequential Connectionist Temporal Logic (SCTL) (Lamb et al. 2007), an extension of CILP catering for the representation of temporal logic programs. Once the relation has been identified, the process can be seen as a simple inference, as described above.

We can show the adequacy of neural-symbolic systems in representing this kind of reasoning considering the same example used in the target article. The initial learning step should build the knowledge represented by a set of clauses, where we are considering two different atoms (propositional variables) for representing the relations: Recognise_R, to represent when a relation R is recognised between two instances, and Apply_R, illustrating the application of R over an instance of the input. The set of clauses for the example is expressed below:

```
Cut_Apple if Whole_Apple and Apply_Cut;
Recognise_Cut if Whole_Apple and Cut_Apple;
...
Cut_Cheese if Whole_Cheese and Apply_Cut;
Recognise_Cut if Whole_Cheese and Cut_Cheese.
```

Also, for the application of this case of (A:B::C:D), the agent should have some knowledge about the task, which can be represented in temporal propositional logic by inserting a clause "(next) Apply_R if Recognise_R," for each relation "R," where "(next)" is the operator referring to the next time point. Therefore, if the system recognises a relation "R" at time point t , it applies the reasoning over the next item presented at time point $t + 1$ using "R." In Figure 1, we show an example of a network based on such a program. After the network is built, we have a first step where information about objects "A" and "B" would activate an output neuron representing the relation between them (Recognise_R). As for the second step, we would have information about "C" as input, together with the relation obtained at the previous time point (e.g., propagated to a context unit [Apply_R] through a recurrent link as done in

Figure 1 (Borges et al.). Recurrent network for analogical reasoning.

Elman [1990]). With "C" and "R," the network is capable of inferring "D" according to the stored knowledge.

A simple example like the one above shows that, with a slight change in perspective, a neural-symbolic system can perform the same analogical reasoning as proposed in the target article. And the benefits of this are to allow for the use of symbolic background knowledge, which is important to model cognitive tasks such as language, and to integrate explicitly robust learning and relational reasoning abilities in the same system, as part of what we consider a more appropriate (or complete) modelling of analogy.

Finally, we corroborate the target article's idea that analogy should be seen as an umbrella to different aspects of cognition, serving also as a way of dealing with the problem of absence of explicitly represented knowledge, even in symbolic cases. Cases like the use of language illustrate the array of possibilities in the development of models of cognitive behaviour. Sound deductive inference and inference by induction, analogy, and even discovery have a role to play in the new logic landscape. This constitutes a challenge for different research areas and in particular computer science, as suggested by Valiant (2003), according to whom the modelling and integration of different cognitive abilities, such as reasoning and learning, is a great challenge for computer science in this century.

ACKNOWLEDGMENTS

Rafael Borges is funded by City University London, UK. Luis Lamb is partly supported by CNPq, Brazil.

The evolution of priming in cognitive competencies: To what extent is analogical reasoning adaptive?

doi:10.1017/S0140525X08004494

Paul Bouissac

Department of French Linguistics, University of Toronto, Victoria College 205, Toronto, Ontario, M5S 1K7, Canada.

paul.bouissac@utoronto.ca

<http://www.semioticon.com/people/bouissac.htm>

Abstract: This commentary questions the general assumptions concerning the cognitive value of analogical reasoning on which the argument developed by Leech et al. appears to rest. In order to better assess the

findings of their meta-analysis, it shifts the perspective from development to evolution, and frames their “how” concern within a broader “why” issue.

This commentary does not bear upon the technicalities of the empirical evidence marshaled by Leech et al. to support their developmental claims concerning analogical reasoning. Nor does it address the validity of the connectionist simulations which are presented as plausible neurological explanations for this cognitive competency, which is a defining feature of human psychology. It concerns instead the general assumptions regarding analogical reasoning on which the authors’ whole argument implicitly rests.

First, the pervasive notion that analogical reasoning can be construed as normative can be questioned. The authors constantly refer to the “correct” analogy in their analysis of the experimental results of the tests based upon the formula: *a is to b as c is to what?* It is noticeable that the “correct” answers that are expected are all ultimately based on lexical units and their semantic values even when the tests are apparently purely visual (e.g., a whole pie and a slice of pie). The most convincing part of the general argument is that the “correct” answers rely on the general knowledge available to the children tested. This also holds for the adults. But general knowledge is necessarily “cultural knowledge.” General or “common sense” knowledge is largely built on stereotyped analogies, which are a part of the cultural endowment imparted to children during their earliest languaging and socializing process. It is in this sense normative and culture-dependent. It is socially adaptive, and it does not have any heuristic value.

But analogical thinking is also the source of innovative thinking when it is applied outside the confines of the cultural box. There is anecdotal evidence that children make “wild” analogies. Analogical thinking can be highly innovative and lead to counterintuitive scientific hypotheses or, alternatively, to unsettling artistic visions as in the creative methods promoted by the Surrealist literary movement of the 1920s: what they called the “image” was the perception of a relation – semantic or morphological – between cognitively “distant” realities (Breton 1924/1969, pp. 20–21). In those cases, analogical thinking is unpredictable and has the power to shatter commonsense knowledge. Do infants (and children) start with such “wild” analogical processing of information, and do they need to be “educated” according to their cultural analogical norms? Training by the caretakers is after all very similar to the tests that are described in the target article (e.g., the kitten is the puppy of the cat), but saying that the apple is the kitten of the apple tree probably would not be accepted as a valid answer, although it would be biologically insightful. Then, what about rocks being analogically construed as the puppies of the mountain?

Analogical thinking may be a defining feature of human cognition, but this evolved mental tool to infer knowledge is definitively double-edged. It can equally be the source of scientific discoveries or the perpetuation of erroneous beliefs. A large part of counterintuitive scientific knowledge was built upon the overcoming of entrenched analogical thinking that sustained intuitive evidence, notably in physics and medicine. This is why it is dangerous to uncritically mix analogical and objective truth as the authors do. They overlook the fact that their implicit notion of “normative” analogical thinking is culture-dependent, and not necessarily adaptive in the evolutionary sense of the term. Their examples show that their reasoning is based on positivistic “scientific” (commonsense) knowledge with the corresponding semantic map that is conveyed to infants and children with gesture and language. Other worldviews in other cultures do inculcate different analogies. And some of these analogies can determine, more often than not, maladaptive behavior.

Without raising the thorny issues of the *evo-devo* debate, it might be productive to evoke the evolutionary significance of analogical reasoning. It seems that the human brain is particularly adept, for better or worse, at relying on analogical reasoning. Inferring knowledge from incomplete information (and acting

upon it) is undoubtedly adaptive in most (or at least some) cases. There seems to be good empirical evidence that all perceptions are based on priming and anticipation (e.g., Glimcher 2003; Gold & Shadlen 2007). But, obviously, like the much-celebrated “theory of mind” (TOM) that evolved as a highly adaptive feature of the human genetic endowment (e.g., Penn et al. 2008), it can extend beyond its adaptive range. If TOM prompts populations to attribute mental states to apparently animated objects such as volcanoes, meteorological events, or virtual entities, and if analogical reasoning construes these objects as bloodthirsty predators that have to be satiated or pacified by sacrificing precious resources, obviously TOM (and analogical thinking upon which TOM ultimately rests) is not unambiguously adaptive.

Finally, it should be pointed out that the two main examples of analogical thinking put forward by Leech et al. are particularly infelicitous – and probably prove my point. Niels Bohr’s popular suggestion in 1913 that atoms were to be thought of as miniature planetary systems proved to be seriously misleading and was soon repudiated by Bohr himself, who recanted as early as 1919 (von Baeyer 2003, p. 63). No serious progress could have been achieved in physics if this analogy had become entrenched. As to the Gulf War, the analogy with World War II explored at length by the authors cannot stand political or historical scrutiny. This analogy was created as a political argument developed by the American administration of the time to convince their European allies to join the fight. The image was indeed powerful through the foregrounding of superficial similarities, but very shallow if closely examined. Its purpose was to achieve an emotional impact on the media. Analogies can provide compelling rhetorical arguments through a drastic simplification of situations and problems but are a very shaky ground for constructing scientific or historical knowledge. At most, they can be heuristic. Priming is an effective shortcut to decision making that is not risk-free and can be very costly. Not all shortcuts lead to survival (e.g., Chittka & Osorio 2007).

Analogy as relational priming: The challenge of self-reflection

doi:10.1017/S0140525X08004500

Andrea Cheshire, Linden J. Ball, and Charlie N. Lewis

Department of Psychology, Lancaster University, Lancaster LA1 4YF, United Kingdom.

a.cheshire@lancaster.ac.uk

l.ball@lancaster.ac.uk

c.lewis@lancaster.ac.uk

<http://www.psych.lancs.ac.uk/people/AndreaCheshire.html>

<http://www.psych.lancs.ac.uk/people/LindenBall.html>

<http://www.psych.lancs.ac.uk/people/CharlieLewis.html>

Abstract: Despite its strengths, Leech et al.’s model fails to address the important benefits that derive from self-explanation and task feedback in analogical reasoning development. These components encourage explicit, self-reflective processes that do not necessarily link to knowledge accretion. We wonder, therefore, what mechanisms can be included within a connectionist framework to model self-reflective involvement and its beneficial consequences.

We commend Leech et al.’s attempt to draw together diverse and apparently contradictory accounts of analogy within a single connectionist model. From our standpoint, the particular strengths of this model are that it makes developmental processes paramount, embeds higher level processes in lower level ones, and captures the way in which analogising can take place through sub-symbolic relational priming. Yet their model neglects the possibility that the transfer from lower- to higher-level cognition might occur through explicit, self-reflective processes. To the extent that connectionist approaches fail to embody such

processes, then, Leech et al.'s model needs to be enhanced to accommodate mechanisms that can interleave self-reflective thinking with activation-based priming and pattern completion. In particular, evidence from microgenetic studies – short term investigations designed to chart or manipulate the process of developmental change – identifies two factors associated with explicit, self-reflective thought that are important to any comprehensive model of the development of children's reasoning skills: the child's *explanation* of strategy use and *feedback* provided on task performance. These factors, which derive from the child's interactions with the broader task and social environments, are not given serious consideration in Leech et al.'s model.

Asking children to provide retrospective explanations of their answers during analogical problem solving tasks has been shown to aid the development of their analogical reasoning skills in several recent studies (e.g., Cheshire et al. 2005; Siegler & Svetina 2002). Leech et al. simulated how the provision of verbal labels for relational terms benefits subsequent analogical reasoning behaviour, but there are key differences between providing such verbal cues and asking children to explain why they chose a particular answer. Verbal labels may augment the representation of spatial or transitive relations and thus facilitate performance, but self-explanation prompts children to be reflective in order to verbalise explicitly why they have selected a response. What seems to be needed is an account of how the explicit process of self-explanation can impact beneficially on analogical reasoning. It may be, for example, that self-explanation somehow increases the likelihood of an appropriate relational item being selected in the future. Such an account, however, is complicated by the finding that improvements in relational mapping that derive from self-explanation are *transitory*, with children reverting to a preference for superficial object similarity when no longer asked to explain (Cheshire et al. 2005). The inherent fragility of benefits deriving from self-explanation seems especially difficult to model via activation-based priming and pattern completion.

Children also benefit from feedback during analogical reasoning and other cognitive tasks (e.g., Cheshire et al. 2005; Muldoon et al. 2007; Siegler 2006; Siegler & Svetina 2002). Again, however, such interaction with the social environment associated with learning is not considered in Leech et al.'s proposals. Our microgenetic data indicate that there is a rapid phase of learning when explicit feedback is given, with such feedback promoting long term improvements in analogical reasoning. A connectionist model could no doubt simulate the rapidity and permanence of such feedback-based learning. But exactly how positive and negative feedback strengthens the connections between the nodes in the network is a complex issue and one that seems unlikely to be addressed simply through an appeal to relational priming and pattern completion mechanisms.

Furthermore, our microgenetic data suggest different trajectories of change for children in different conditions dependent upon the presence or absence of feedback and explanation. Figure 1 compares the effects over seven sessions of repeated testing without feedback or explanation (i.e., simple priming), providing feedback, asking children to explain their decision, and a combination of the two (see Cheshire et al. 2007). Asking children to provide explanations enhances performance in terms of a preference for relational mappings over superficial object matches, but, as noted above, this improvement is transitory, and children subsequently revert to a preference for object similarity when no longer asked to explain (Cheshire et al. 2005). Although the effects of explanation versus feedback on analogising are qualitatively different to each other (i.e., explanation leads to more transitory effects; feedback produces more enduring changes), the combination of the two leads to greater accuracy and more permanent learning. The varied developmental pathways depicted in Figure 1 attest to the subtle and complex influences of self-reflective processes on the development of analogical reasoning. We contend that such subtleties and complexities need to be matched by equally sophisticated mechanisms

Figure 1 (Cheshire et al.). Different paths of change in a microgenetic study of the development of analogical reasoning (Cheshire et al. 2007). Sessions involved 22 matrix-completion trials (essentially involving analogies of the a:b::c:d form) whose solutions entail relational mapping processes. Sessions 1, 6, and 7 were not associated with any experimental manipulation. Session 7 occurred approximately 8 weeks after Session 6.

within a connectionist framework of the type proposed by Leech et al.

In conclusion, Leech et al.'s model falls foul of a long-standing problem that in concentrating on the development of internal structures such models are solipsistic (Frawley 1997). The development of analogical problem solving is not just due to repeated, passive exposure to a problem and general knowledge accretion. Through their active commitment to learning children partly train their own networks by strengthening different connections based on their individual experiences (which may include explanation and feedback). Can the interplay between implicit networks and explicit, self-reflective thought be modelled in a detailed, psychologically plausible, and testable manner?

Analogy is priming, but relations are not transformations

doi:10.1017/S0140525X08004512

Fintan J. Costello

School of Computer Science and Informatics, University College Dublin, Belfield, Dublin 4, Ireland.

fintan.costello@ucd.ie

<http://www.cs.ucd.ie/Staff/AcademicStaff/fcostello/>

Abstract: Leech et al. make two proposals: that relational priming is central to analogy, and that relations between objects are best represented as transformations of those objects. Although their account of analogy as relational priming is a useful contribution to our understanding of analogical development, in this commentary I show that relations in general cannot be represented by transformations.

In the “additive transformation” model of relations proposed by Leech et al., a relation R holding between two objects, *Object 1* and *Object 2*, is represented by a transformation vector T , which, when added to the feature vector representing *Object 1*, produces the feature vector representing *Object 2*. The transformation vector T defines the relation R : if any pair of objects are related by relation R , that means the first object can be transformed into the second object by adding the vector T ; conversely, if one object can be transformed into another object by adding the vector T , that means those two objects can stand in relation R to each other. (The zero vector has a special meaning in this model: it represents the auto-association relation, which simply transforms an object into itself.) Leech et al. use this transformational representation of relations to simulate and test a relational priming account of analogy. These simulations involve training a

connectionist network to learn relations encoded as transformations. To train the network on relation R the network is presented with pairs of objects in which the first object's feature vector differs from the second by the transformation vector T . These object pairs are presented in conjunction with a feature vector that labels the relation in some way. Over the course of training, the network learns to associate the relation label with the transformation required to get from one object in a pair to the other; in other words, the network learns the transformation. Leech et al. show that networks trained in this way can learn a number of different relations and can, by allowing objects to prime relations, carry out simulations of something very like analogical reasoning. These simulations produce results that agree with those seen in experimental studies of analogy, and thus support a relational priming model of analogy.

Although this additive transformation model is plausible for causal relations such as the *cutting* relation used by Leech et al., in general relations cannot be represented by transformations. The problem is that many relations are bidirectional: they can hold both between *Object 1* and *Object 2* and also between *Object 2* and *Object 1*. The relation *threatens*, for example (which Leech et al. also use in their simulations), is bidirectional: it can hold both between *Iraq* and the *United States* (i.e., Iraq threatens the United States) and between the *United States* and *Iraq* (i.e., the United States threatens Iraq). Both directions of a bidirectional relation, however, cannot be represented via the same additive transformation, and so such relations cannot be represented in an additive transformation model of relations.

To see why bidirectional relations cannot be represented in the additive transformation model proposed by Leech et al., assume that relation R is bidirectional and that R holds both between *Object 1* and *Object 2* and between *Object 2* and *Object 1*. Let $[Object\ 1]$ represent the feature vector describing *Object 1*, and $[Object\ 2]$ denote that describing *Object 2*. By definition, for the relation R to hold in both directions, the two expressions

$$(1) [Object2] = [Object1] + T$$

$$(2) [Object1] = [Object2] + T$$

must both be true (where T is the transformational vector representing relation R , as before). However, by substituting Equation 1 into Equation 2 we can see that this means

$$(3) [Object1] = [Object1] + T + T$$

must be true. This can only be true when

$$(4) T = 0.$$

Thus the bidirectional relation between two objects can only be represented via transformation if the transformational vector T applied to convert one object to the other is simply the zero vector. But the zero vector only "converts" one object into itself (auto-association, in Leech et al.'s terms). The additive transformation model is thus unable to represent bidirectional relations between two different objects. This means that if a transformational system of relations is set up to represent "Iraq threatens the United States," that system is necessarily unable to represent "the United States threatens Iraq."

Leech et al.'s additive transformation model of relations, then, can only apply to relations which are not bidirectional: relations which, if they hold between *Object 1* and *Object 2*, cannot hold between *Object 2* and *Object 1*. The causal relations used by Leech et al. in their first simulation are clearly not bidirectional because they involve irreversible physical change to an object: the *cutting* relation transforms an *apple* into a *cut apple*, but cannot transform a *cut apple* back into an *apple*. Leech et al. have thus provided a relational priming model of analogy which works for certain types of relations (relations that are not bidirectional), but not for relations in general.

This difficulty with bidirectional relations does not take away from the success of Leech et al.'s account of analogy as relational

priming. Their account applies well to an important, and perhaps a fundamental, type of relation (simple causal relations), and gives a good account for the developmental results on analogies using relations of that type (Goswami & Brown 1989; Goswami et al. 1998). It may be possible to extend this account by using a different transformational model of relations rather than the additive model used by Leech et al. – for example, a transformational model based on modular arithmetic may be able to successfully represent bidirectional relations. It is hard to see how modular arithmetic could be implemented in a connectionist network, however. It may be better to accept that the transformational model of relations will only work well for relations that actually involve transformations: for relations that are not transformations, a different model will be needed.

ACKNOWLEDGMENT

The author was supported by the FP6 NEST Programme of the European Commission (ANALOGY: Humans – the Analogy-Making Species: STREP Contr. No 029088).

Toward extending the relational priming model: Six questions

doi:10.1017/S0140525X08004524

Eric Dietrich

Philosophy Department, Binghamton University, Binghamton, NY 13902-6000.

dietrich@binghamton.edu

http://bingweb.binghamton.edu/~dietrich

Abstract: Six questions are posed that are really specific versions of this question: How can Leech et al.'s system be extended to handle adult-level analogies that frequently combine concepts from semantically distant domains sharing few relational labels and that involve the production of abstractions? It is Leech et al. who stress development; finding such an extension would seem to have to be high on their priority list.

I begin with my first question: Why is *apple:cut apple :: bread:cut bread* an analogy and not a simple categorization? After all, they are both from the food domain. If I show a child a paradigmatic dog – for example, a beagle – and then show that child an unusual dog – for example, a hairless Chinese Crested – then if the child correctly identifies the Chinese Crested as a dog, we would say that the child has *categorized* the Crested, not that the child has made an analogy. Presumably then, what makes an analogy is mapping *relational structure*. (I am happy to grant this, but it is a gift, since, for all we know, most concepts have relational structure. So, appealing to such mapping, even mapping *systems* of structures, may not actually distinguish categorization from analogy.) But is any relational categorization an analogy? Perhaps relational mapping isn't all there is to analogy. Though the notion is hard to define experimentally or even theoretically, the "semantical distance" between the two analogue domains is also important: the greater the distance, the deeper the analogy (Leech et al. make some interesting comments related to this in sect. 5.1.2). So, now my second question: Could Leech et al.'s connectionist network make an analogy (with appropriate changes) between *apple:cut apple* and *team:cut from the team*, or *cut of a deck of cards*? Or, since the result in their examples of cutting an apple and a loaf of bread is a slice, could their network be trained in such a way as to make an analogy with *slice of life*?

Admittedly, analogies between cutting an apple and being cut from a team are more complex, "adult-level" analogies, and Leech et al. admit that their model has a ways to go before it can make adult analogies; yet, though they do an admirable job of arguing that their model can be extended to make such analogies, one still wonders whether their three central notions of pattern completion, relational priming, and implicit representations are

up to the task. This leads to my third question: Can their implicit relational representations account for both the abstraction and conscious knowledge seen in adult analogies? This third question is really comprised of three, more specific, questions, to which I now turn.

Two crucial aspects of adult analogies are that they create abstractions on the fly (the relational structures of the two analogues are abstracted; this is sometimes called “relational change”; see Dietrich et al. 2003) and they bring to consciousness the fact that the two analogues are indeed analogous (one can always say, “X is like Y” or “X reminds me of Y”; see Dietrich 2000; and cf. Leech et al., sect. 5.2, first para.). Rutherford’s crucial analogy between the deflected trajectories of alpha particles and the orbital paths of comets is a classic example of the sort of abstraction that I mean (the abstraction is due in part to the fact that deflected trajectories are by no means orbits – Rutherford 1911; also see “Rutherford Scattering” and “Gold Foil Experiment” in Wikipedia). This analogy (not the data behind it) doomed the reigning plum-pudding model of the atom and led eventually to the better Bohr model of the atom. Science is rife with such examples.

It is hard to see how these two aspects of analogy could be realized in Leech et al.’s model. How can a relational structure be abstracted if it’s only implicitly represented as a state transformation? That would seem to require abstracting the states. Is that possible on their model? And how can the knowledge of something be brought to consciousness if it is not explicitly represented?

Models of analogy that use explicit, discrete representations seem to be able to handle these two aspects. Structure-mapping theory (Gentner 1983) is a good example. However, structure-mapping theory is not the final word on analogy, as everyone, including Gentner, admits (though, of course, there is disagreement on *why* it is not the final word). It is clear that Leech et al.’s developmental model of analogy brings important new insights to the table. What is not clear is how to have one model that does justice to the developmental data, the obvious constraint that analogy would seem to have to arise from simpler mechanisms such as priming, and the fact that analogies produce abstractions to which we have conscious access (as we do to the analogies themselves). This problem of different, incompatible models explaining different aspects of a single cognitive phenomenon is quite common in cognitive science. This takes me to my final question.

In their Gulf War/World War II model, Leech et al. use the exact same relations in each domain (see Table 3 in sect. 4.1.2). This seems to weaken their case that their relational priming model could be extended to adult-level analogies, for it is very unlikely that the exact same relations would obtain in different domains. It is even very unlikely that the same labels (names) would obtain in different domains. This problem is related to a problem within structure-mapping theory (SMT). According to SMT, analogies are isomorphisms between high-level structures. In an important sense, the analogous concepts simply share one structure that funds their being analogous. Furthermore, SMT assumes that these isomorphisms obtain *before* a given analogy is made (indeed, the isomorphisms explain *why* the analogy was made). The probability that such isomorphisms obtain before an analogy is made is very low – too low to account for the quantity of analogies an individual produces (Dietrich 2000). Therefore, on any SMT-like model, the relevant representations in an analogy have to change at the time of retrieval in order to forge the needed isomorphism (Dietrich et al. 2003). Perhaps relational priming could solve this problem without invoking retrieval-based change.

It seems as though the best extension of Leech et al.’s system for handling adult-level analogies would be a system combining their insights regarding relational priming, a process for rendering explicit some implicit relational representations, together with a process for abstracting those explicit representations so that

semantically distant analogies can be made. Is such an extension compatible with how Leech et al. see their future work unfolding?

Developing structured representations

doi:10.1017/S0140525X08004536

Leonidas A. A. Doumas^a and Lindsey E. Richland^b

^aDepartment of Psychological and Brain Sciences, Indiana University, Bloomington, IN 47405; ^bDepartment of Education, University of California, Irvine, Irvine, CA 92697-5500.

adoumas@indiana.edu

lerich@uci.edu

<http://www.gse.uci.edu/richland>

Abstract: Leech et al.’s model proposes representing relations as primed transformations rather than as structured representations (explicit representations of relations and their roles dynamically bound to fillers). However, this renders the model unable to explain several developmental trends (including relational integration and all changes not attributable to growth in relational knowledge). We suggest looking to an alternative computational model that learns structured representations from examples.

Leech et al. propose relational priming as the mechanism for the development of analogy making in children, and they provide simulation results that follow several known trends observed in children’s early analogical reasoning. However, there are limitations in the present model that undercut its feasibility as a general explanation for the development of analogy. First, the authors claim that rather than explicitly representing relations as structured representations, it may be sufficient to conceptualize relations as transformations on objects. A consequence of this approach is that it appears to leave the model unable to integrate multiple relations when making analogies. For example, it is unclear how the model could sort out the mappings between *kicks* (x, y) and *kicks* (y, z) and *hits* (b, c) and *hits* (a, b). Although the authors acknowledge that their current model does not explain all aspects of relationally complex adult analogical reasoning, in fact the ability to represent integrated relations is crucial to a developmental model as well. By 5–6 years of age, children are well above chance in tasks requiring integration of multiple relations (Andrews & Halford 2002; Holyoak et al. 1984; Richland et al. 2006), and even children ages 3–4 show some ability to integrate two relations (e.g., Rattermann & Gentner 1998a; Richland et al. 2006).

Second, all developmental changes in the model derive from knowledge accretion, but knowledge has been shown unable to predict all behavioral patterns of development (see Goswami et al. 1998; Richland et al. 2006). In Richland et al. (2006), children were asked to map analogical correspondences between sets of two pictures in which the same relation was represented (e.g., picture A showed dog chases cat chases mouse; picture B showed mom chases boy chases girl). A test of the stimuli showed that children ages 3–4 had sufficient relational knowledge to identify the primary relation at 90% accuracy across all the pictures, indicating that adequate relational knowledge could not be the explanation for condition or age related differences in performance. Although 3–4 year olds could solve simple one-relation analogies between the scenes (e.g., dog chases cat; boy chases girl), the same participants (with the same relational knowledge) showed significantly lower relational mapping accuracy when there was an additional processing load. This was either in the form of an object similarity distracter in the target scene, or when the mapping task required a higher level of relational complexity. All pictures were counterbalanced across conditions. With age, children improved in these conditions in timing that coincides with maturation of working memory and inhibitory control systems. These data indicate that while acquisition of adequate

relational knowledge is likely a prerequisite to relational reasoning, it is not sufficient to explain established developmental patterns.

The authors motivate the approach of using priming rather than explicit relational representations because no current models account for how explicit relational representations can be learned. This is, however, an oversight. The Discovery of Relations by Analogy model or DORA (Dumas & Hummel 2005a; 2007; Dumas et al. 2008) is a connectionist model that learns explicit representations of relations from unstructured examples. Beginning with simple distributed representations of objects as feature vectors, DORA uses comparison to bootstrap learning explicit representations of object properties and relational roles. DORA then uses mapping to link sets of relational roles into complete multi-place relations. Importantly, DORA uses time to dynamically bind these representations of object properties, roles, and relations to arguments (see, e.g., Dumas et al. 2008; Hummel & Holyoak 1997; 2003). In other words, DORA learns explicit representations of relations that explicitly specify bindings between relational roles and fillers. The resulting representations allow DORA to account for numerous phenomena from children's and adults' relational reasoning including all those accounted for by the target model, as well as numerous phenomena the present model cannot account for (such as relational integration; see Dumas et al. 2008).

Relational processing in conceptual combination and analogy

doi:10.1017/S0140525X08004548

Zachary Estes^a and Lara L. Jones^b

^aDepartment of Psychology, University of Warwick, Coventry CV4 7AL, United Kingdom; ^bDepartment of Psychology, Francis Marion University, Florence, SC 29501.

Z.Estes@warwick.ac.uk

<http://www2.warwick.ac.uk/fac/sci/psych/people/academic/zestes>

ljones@fmarion.edu

<http://acsweb.fmarion.edu/ljones/>

Abstract: We evaluate whether evidence from conceptual combination supports the relational priming model of analogy. Representing relations implicitly as patterns of activation distributed across the semantic network provides a natural and parsimonious explanation of several key phenomena observed in conceptual combination. Although an additional mechanism for role resolution may be required, relational priming offers a promising approach to analogy.

Leech et al. propose that analogies are understood by relational priming. For instance, PUPPY:DOG::KITTEN:??? is completed by applying the relation between the base concepts (a PUPPY is the *offspring* of a DOG) to the target concepts (a KITTEN is the *offspring* of a CAT). In addition to its occurrence in analogy, relational priming also occurs regularly in common language use. Indeed, object concepts are combined frequently in language (e.g., BIRD NEST), and like analogies, such conceptual combinations are understood by retrieving or inferring some relation (e.g., *habitation*) between the given concepts. In conceptual combination, relational priming occurs when one phrase (e.g., BIRD NEST) facilitates comprehension of a subsequent phrase (e.g., FISH POND) that instantiates the same relation (Estes 2003; Estes & Jones 2006; Gagné 2001; Spellman et al. 2001). Given this fundamental similarity between analogy and conceptual combination, then, research on conceptual combination may serve as a useful tool for evaluating models of analogy.

Leech et al. posit that relations are represented as transformations between activation states. More specifically, a relation is represented as the pattern of activation required to transform

an input object (e.g., APPLE) into an output object (e.g., CUT APPLE). In terms of conceptual combination, this corresponds to a simple concept (e.g., NEST) being transformed into a compound concept (e.g., BIRD NEST). Because such transformations are carried out within the hidden layer of the model, a relation is represented implicitly as a pattern of activation within the semantic network.

This transformational model of relational representation naturally explains several of the key observations in research on conceptual combination. First, familiar combinations (e.g., BIRD NEST) and novel combinations (e.g., TURTLE CAGE) are understood via the same processes. Intuitively, it seems that familiar combinations would be understood by simply retrieving the compound concept from memory, whereas novel combinations necessitate a relational inference. However, the evidence suggests that familiar and novel combinations undergo the same computations (e.g., Gagné & Spalding 2004). The transformational model provides a straightforward explanation for this otherwise counterintuitive observation: The relational inference entails a transformation from simple concepts to a compound concept, regardless of the familiarity of the compound. Although the relational transformation may proceed more quickly for familiar compounds than for novel compounds, it nevertheless must occur in both cases.

Second, relational representations are independent of the concepts that instantiate them. If relational representations were concept bound, then relational priming should only occur when the base and target exhibit lexical repetition (e.g., BIRD CAGE → BIRD NEST; Gagné 2001). In actuality, however, relational priming also occurs in the absence of lexical repetition (Estes 2003; Estes & Jones 2006; Raffray et al. 2007; Spellman et al. 2001). For example, FISH POND facilitates the comprehension of BIRD NEST because both combinations utilize the same relational representation. Because the transformational model posits that relations are represented as unique patterns of activation that may be triggered by multiple input objects, the model clearly predicts relational priming without lexical repetition (otherwise, it couldn't possibly explain analogy). The independence of relational representations is further evidenced by the facilitative effect of relational labels on analogy completion. Although a relational label does not add new information to the network, it effectively cues the relational transformation, regardless of the concepts that instantiate it.

Finally, relational representations are somewhat specific. To illustrate, BIRD NEST, TURTLE CAGE, and COOKIE JAR all nominally instantiate a general *location* relation. So if relational representations were this general, then TURTLE CAGE and COOKIE JAR should both facilitate comprehension of BIRD NEST. But in actuality, TURTLE CAGE facilitates comprehension of BIRD NEST, but COOKIE JAR does not. This selectivity of relational priming indicates that relational representations are specific, more like *habitation* (i.e., TURTLE CAGE and BIRD NEST) and *containment* (i.e., COOKIE JAR). The transformational model explains this selectivity of priming as a consequence of relational similarity. That is, the pattern of activation required to transform NEST into BIRD NEST is highly similar to that required to transform CAGE into TURTLE CAGE, but is relatively dissimilar to the transformation from JAR to COOKIE JAR. Without sufficient relational similarity, relational priming cannot occur (Estes & Jones 2006).

An important issue that may pose a challenge for the transformational model is role resolution. That is, once a relation between concepts is inferred, those concepts must also be assigned to appropriate roles in that relation. Otherwise, the relational inference would lead to frequent errors in interpretation (Hummel & Holyoak 2003). Consider the causal relation, for which it is crucial to distinguish cause from effect (see Fenker et al. 2005). WIND EROSION and GROWTH HORMONE both instantiate the causal relation, but note that the ordering of cause and effect is reversed in the two combinations.

Because the transformational model has only implicit relational representations, with no explicit provision for role resolution, it is unclear how the model will account for such differences. To demonstrate the issue with analogy, consider WIND:EROSION::SMOKE:???. Once the base pair activates the causal transformation, the target pair will tend to undergo the same transformation. But on what basis will the model correctly produce an effect of SMOKE (e.g., COUGH) rather than a cause (e.g., FIRE)? A simple solution is to stipulate that each relation has two distinct transformations, one for each possible ordering of role assignments (e.g., cause → effect and effect ← cause). However, the cost of this relational proliferation may essentially offset the benefit of representing relations implicitly. Thus, we view role resolution as an important issue requiring explicit elaboration in the model.

In summary, the transformational model parsimoniously explains several key phenomena of conceptual combination. Although important issues remain to be addressed, we believe that Leech et al. have provided a promising framework for modeling analogy and other relational processes, such as conceptual combination.

Relational priming is to analogy-making as one-ball juggling is to seven-ball juggling

doi:10.1017/S0140525X0800455X

Robert M. French

LEAD-CNRS UMR 5022, University of Burgundy, 21065 Dijon, Cedex, France.

robert.french@u-bourgogne.fr

<http://www.u-bourgogne.fr/lead/people/rfrench.html>

Abstract: Relational priming is argued to be a deeply inadequate model of analogy-making because of its intrinsic inability to do analogies where the base and target domains share no common attributes and the mapped relations are different. Leech et al. rely on carefully handcrafted representations to allow their model to make a complex analogy, seemingly unaware of the debate on this issue fifteen years ago. Finally, they incorrectly assume the existence of fixed, context-independent relations between objects.

Although relational priming may, indeed, play some role in analogy-making, it is an enormous – and unjustified – stretch to say that it is “centrally implicated in analogical reasoning” (sect. 2, para. 2). This is a very strong statement, and the authors have not shown at all that this mechanism – as opposed to, say, explicit structure-mapping (Gentner 1983) or slippage (French 1995; Mitchell 1993; Mitchell & Hofstadter 1990) or constraint satisfaction (Holyoak & Thagard 1989) or dynamic binding (Hummel & Holyoak 1997), to name a few – is at the heart of high-level analogy-making.

The proportional analogy that dominates almost three-quarters of the target article, both in the theoretical discussion and the simulations, is:

BREAD is to CUT BREAD as APPLE is to ?

The possible answers are:

CUT APPLE, APPLE, BRUISED APPLE, and CUT BANANA.

A simple connectionist attractor network produces the correct answer (CUT APPLE). The authors claim that “relational priming” – supposedly what the network is doing – is the key to analogy-making and that this mechanism can be extended to full-blown analogy-making. Unfortunately, both claims are dubious, at best.

Let us start by considering a run-of-the-mill analogy that occurred to me recently. Ty Cobb, one of baseball’s greatest players ever, did not hit over-the-wall home runs like his arch-rival, Babe Ruth. However, in May 1925, after being egged on

by reporters, he said, “You want me to hit home runs like Babe Ruth? No problem.” So, in the next game he hit three over-the-wall home runs and two more in the following game, after which he said, “I told you I could hit home runs. Now I’m going back to playing baseball the way it was meant to be played.” It occurred to me that this is analogous to a professor who publishes only books and who is criticized for never publishing in journals. One day, in response to his critics, he says, “I could publish in journals if I wanted to, I simply chose not to,” and, to prove his point, he rapidly racks up a number of publications in the top journals in his field. Thereafter, he returns to writing books.

Here we have two situations in which the objects have no features in common (*base hits* and *books*; *home runs* and *journal articles*) and where the relations are semantically miles apart (*hitting* and *writing*). And yet, the analogy works perfectly. This is the heart of analogy-making, and it is not in the least clear how relational priming, as implemented by the authors, could begin to deal with this problem.

Late in their article, the authors finally come to grips with the necessity of explaining “how such a mechanism [relational priming] might explain the more difficult and complex analogies used to test adults” (sect. 3.4, para. 10). They then discuss how this would occur by considering an analogy between the 1991 Gulf War and World War II (Spellman & Holyoak 1992). But where did the relations and objects on which they train their network come from? This is a textbook case of the Problem of Representation (Chalmers et al. 1992), a problem that, incredibly, they never mention, and one that, arguably, was the greatest problem of traditional artificial intelligence (AI). Once you have handcrafted representations for each situation with a limited number of relations, finding mappings between them is, relatively speaking, a piece of cake.

This careful handcrafting of objects and relations in the two situations into exactly the representations that are needed for an analogy to be found is precisely what certain members of the analogy-making community (Chalmers et al. 1992) have been railing against for years. To be completely clear: There are literally *millions* of relations and objects that could be used to describe the 1991 Gulf War and just as many that characterize World War II. But for their example, Leech et al. have selected *only* those relations that make their analogy work. This is wholly unacceptable as a way forward in analogy-making and, what’s more, tells us essentially nothing about how real analogy-making works, since *finding* the relations is part-and-parcel of the process. It cannot be separated out.

Finally, the authors write as if the relation between two objects exists in a context-independent manner and can thus be primed by the presence of the objects themselves. They write, for example, “we propose that exposure to the *a* (e.g., *puppy*) and *b* (e.g., *dog*) terms of an analogy primes a semantic relation (e.g., *offspring*), which then biases the *c* term (*kitten*) to produce the appropriate *d* term (*cat*)” (sect. 2.1, para. 1). But this cannot be right. Consider the following example: *puppy:dog::watch:?* Most people would say: *clock*. But is a watch an offspring of a clock? Of course not. The point is that the word *watch* helps determine the relationship between *puppy* and *dog*. There is no a priori intrinsic relation between *puppy:dog* that can be used for all analogies, as the Leech et al. model needs to assume. In the latter analogy, the germane relationship was *bigger than* and certainly not *offspring of*. The authors fail to understand this absolutely crucial point about the context-dependence of analogy-making. This seems to be a deep, and, in my opinion, irreparable, flaw in their model. This point is not one that can be simply glossed over or easily patched up. The whole manner in which their model is trained up requires there to be an a priori relationship between *a* and *b*, which is then transferred to *c*. But this is deeply wrong. Analogy-making, as the above example clearly shows, doesn’t work that way. And one can come up with examples like this all day. For a detailed discussion of this point, see Chalmers et al. (1992).

Conclusion. The ideas presented by the authors in their model are not central to the key principles underlying the mechanisms of analogy-making. Their model is capable of limited relational learning, that is all; something that was done by earlier connectionist models (e.g., Chalmers 1990) almost twenty years ago. In short, as a model of analogy-making, the present model is woefully inadequate. There is much, much more to analogy-making than relational priming.

ACKNOWLEDGMENT

This work was supported in part by a grant FP6-NEST-029088 from the European Commission.

Analogy and the brain: A new perspective on relational primacy

doi:10.1017/S0140525X08004561

Usha Goswami

Centre for Neuroscience in Education, Faculty of Education, Cambridge, CB2 8PQ, United Kingdom.

ucg10@cam.ac.uk

<http://www.educ.cam.ac.uk/research/centres/neuroscience/>

Abstract: Leech et al.'s demonstration that analogical reasoning can be an emergent property of low-level incremental learning processes is critical for analogical theory. Along with insights into neural learning based on the salience of dynamic spatio-temporal structure, and the neural priming mechanism of repetition suppression, it establishes relational primacy as a plausible theoretical description of how brains make analogies.

At last, a truly developmental account of a fundamental cognitive "skill," analogy! Leech et al.'s computational demonstration that reasoning by analogy can *in principle* arise from pattern completion by simple memory processes is extremely important for analogical theory. Recent advances in cognitive neuroscience and connectionism are helping us to understand how the brain builds a complex cognitive system on the basis of simple incremental learning mechanisms that begin with sensory input (Goswami 2008). Indeed, aspects of development long considered to require innate "pre-knowledge," such as syntax acquisition, can *in principle* be acquired by networks that learn according to simple statistical algorithms (Freudenthal et al. 2006). In developmental psychology, connectionist demonstrations have been crucial both for dealing with the "poverty of the stimulus" argument (Elman 2005) and for establishing that relational information can be computed from instance-based learning. Leech et al. achieve the same outcomes for analogy. Simple sensory input (apples, cut apples, knives) can be rich in relational information, and instance-based learning can yield complex cognition.

In my view, the question of whether the brain achieves analogy on the basis of relational priming is secondary to the demonstration that analogies can in principle be achieved by low-level, automatic, and incremental learning processes. Indeed, cognitive neuroscience and developmental psychology both provide extensive empirical demonstrations that relations are represented as more than "transformations between items" (sect. 2.2, para. 2). Instead, low-level sensory processing represents spatio-temporal dynamic structure. In adult neuroscience, it can be shown that when cross-modal input has temporal patterning, the brain will abstract these patterns and alter uni-modal sensory processing in terms of the higher-order dependencies in the patterns (e.g., Noesselt et al. 2007). Regarding babies, 3-month-olds who view abstract motion patterns (point light displays) that specify either vehicles or animals distinguish the two inputs on the basis of this dynamic information alone (and do so as well as they distinguish actual pictures of vehicles and animals; Arterberry & Bornstein 2001). To babies, relations specify *as much information* as the objects themselves.

Our sensory processing systems therefore end up prioritising spatio-temporal dynamic structure (abstracted dependencies) over instance-based featural (object) information, and this is also revealed by sensory "illusions" (Riecke et al. 2007). The primacy given to these abstracted dependencies is very important theoretically, as it means that the way in which our brains process sensory structure can *in principle* yield the "relational primacy" that I argued for in 1991 (see target article, sect. 3.4, para. 7). In essence, sensory systems are representing underlying structure (traditionally discussed as "prototypes," "naïve theories," "innate biases," or "core knowledge"; e.g., Rosch 1978; Spelke 1994). Incremental learning by sensory neural networks that represent dynamic relations automatically represents relational structure. These emergent knowledge systems are enriched and transformed as the child acquires language (Vygotsky 1978). In analogy, as in conceptual development, verbal labeling supports structural similarity over perceptual similarity (Gelman & Coley 1990).

The modelling conducted by Leech et al. is thus compelling in establishing two central developmental phenomena: Analogical completion is an emergent property of the way that relational information is represented in a (neural) network; and incremental learning processes can yield developmental effects previously explained by symbolic theories. Do we need the extra assumptions, that relations are transformations, and that consequently a strong test of the relational priming account is whether semantic relational priming is found in young children (target article, sect. 5.2, para. 4)? I doubt it.

Firstly, relational priming as discussed by Leech et al. and as tested in the studies on adults they cite is rather narrow in scope (e.g., "apple" and "cake" priming "made of"). If children didn't show these automatic effects (which also require adequate reading skills, noise-free reaction times [RTs], and relevant domain knowledge), I am not sure it would matter. Secondly, priming in the cognitive neuroscience literature offers a general tool for studying the nature of the code in a given brain region, via repetition suppression (Dehaene et al. 2001). This latter conceptualisation of priming seems more relevant to analogy as an emergent phenomenon, and gives inhibition "for free." Neural repetition suppression techniques have already been used to investigate the coding of relational information (e.g., numerical quantity; see Naccache & Dehaene 2001). Hence, to show that Leech et al.'s model is biologically plausible, all that is required is evidence that young children show neural repetition suppression to relational information. Even simple causal information (e.g., launching events) would suffice for such a test. Priming effects seem unlikely to be isolated to the temporal cortices (sect. 5.4, para. 7), as semantic memory is no longer understood as a distinct symbolic system. Rather, the activation of particular concepts (by adults) produces neural activation in the sensory modalities associated with those concepts and in association areas recording the conjunctions of particular sets of sensory information (e.g., Barsalou et al. 2003). Studying the repetition suppression of such conjunctions appears the most productive way to understand analogy as a form of neural priming.

Finally, what of the "unexpected consequence" of the Leech et al. model (sect. 5.1.1, para. 2), that there is no *necessary* relational shift for any given relation in a child's similarity judgements? This seems highly likely, and a relational primacy account must predict that for some relations, children might show an initial bias for relations over objects. In fact, 3-month-old babies do just this. In the 1980s, Rovee-Collier and her colleagues conducted a series of experiments on infant memory, using a "conjugate reinforcement" paradigm. Three-month-old babies were trained to kick when they saw a distinctive mobile, and were reinforced because their kicking set the mobile in motion (a causal relation). These memories lasted for months, and the most effective retrieval cues were other mobiles, including mobiles that were visually completely different from the distinctive training mobile (e.g., a mobile with one butterfly on a ring as reminder cue, a mobile of hanging dice during relational learning; see Greco et al. 1990). Thus shared

functional relations and not object similarity was the core retrieval cue – consistent with a relational primacy account and with Leech et al.'s model.

Implicit analogy: New direct evidence and a challenge to the theory of memory

doi:10.1017/S0140525X08004573

Anthony J. Greene

Department of Psychology, University of Wisconsin–Milwaukee, Milwaukee, WI 53223.

ag@uwm.edu

<http://www.uwm.edu/~ag>

Abstract: The authors propose that analogical reasoning may be achieved without conscious or explicit deliberation. The argument would be strengthened by more convincingly demonstrating instances of analogy that do not require explicit deliberation. Recent findings demonstrate that deliberative or explicit strategies are not necessary for flexible expression under novel circumstances (Greene et al. 2001) to include analogical transfer (Gross & Greene 2007). This issue is particularly critical because the existence of relational priming poses a serious challenge to the widely held notion that flexible expression of learned relations requires deliberative processes.

Leech et al. make a compelling argument that analogical transfer may not require a deliberative or explicit mapping process. However, they provide no direct evidence showing successful implicit analogy. The analogy tasks reviewed in the target article do not provide evidence which rules out the use of deliberative strategies. In one such task the intent was to show that analogical transfer did not require a mapping of base to target (Ripoll et al. 2003). Although the finding tends to confirm that subjects perform analogical transfer without deliberative mapping, it does not rule out the possibility that analogical transfer was itself accomplished deliberately. Indeed, participants were all explicitly informed that their task was to complete an analogy. In another task, pre-exposure to a problem solution in the base task biased participants to employ a similar strategy in the target task (Schunn & Dunbar 1996). However, the only evidence that participants employed an implicit strategy was that they did not mention the base problem as a foundation for the solution to the target problem. Although these findings are both consistent with the use of implicit strategies, neither directly rules out the possibility that deliberative strategies were used.

Recently, direct evidence showed that analogical transfer occurs in the absence of deliberative strategies (Gross & Greene 2007). Participants learned a transverse patterning sequence and transferred the learned relations to a novel stimulus set. For transverse patterning, given a choice between two faces, A and B, participants learn by trial and error to select A ($A > B$), and likewise they learn $B > C$ and $C > A$. To test for analogical transfer, participants simultaneously learn a partial set of relations among different faces, $X > Y$ and $Y > Z$, and are then tested on the novel pairing $X ? Z$ (the question mark indicated that no choice was ever reinforced). In the control condition, the training on the transverse patterning set is omitted, and participants uniformly infer that $X > Z$. However, with exposure to the transverse patterning set, a significant portion of participants apply the transverse patterning relations and select $Z > X$. This analogical transfer occurred without intent or awareness of the transfer. Two tests of awareness were used. First, a post-experimental questionnaire was employed with increasingly leading questions about the use of analogical transfer. Only those participants who asserted – in a forced-choice question – that no analogical reason existed for the $Z > X$ choice were categorized as unaware. In addition, it was noted during the debriefing that many participants were surprised that they had not been trained on the $Z ? X$ choice. To test this, a

recognition task revealed that most participants did not recognize $Z ? X$ as a novel configuration. Had a deliberative strategy been the foundation for analogical transfer, participants would have had to recognize the novel configuration to apply the mapping.

Analogical transfer is but one instance of relational priming. Evidence that complex relations can be learned implicitly is found in the contextual cueing task (Chun & Phelps 1999; Greene et al. 2007) and the relational manipulation task (Ryan et al. 2000). Furthermore, evidence that implicit relational learning can be expressed under novelty is demonstrated by the transitive inference task (Ellenbogen et al. 2007; Greene et al. 2001; 2006).

These additional relational priming tasks may also assess the semantic priming hypothesis put forth in the target article. The authors propose that relational priming is a facet of semantic priming (e.g., “chicken is to hen” may semantically prime the term “female” so that “horse is to mare” may be correctly selected), suggesting that implicit relational learning is verbally mediated. However, because several implicit relational learning tasks, including analogy, are nonverbal tasks, it is difficult to hypothesize that such verbal primes necessarily mediate performance, particularly in the spatial tasks (contextual cueing and relational manipulation). In fact, some evidence suggests that verbal strategies may bias the use of explicit strategies when implicit strategies could otherwise be employed (for a review, see Greene 2007), suggesting that it may be more fruitful to explore the relational priming hypothesis using nonverbal stimuli.

The evidence that analogical transfer and other forms of relational priming can be accomplished implicitly forces a major change in the theory of memory. It has previously been asserted that only deliberative processes could support learning that is sufficiently flexible for abstract application under novelty (e.g., Cohen et al. 1997; Reber et al. 1996). It is now evident that such relational flexibility is not solely the domain of declarative memory (Chun & Phelps 1999; Ellenbogen et al. 2007; Greene 2007; Greene et al. 2001; 2006; 2007; Gross & Greene 2007; Ryan et al. 2000). The consequence of the discovery of such implicit tasks is that the declarative memory model (e.g., Squire 1992) must be considered inadequate in its current form.

No way to start a space program: Associationism as a launch pad for analogical reasoning

doi:10.1017/S0140525X08004585

Keith J. Holyoak^a and John E. Hummel^b

^aDepartment of Psychology, University of California, Los Angeles, Los Angeles, CA 90095-1563; ^bDepartment of Psychology, University of Illinois at Champaign–Urbana, Champaign, IL 61820.

holyoak@lifesci.ucla.edu

<http://reasoninglab.psych.ucla.edu/>

jehummel@cyrus.psych.uiuc.edu

<http://www.psych.uiuc.edu/people/showprofile.php?id=543>

Abstract: Humans, including preschool children, exhibit role-based relational reasoning, of which analogical reasoning is a canonical example. The “role-less” connectionist model proposed in the target article is only capable of conditional paired-associate learning.

Here's a source analogue for the model proposed in this target article: A politician wants to start a space program. Lacking typical prerequisites such as rockets, he gets his assistant to climb the highest lookout tower in the land. We have lift-off!

Analogical reasoning is a canonical case of role-based relational reasoning (RRR), a capability (perhaps uniquely human) that typically emerges around age 5 years (Doumas & Hummel 2005b; Halford 1993; Penn et al. 2008). Consider the example of 4-year-old Lori, a participant in one of the earliest experiments on analogy development (Holyoak et al. 1984). Lori was read a fairy

tale about a genie faced with the problem of transferring a number of jewels from his current bottle to a new home in a different bottle. The genie's solution was to command his magic carpet to roll itself into a tube, place it between the two bottles, and then roll his jewels through it. Without mentioning any connection with the fairy tale, the experimenter asked Lori to figure out how to transfer some small balls from one bowl to another at some distance (while remaining seated). A variety of objects were available, including a large rectangular sheet of paper. Lori referred to the balls as "pretend jewels." Looking at the sheet of paper she said, "That will be a magic carpet." She laughed as she picked it up and rolled it. "I did it just like the genie!"

Lori's reasoning exploited systematic correspondences between objects filling parallel roles. For example, the genie wants to move his jewels from one bottle to another just as Lori wants to move balls across bowls. More generally, RRR implies the ability to draw inferences about entities based on the roles they fill in relations, where the roles are not predictable by features of the entities and the relations cannot be coded as role-less chunks (Penn et al. 2008).

The target article illustrates an all-too-common approach to connectionist modeling of those cognitive processes most central to human intelligence: suck the essence out, then force-fit what's left into an associationist straitjacket. The project begins by reducing analogical reasoning to a glorified paired associates task. Leech et al. focus on four-term analogies, which have the virtue (for associationist purposes) of providing a highly stereotyped format. Their three-layered network is trained with "facts" glossed along the lines "apple + knife → sliced apple" and "bread + knife → sliced bread." After extensive training, the model is given the inputs "apple" and "sliced apple," thereby activating "knife," which is then clamped so that when "bread" is added as input it combines with "knife" to yield "sliced bread." Analogy solved?

Leech et al. talk about such problems using concepts such as "causal transformations," but the model itself simply learns bidirectional conditional paired associates. Let "object at time 1" be S_i , "causal relation" be C_i , and "object at time 2" be R_i . The model is trained with triples, including $\langle S_1, C_1, R_1 \rangle$ and $\langle S_2, C_1, R_2 \rangle$. The key to the model's performance is that it initially learns all pairwise conditional associations, including $\langle S_i, C_i \rangle \rightarrow R_i$ and $\langle S_i, R_i \rangle \rightarrow C_i$. At test, the input $\langle S_1, R_1 \rangle$ outputs C_1 ; then S_2 is added, and $\langle S_2, C_1 \rangle$ outputs R_2 .

The model operates without explicit representations of relational roles, such as "cause," "effect," or "instrument." In order to fit the four-term analogy format, the modelers hand-code the assignments of roles to predetermined banks in the input and output layers. But although four-term analogies are indeed stereotyped, they are nonetheless richer than conditional paired associates. Consider a couple of small variations. People who understand what a knife does could also solve the analogy

sliced-apple : apple :: sliced-bread : ??,

where the role assignments are reversed. We suspect that the Leech et al. model will be unable to solve this simple variation without additional training, as it has never previously seen, for example, "sliced apple" assigned to its input layer. The only way the model could solve this rearranged analogy is if Leech et al. hand-code the familiar role assignments for it. But if such hand-holding is acceptable, it seems the model will be led into a different, equally embarrassing error. Given the problem

apple : knife :: bread : ??,

people will likely reject it or else complete it with "knife." However, if Leech et al. help their model along by placing "knife" in its familiar bank and clamping the output as usual, it seems that the model will output "sliced bread" as a fine "analogical completion." Perhaps Leech et al. can provide simulation results for these examples in their response.

The model's role-less, paired-associate-style representations, inadequate for even the simplest four-term analogies, render it

incapable of solving any problem requiring integration of information across multiple roles. Such capabilities, illustrated by the protocol from Lori, are present in preschool children (Halford 1993). In a gesture toward extending the model to adult performance, Leech et al. apply it to "a large, complex analogy" (sect. 4.2, para. 1): that between World War II and the first Gulf War (Spellman & Holyoak 1992). The most interesting data reported by Spellman and Holyoak, showing that people were systematic in mapping leaders and countries in pairs (either Churchill, Britain → Hussein, Iraq, or else FDR, US → Hussein, Iraq), can only be explained by a model capable of relational integration (Hummel & Holyoak 1997). But rather than modeling how a reasoner could sort out the interrelationships among two source countries (Britain and the United States) and their leaders with respect to Hitler's Germany, Leech et al. simply eliminate this difficulty – by adopting representations of World War II that leave the United States out of it.

Relational priming is an important phenomenon. But as Spellman et al. (2001) reported in demonstrating it, it requires attentional resources (contrary to Leech et al.'s claim that "relational priming is a robust psychological phenomenon that does not require explicit strategic control"; sect. 2.1, para. 4). Neuroimaging data indicate that evaluating causal (as opposed to merely associative) relations activates the prefrontal cortex (Satpute et al. 2005), a brain area that is slow to develop. In fact, we are unaware of any demonstration of relational priming in children (and Leech et al. do not mention any). In the absence of evidence that young children actually exhibit relational priming, it seems premature to assume that it precedes, rather than follows, development of relational roles. For a recent model of how relational roles could be acquired from experience, see Doumas et al. (2008).

Returning to the "space program" analogy: Lacking any sense of relational roles, the target model won't get the point of the analogy. Do you?

Dynamic sets of potentially interchangeable connotations: A theory of mental objects

doi:10.1017/S0140525X08004597

Alexandre Linhares

Getulio Vargas Foundation, Rio de Janeiro 22250-900, Brazil.

linhares@clubofrome.org.br

<http://www.intuition-sciences.com/linhares>

Abstract: Analogy-making is an ability with which we can abstract from surface similarities and perceive deep, meaningful similarities between different mental objects and situations. I propose that mental objects are dynamically changing sets of potentially interchangeable connotations. Unfortunately, most models of analogy seem devoid of both semantics and relevance-extraction, postulating analogy as a one-to-one mapping devoid of connotation transfer.

Leech et al. provide an ambitious framework with which to view analogy as relational priming. I find the idea full of potential, and I fully agree with a number of points. Having worked over the last decade on the question of analogy and abstractions in chess cognition, I share with the authors the following conclusions: (1) that analogies lie at the core of cognition (Hofstadter 2001); (2) that analogies let us expand our knowledge; and (3) that, as the target article proposes, relations can be viewed as transformations. These are certainly true in a chess player's mind (Linhares 2005; submitted; Linhares & Brum 2007).

There are, however, implicit, unstated, assumptions that should be seriously faced by any computational model of cognition, as failure to do so generally leads to simplistic and semantically vacuous models.

First, there is the *representational module assumption*. The model presented by Leech et al. supposedly makes analogies

between apples, lemons, Hitler, Hussein, wars, and the like. The representational module assumption presupposes that a separate representation module would provide these mental objects (e.g., entities in short-term memory [STM]) as input to Leech et al.'s computational model. It is an unfeasible idea for which there is no space to repeat previous arguments (Chalmers et al. 1992; French 2000; Hofstadter & the Fluid Analogies Research Group 1995). Hence, I will consider a related, second assumption: metaphysical realism, a doctrine which posits that objects (bounded entities), and their properties and relations, are independent of observers and, hence, can have one static description (Linhares 2000; Putnam 1981; Smith 1996).

As anyone who has tried to produce a robot navigating real-world settings knows, such robots cannot see people, animals, chairs, or Coke cans, as humans do effortlessly. Robots deal with wave and signal processing. It takes effort to carve waves into mental objects: This process involves setting boundaries on waves (e.g., speech recognition, visual segmentation, etc.), finding properties of these newly bounded entities, finding relations to other bounded entities, and classifying these into concepts.

Objects, properties, and relations are context- and observer-dependent. Consider DNA: DNA is like a staircase. DNA is like a will. DNA is like a fingerprint. DNA is like a zipper. DNA is like a computer program. DNA is like a tape. DNA is like a train track that is pulled apart by the train. DNA is like a tiny instruction manual for your cells.

How can a molecule acquire the properties and relations of objects so semantically apart as are train tracks, computer code, fingerprints, staircases? What are mental objects made of? I suggest that mental objects are dynamic sets of connotations, and that connotations are potentially interchangeable – two characteristics which are ignored by most cognitive models (including Leech et al.'s).

Mental objects, I suggest, are dynamically built. Each concept, and each instance of a concept, has a set of connotations attached to it. But this set is not fixed. It changes dynamically because of contextual pressures. And what are such connotations like? They are either (i) rules for carving waves (sounds into separate spoken words, an image into a set of objects, etc.), (ii) relations between mental objects (a chess queen that pins a king), or (iii) properties of particular objects (red-ness). Most importantly, these connotations are *potentially interchangeable between objects*. This is why DNA as a mental object can acquire so many characteristics that are found far beyond the realm of chemistry.

Analogies are mechanisms by which a mental object acquires connotations from different mental objects. This theory stems from Hofstadter and the Fluid Analogies Research Group (1995), is close to the model of Fauconnier and Turner (2002), and is different from a one-to-one mapping.

Leech et al.'s model has promising ideas, but does not account for this. It assumes that the *perfect set of mental objects* has been constructed *a priori, independently of the analogy in question*. Hitler had a mustache. Hussein had a mustache. Why doesn't their analogy consider this mapping? Answer: because it is irrelevant. But *who is to say so?* Why is this irrelevant, even though it would map beautifully? War in Iraq is hard because it is hot; roads have literally melted as American tanks drove by (Pagonis & Cruikshank 1994). War in Germany is hard because it is cold, and night fires are enemy-attracting (Ambrose 1998). Who is to decide that this is not relevant?

By providing perfectly built objects *a priori*, this work does not reflect the dynamic nature of mental objects. It succumbs to the Eliza effect: nothing is known about Hitler besides a token. Although readers see the full imagery of Adolf Hitler come to mind, with the enormous set of powerful connotations that name brings up (Nazism, the Aryan race, the swastika, the propaganda, WWII, Auschwitz, etc.), the model is satisfied with a single token without any connotations attached. I invite readers to swap all tokens (apple, Hitler, etc.) with randomly chosen Greek letters and reread the target article. However interesting

the psychological constraints posed by the authors, and however rigorous their attempt to remain close to biological plausibility, the model never makes, in any significant sense of the word, the analogies humans effortlessly do. The one-to-one mapping model has no connotation transfer and does not reflect the dynamic nature of mental objects.

Analogical inferences are central to analogy

doi:10.1017/S0140525X08004603

Arthur B. Markman and Jeffrey P. Laux

Department of Psychology, University of Texas, Austin, TX 78712.

markman@psy.utexas.edu

laux@mail.utexas.edu

<http://www.psy.utexas.edu/psy/FACULTY/Markman/index.html>

Abstract: It is important to take a developmental approach to the problem of analogy. One limitation of this approach, however, is that it does not deal with the complexity of making analogical inferences. There are a few key principles of analogical inference that are not well captured by the analogical relational priming (ARP) model.

The developmental psychology literature has necessarily used simple tasks to study analogical reasoning ability in children. This research has demonstrated that children can complete simple A:B::C:D analogies when the relations are simple and well known (e.g., Gentner & Toupin 1986; Goswami & Brown 1989). An unfortunate side-effect of these studies is that they focused research primarily on factors that influence children's abilities to form correspondences between situations on the basis of relational similarities.

The ability to form relational correspondences is crucial to analogical reasoning, but it is only one subcomponent of the process. Analogies are central to cognitive processing because they allow people to extend their knowledge of one domain by virtue of its similarity to another domain (Clement & Gentner 1991; Markman 1997). This extension of knowledge is accomplished via analogical inferences.

Analogical inference occurs when people take facts about a base domain that are connected to the match between the base and target and posit that those facts are also true of the target domain. Although analogical inference has not been studied extensively in development, it is clear that children draw analogical inferences frequently. For example, in their seminal studies of children's mental models, Vosniadou and Brewer (1992) found that children's beliefs about the earth were strongly influenced by simple analogies. For example, some children knew that the Earth is round, but believed it to be round like a pancake, and so they posited a flat round Earth with people living on the top. Other children knew that the Earth was round like a ball, but assumed that the people lived inside the ball with stars painted on the top. In each case, children were using elements of a known base domain (e.g., pancakes and balls) and transferred knowledge from that base to the less well-known domain of the solar system.

The relational priming model is too limited to account for analogical inference. Obviously, as the model stands, it has no mechanisms for making inferences. More importantly, it is not obvious how such mechanisms could be added in a way that would respect what is known about the inferences people make.

It is crucial that analogical inferences are constrained in some way because this prevents analogies from positing that every fact that is true about the base is also true of the target. (Thus, while a child might believe that the Earth is flat like a pancake, that child is unlikely to think that the Earth would taste good with syrup.) In their WWII–Gulf War analogy simulation, the authors coded only the facts that were relevant, and this was critical to the model's success. The simulation also appears to have benefited from some external control structure that always suppressed

the appropriate layers and interleaved the appropriate control process at just the right time to ensure that each of the necessary relations was picked out in turn. How else could it be that the model never cycled through the same relation twice, or searched for a nonexistent element and became stuck?

In contrast, structural alignment assumes that inferences involve facts from the base that are connected to matching higher-order relations between base and target (Clement & Gentner 1991). This systematic relational structure and the preference for systematicity thus provide constraints on inferences such that structural accounts can function even with rich natural concepts and without any external direction. In addition, the inferences can easily be incorporated into the representation of the target domain.

The authors of the target article try to head off criticisms of this variety by suggesting that explicit mappings (and presumably inferences) could be carried out by different processes than the more implicit processes that find correspondences between domains. The authors use the example of semantic priming in language to illustrate their point. If their suggestion turns out to be correct, then it is those processes that could form the basis for a new theory of analogy. Therefore, the theory posited by the authors may help us to understand some of the sub-processes that are recruited during analogical processing, but it is not actually a theory of analogical processing itself. Indeed, it is worth noting that semantic priming is not taken to be a theory of language; rather, it is understood to be a sub-process that is used in language.

If there were no computational models of analogical reasoning that encompassed both mapping and inference processes, and if those models had never been applied to both developmental and adult data, then it might be reasonable to divide these processes into separate components and assume that two distinct models are required to account for them. However, models like the Structure-Mapping Engine (SME) (Falkenhainer et al. 1989) and Learning and Inference with Schemas and Analogies (LISA) (Hummel & Holyoak 1997) are designed to account for both analogical mapping and inference, and both models are able to make use of higher-order relations in their domain representations. Furthermore, as the target article notes, SME has been applied to developmental tasks (Gentner et al. 1995). Thus, it seems unparsimonious to assume that analogical reasoning abilities begin with processes that cannot ultimately perform the variety of tasks that are clearly part of the repertoire of older children and adults.

Although a developmental approach to analogy has the potential to offer great value, it must ultimately point the way toward adult analogical competence in order to actually deliver that value. That is, to be a successful developmental account, a theory must begin at a reasonable starting point and demonstrate the path/process through which the system progresses to reach the known end state. The ARP theory does not explain full competence, and cannot, in principle, be extended to do so without it becoming a part of a larger theory.

Neurocognitive process constraints on analogy: What changes to allow children to reason like adults?

doi:10.1017/S0140525X08004615

Robert G. Morrison^a and Soohyun Cho^b

^aDepartment of Psychology, Northwestern University, Evanston, IL 60208-2710; ^bDepartment of Psychology, University of California, Los Angeles, Los Angeles, CA 90095-1563.

morrison@northwestern.edu

http://www.morrisonlab.org

soohyun@psych.ucla.edu

http://shcho.bol.ucla.edu

Abstract: Analogy employs a neurocognitive working-memory (WM) system to activate and bind relational representations, integrate

multiple relations, and suppress distracting information. Analogy experiments exploring these processes have used a variety of methodologies including dual tasks, neuropsychology, and functional neuroimaging, as well as experiments with children and older adults. Collectively, these experiments provide a rich set of results useful in evaluating any model of analogy and its development.

Analogy involves a structured comparison, or *mapping*, between one situation (*source*) and another (*target*). For instance, a reasoner may be given a problem such as:

bird:nest::bear: ? =

and be asked which word, CAVE or HONEY, completes the analogy. To choose CAVE, the participant would need to realize that birds *live in* nests as bears *live in* caves while not being distracted by the fact that bears *eat* honey. Using several priming tasks, Spellman et al. (2001) investigated whether analogy might just be a consequence of the organization of concepts in semantic memory. They found that unlike traditional semantic priming, “analogical” priming was not automatic and instead required the participant to direct attention to relations between word pairs. This suggested that controlled retrieval of a bound relation into working memory (WM) may be a necessary process for analogical reasoning. Subsequent experiments demonstrated that WM was indeed important for analogical mapping (e.g., Morrison et al. 2001), as well as relational binding (see Morrison 2005), a finding confirmed using functional magnetic resonance imaging (fMRI; Bunge et al. 2005).

WM is also important for suppressing distracting information, such as irrelevant semantic associates or featural similarities likely to enter WM during analogical retrieval and mapping. Waltz et al. (2000) demonstrated that adults performing a semantically rich scene-analogy task shifted from preferring analogical to featural mappings under WM dual-tasks. Using the same task, Morrison et al. (2004) found that frontal patients with damage to WM areas showed a similar pattern. Morrison et al. also developed an A:B::C:D or D' verbal analogy task that required participants to choose between D (analogically correct choice) and D' (foil), which were both semantically related to the C term of the analogy. When the foil was more semantically associated to the C term than was the correct choice, frontal patients performed near chance. In contrast, semantic dementia patients who exhibited profound decrements in relational knowledge performed poorly on all of the verbal analogies regardless of the degree of semantic association between C:D and C:D'. Using the same task, Cho et al. (2007b) found that individuals who scored higher on the Raven's Progressive Matrices (RPM) showed greater fMRI activation increase in neural areas, including the prefrontal and visual cortices, on trials in which reasoners had to reject foils that were highly associated with the C term. This finding suggests that there are neural regions whose level of activation for interference resolution during analogical reasoning relates to individual differences in fluid intellectual capacity.

Many real-world analogies, as well as reasoning tasks developed for psychometric purposes such as the RPM and People Pieces Analogy task (PPA; Sternberg 1977b), require integration of multiple relations to map more relationally complex analogies. Numerous fMRI studies (e.g., Christoff et al. 2001; Kroger et al. 2002) have shown increasing levels of activation in anterior prefrontal cortex for more relationally complex RPM problems, a finding consistent with a neuropsychological study with frontal patients (Waltz et al. 1999). Using an adaptation of the PPA task, Viskontas et al. (2004) found that older adults showed decrements in both relational integration and relational distraction. Using this same task, Cho et al. (2007a) found that executive resources are shared between relational integration and interference resolution during analogical reasoning. In an fMRI follow-up study, Cho et al. (2007c) found partially overlapping but distinct regions within inferior

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 function¹ $b = F_R(a)$. 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* $\rightarrow$ *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).

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

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.

of a set of stimuli in novel relationships. Suppose, for example, in one task instance (T1) *square predicts circle*, *circle predicts triangle*, and *triangle predicts square*; in the next task instance (T2) *cross predicts star*, *star predicts bar*, and *bar predicts cross*; and so on. The transverse patterning structure and the fact *cross predicts star* (information trial) are sufficient to correctly predict the responses to *star* and *bar* in the other two trials. Even on more complex structures involving more objects and more information trials, adults reach the point of correctly predicting the responses on the remaining trials (Halford et al. 1998).

The target authors' model fails to account for this sort of abstract analogy because the system can only utilize relations between objects that have already been learned as transformation functions on the basis of prior experience. Analysis of internal representations by the authors revealed that the developed network groups objects in hidden unit activation space by the relations that transform them. The input/hidden-to-output connections effectively implement a mapping whose domain is partitioned into subdomains, one for each causal relation (e.g., *cut*, *bruised*, etc.). The input-to-hidden connections implement a mapping from object pairs to points located within the subdomain corresponding to the relationship between the two objects, effectively providing an index to the objects' relation. For example, *apple* and *cut apple* are mapped to a point in hidden unit space contained in the subdomain for the *cut* transformation function. This point provides the context for mapping the next object, say, *banana* to *cut banana* (assuming that this transformation was also learned) to complete the analogy. The same sequence of steps may also be applied to transverse patterning, assuming that the network has learned all the required mappings: For example, *cross* and *star* would map to a point in the subdomain corresponding to the task relation T2; and *star* in the context of T2 would map to *bar*. Unlike adults, however, the network must be trained on all possible transformations to make this inference.

Notice that the problem with Leech et al.'s model is not about a complete failure to generalize. Suitably configured, some degree of generalization may be achievable using a learned internal similarity space of object representations. All fruit, for example, could be represented along a common dimension, and the various causal relations could be orthogonal projections that systematically translate the representations of *fruit* to *cut fruit*, or *bruised fruit*, and so on. Learning to complete analogies for some instances of fruit and cut fruit may generalize to the other instances, assuming the number of parameters (weights) implementing the mappings is sufficiently small compared to the number of fruit examples. But the elements of a transverse patterning task may not be systematically related in any way other than via the transverse patterning structure; they need not belong to the same category of objects, and they may even contradict mappings learned from a previous task instance (e.g., *cross* may predict *bar* in a new instance of the task). Thus, there is no basis on which the needed similarity could have developed. The problem is that the capacity for abstract analogical inference transcends specific object relationships.

Despite this pessimistic assessment, perhaps an explanation for analogy could be based on transformations augmented with processes that represent and manipulate symbols. Assuming a capacity to bind/unbind representations of objects to representations of symbols, abstract analogies such as transverse patterning may be realized as the transformation of symbols (e.g., symbol *a* maps to *b*, *b* maps to *c*, and *c* maps to *a*), instead of specific object representations. However, hybrid theories are to be judged at a higher explanatory standard (Aizawa 2002). Not only are they required to explain each component (e.g., an object transformation account for concrete analogies and a symbolization account for abstract analogies), but they also need to explain why the components are split that way.

Indeed, Aizawa's detailed analysis of the systematicity problem (Fodor & Pylyshyn 1988) and its proposed "solutions" (for a review, see Phillips 2007) signpost a general developmental theory of analogy. To paraphrase, the problem is not to show how analogy is possible under particular assumptions, but to show how analogy is a necessary consequence of those assumptions. The capacity for analogy, like the property of systematicity, is a ubiquitous product of normal cognitive development. If a developmental connectionist explanation depends on a particular network configuration, then why does it get configured that way? And if the answer is an appeal to error minimization, then what preserves this configuration in the face of optimization over numerous stimulus relations that may have nothing to do with analogy? Answers to these sorts of questions without relying on what Aizawa distinguishes as ad hoc assumptions would help to shift Leech et al.'s account from one that is simply compatible with the data to one that actually explains it.

Leech et al.'s developmental approach may yield valuable insights into the early acquisition of a capacity for concrete analogical inference. But to expect that it will lead directly to higher cognition seems more like wishful thinking.

Relation priming, the lexical boost, and alignment in dialogue

doi:10.1017/S0140525X08004640

Claudine N. Raffray, Martin J. Pickering, and Holly P. Branigan

Department of Psychology, University of Edinburgh, Edinburgh, EH8 9JZ, United Kingdom.

claudine.raffray@ed.ac.uk

martin.pickering@ed.ac.uk

holly.branigan@ed.ac.uk

http://www.psy.ed.ac.uk/people/s0340872/index_html

http://www.psy.ed.ac.uk/people/martinp/index_html

http://www.psy.ed.ac.uk/people/holly/index_html

Abstract: The authors' claim that analogical reasoning is the product of relational priming is compatible with language processing work that emphasizes the role of low-level automatic processes in the alignment of situation models in dialogue. However, their model ignores recent behavioral evidence demonstrating a "lexical boost" effect on relational priming. We discuss implications of these data.

Leech et al. present a connectionist model of analogical reasoning based on relation priming, rather than on explicit structure-mapping processes. Their core idea is that priming is itself a mechanism for producing analogy, and from it ultimately emerges the relation that is critical for establishing the similarity between a pair of terms in one domain and a pair of terms in a second domain. This claim is compatible with recent work in language processing that emphasizes the role of "low-level" priming in the development of semantic representations. This is most apparent in work on dialogue, in which interlocutors prime each other to produce equivalent situation models that form the basis of mutual understanding (Pickering & Garrod 2004). For example, interlocutors tend to repeat each other's choice of reference frames or ways of interpreting complex arrays (Garrod & Anderson 1987; Schober 1993). Clearly, alignment of analogical structures constitutes an important part of such situation models.

Critically, Pickering and Garrod's (2004) framework suggests that alignment takes place at many linguistic levels, and that repetition at low levels such as words enhances alignment at higher levels, such as the situation model. It follows that lexical repetition should enhance relational priming, and therefore analogical reasoning. Raffray et al. (2007) directly addressed the issue of the effects of lexical repetition (of the head or modifier) on

relation priming of noun-noun combinations such as *dog scarf*. Three expression-picture matching experiments investigated whether relation priming occurred in the context of head repetition, modifier repetition, or both, and allowed direct comparison of the effects of head and modifier repetition. Results showed that participants were more likely to interpret *dog scarf* as a scarf decorated with a picture of a dog (i.e., *dog DESCRIBES scarf*) than as a scarf worn by a dog (i.e., *dog POSSESSES scarf*) after interpreting another expression involving the description relation rather than the possession relation; but the priming was greater when one term was repeated (e.g., *dog T-shirt* or *rabbit scarf*) than if neither was repeated (e.g., *rabbit T-shirt*). In sum, while conceptual relations were independently primed, the level of activation that a given relation received was enhanced where there was repetition of lexical items between prime and target.

We propose that such “lexical boost” effects, similar to those found in syntactic priming studies (Pickering & Branigan 1998), mean that priming of analogical relations should be enhanced by any repetition of terms. In Goswami and Brown (1989), the participant infers that *lemon* is to *cut lemon* as *bread* is to *cut bread*. Importantly, the concept of a lexical boost within analogical reasoning only makes sense in the context of two- (or more) place relations. That is, to get a lexical boost we would need to consider analogies such as *boy & ball* is to *boy kicks ball* as *man & stone* is to *man kicks stone*. In this case, the lexical boost predicts that participants should find it easier to resolve analogies containing repeated terms, such as *boy & ball* is to *boy kicks ball* as *man & ball* is to *man kicks ball*, or similarly *boy & ball* is to *boy kicks ball* as *boy & stone* is to *boy kicks stone*. For more complex analogies, the prediction is that any repetition of concepts will enhance analogy. To take the authors’ example, it should be easier to draw the analogy from *World War II* to *World War I* than to the *Gulf War*, because more of the objects (e.g., *Germany*) are repeated (see Table 3 of the target article, sect. 4.1.2). Whereas the analogy between *Churchill orders_attack_of Germany* and *Bush orders_attack_of Iraq* involves different Object 1s and Object 2s, the analogy between *Churchill orders_attack_of Germany* and *Lloyd George orders_attack_of Germany* involves different Object 1s but the same Object 2. If such analogy works like the priming effects we have discussed, then lexical repetition should facilitate analogical reasoning.

There is also evidence for a semantic boost to syntactic priming (Cleland & Pickering 2003), so that priming is stronger when terms are semantically related than when they are not. For example, participants are more likely to describe a red sheep as *The sheep that is red* after hearing *The goat that is red* than after hearing *The door that is red*. It might similarly be the case that priming of analogical relations is enhanced by the inclusion of semantically related terms. That is, *boy & stone* is to *boy kicks stone* as *man & pebble* is to *man kicks pebble* might be easier to process than an analogy that contains semantically unrelated terms, such as *boy & stone* is to *boy kicks stone* as *man & ball* is to *man kicks ball*. Such effects should affect both the speed and the likelihood of obtaining a particular analogy.

Pickering and Garrod’s (2004) model of dialogue assumes that processes by which interlocutors align their models of the situation are largely the result of automatic priming. In particular, repetition at one level of representation enhances repetition at other levels. For example, Branigan et al. (2000) showed that interlocutors are more likely to repeat each other’s grammatical choices if they also repeat each other’s choice of verbs. More generally, we assume that if two people start to use the same words, they start to take on board the same conceptualizations. One of the functions of analogical reasoning is to assist in the process by which interlocutors end up with equivalent situation models, and therefore we predict that lexical and other repetition will enhance this process. Therefore analogical reasoning can be seen as part of the mechanism of alignment and will be affected by the processes affecting priming.

Child versus adult analogy: The role of systematicity and abstraction in analogy models

doi:10.1017/S0140525X08004652

Angela Schwering and Kai-Uwe Kühnberger

Institute of Cognitive Science, University of Osnabrück, D-49076 Osnabrück, Germany.

aschweri@uni-osnabrueck.de

<http://www.cogsci.uni-osnabrueck.de/~aschweri/>

kkuehnbe@uni-osnabrueck.de

<http://www.cogsci.uni-osnabrueck.de/~kkuehnbe/>

Abstract: The target article develops a computational connectionist model for analogy-making from a developmental perspective and evaluates this model using simple analogies. Our commentary critically reviews the advantages and limits of this approach, in particular with respect to its expressive power, its capability to generalize across analogous structure and analyze systematicity in analogies.

Leech et al. present a computational (connectionist) approach to explain analogy-making from a developmental perspective. At the outset we would like to emphasize that this is a very compelling and advanced approach: Tackling the problem from a developmental point of view enables the authors to highlight completely new aspects of analogy-making. This approach reflects the infantile learning process, from comparisons based mainly on superficial similarity to a controlled and advanced strategy of analogical comparison based on structural systematicity. So far, analogy models have always been inspired by adult analogical reasoning. Emphasizing the learning process, it makes sense to approach this problem from an infantile developmental perspective: Leech et al.’s model is based on a neural network implementing a Hebbian learning algorithm which is able to enhance bit by bit the strategy for analogy-making and explicitly model the development from superficial similarity to structural similarity. More precisely, the network can learn causal relations using a transformation of an object (e.g., “apple”) by a causal agent (e.g., “knife”) to achieve a representation of a transformed object (e.g., “cut apple”). The network extends this ability step by step to different domains and cross-mapping analogies; it can model the relational shift from surface similarity to relational similarity; and finally, it is trained on analogies involving multiple objects and multiple relations. There exists no other comparable analogy model modeling strategic learning – current analogy models can only model analogical learning by analogical transfer.

Another interesting capability is the creative potential of the analogy model: A trained network can creatively construct completely new objects when a relation is applied to a new (target) object. However, this capability must also be seen critically: Any relation can be applied to any (suitable or unsuitable) object and always leads to some result, which might be completely meaningless and absurd.

Inspired by research on infantile development, the authors investigate mainly analogies used in previously conducted analogy experiments with children. These are typically proportional analogies, that is, *a-is-to-b-as-c-is-to-what* analogies. All of these analogies are based only on a single, common relation, which is the same in source and target. We argue that such analogies are oversimplified – the task in these examples is applying the same relation to a new target object rather than making an analogous transfer. The target object is in fact very similar to the source object with respect to the applicability and the outcome of the relation. The “analogical” mapping required to solve the analogy is very small. We do not deny that such oversimplified analogies are necessary to investigate the initial analogical abilities of very young children; however, an analogy model (if it is not limited to modeling the analogy-making capability of 1-to-5-year-old infants, who anyway have only a very limited ability of analogy-making) must foremost have the capability to solve analogies with

a minimum analogous structure. Therefore, the authors extend their investigation to so-called adult analogies. But in fact, their “adult analogy” differs from the “child analogy” only in that it comprises not just one single identical relation, but a *set* of single identical relations. The crucial characteristic of an analogy is, however, the common structure, that is, the analogous relations which are part of a common higher-order structural system. The analogy research community agrees on the fact that “systematicity” is the decisive point that is crucial for an analogy. Interesting analogies usually contain mappings that are only partially consistent. The sets of common relations considered by the authors are only connected via having the same source or target domain (see Fig. 16 in sect. 4.2 of the target article) and therefore have a flat and not a higher-order hierarchic structure. Principally it would be possible to extend the proposed approach to cover higher-order structures as well, by allowing the input of the neural network to represent composed structures – for example, as in the Structure-Mapping Engine (SME), where not only the object sun, but also composed structures such as the mass of the sun $mass(sun)$ and the greater than relation between the masses of sun and planet $mass(sun) > mass(planet)$ can each be represented by a single node. But still, this extended approach would only allow for composing relations with objects, not for decomposing them again.

We also see a problem in the expressive power, which arises in almost all analogy models: Except for logic-based analogy models (such as heuristic-driven theory projection [HOTP]; Gust et al. 2006), no analogy model is able to express complex rules (often requiring quantified variables and logic operations). Many analogies in the physical and mathematical domains require this expressivity to be correctly modeled (for examples, see Gust et al. 2007).

A further aspect that should be mentioned is a restriction imposed on the involved relations: all simple analogies discussed in the target article map only on the same relations. If semantically similar relations (e.g., “Iraq occupied Kuwait,” but “Germany annexed Austria”) are represented closely in the neural network, we can assume that semantically similar relations can also be mapped on each other (which is an advantage of the neural network approach to symbolic approaches). However neural network approaches cannot map two analogous, but semantically different, relations on each other. In the case of the Rutherford analogy, the gravitational force in the solar system and the Coulomb force in the atom are analogous, but they cannot be mapped on each other. At a more general level, both relations could be described as two attracting forces. Many analogy models avoid this mapping problem by representing relations already at the required level of abstraction. However, the generalization of gravitation and Coulomb force to attracting force is part – actually it is the *result*!! – of the analogical reasoning and therefore must not be skipped. Like most other analogy models, the approach discussed here does not support this level of abstraction.

Most of our criticism does not refer specifically to the analogy model by Leech et al., but rather is directed at general shortcomings of most (subsymbolic) approaches. However, subsymbolic approaches have certain advantages in handling vagueness and uncertainty and in modeling learning. To bring together the strengths and overcome the shortcomings we see a lot of potential in connecting both approaches via neuro-symbolic integration.

Analogy is to priming as relations are to transformations

doi:10.1017/S0140525X08004664

Vladimir M. Sloutsky

Center for Cognitive Science and Department of Psychology, Ohio State University, Columbus, OH 43210.

Sloutsky.1@osu.edu

<http://cogdev.cog.ohio-state.edu>

Abstract: The commentary discusses three components of the target proposal: (1) analogy as a host of phenomena, (2) relations as transformations, and (3) analogy as priming. The commentary argues that the first component is potentially productive, but it has yet to be fully developed, whereas the second and third components do not have an obvious way of accounting for multiple counterexamples.

Understanding the mechanism underlying analogy is a notoriously difficult problem. As argued by French (2002, p. 204): “The lessons of almost 40 years of research in the computational modeling of analogy-making have, more than anything else, shown us just how hard the problem is.” Analogy involves using one situation for understanding another situation, and often the only commonality across the two situations is a common relation among the constituent elements. Given that the base and the target relation (or relations) is the only commonality, most approaches have posited that analogy requires explicit representation of relations. However, several findings could be potentially problematic for this idea. Most importantly, there is evidence that organisms whose ability to represent relations explicitly is unclear (such as nonhuman primates and very young children), can perform analogy (Chen et al. 1997; Thompson et al. 1997; see also Goswami 2001, for a review). The target article attempts to attack the problem head-on: that is, by considering alternatives to the idea that analogy requires explicit representation of relations. This is a laudable attempt – providing a low-level explanation of a process that has been traditionally explained by invoking higher-level processes would be an important contribution.

The target article’s proposal has three components: (a) analogy could be a host of phenomena, rather than a unitary phenomenon; (b) relations are construed as transformations between states rather than as explicit symbols; and (c) analogy is the priming of a target relation by a base relation. In this commentary, I focus on each of these components. While I am very sympathetic with the search for alternatives to explicit representation of relations, I am not convinced that the proposed account solves the problem.

The construal of analogy as a set of disparate phenomena is a potentially interesting idea, which attempts to simplify the problem by suggesting a “divide and conquer” approach to analogy. In particular, if analogies performed by nonhuman primates, human infants, and adults are in fact different phenomena, then there is no need to develop a unified account of analogy, or to explain how “analogy as priming” gives rise to more complex forms of analogical reasoning. However, a potential downside of this approach is that in the absence of a clear taxonomy of these putatively disparate phenomena (what these phenomena are, how they are interrelated, and how they differ), it is not clear what the theory of analogy purports to explain. In my view, it remains to be seen whether the idea is truly productive, or whether it offers more “divide” than “conquer” in the study of analogy.

The idea of relation-as-transformation, although promising, seems to cover only a small set of situations, such as some of the tasks described by Goswami and Brown (1989). At the same time, even early in development, there are multiple situations where the relation-as-transformation construal is less obvious. These examples include the “animal-habitat” relation, such as “bird to nest as dog to doghouse” (Goswami & Brown 1990); the class inclusion relation (Goswami & Pauen 2005); spatial relations, such as “in the middle” (Lowenstein & Gentner 2005); quantitative relations, such as “monotonic increase” (Kotovsky & Gentner 1996); and identity relations (Thompson et al. 1997). The fact that young children (and even nonhuman primates) successfully perform analogies based on relations that are difficult to construe as transformations suggests that the construal of relations-as-transformations is at least unnecessary to account for the early analogy. In addition, even if some relations could be construed as transformations, surprisingly little is said by Leech et al. about how relations-as-transformations could be learned. At the same time, learning of relations is a nontrivial problem: It has been argued that different

instances of the same relation lack statistical density, and, as a result, it is difficult to learn the categories implicitly, without some sort of supervision (Kloos & Sloutsky 2008).

The third major component of the proposed account is relational priming. Whereas priming is a relatively basic phenomenon, relations that are putatively primed are not. Therefore, it is critically important to understand how relations are represented. Leech et al. suggest that a key to representations of relations (which are construed as transformations) is the causal agent that is responsible for the transformation. For example, the relation "cutting" could be accessed by "knife" as a putative causal agent. However, whereas this idea is relatively clear for "cutting" it is less clear for "melting," which depends on an invisible causal agent (i.e., heat), or for "breaking," which could be implemented by multiple causal agents. If a causal agent is critical for construing relations as transformations, then analogies with more obvious causal agents (e.g., "knife") should be substantially easier than analogies with invisible causal agents (e.g., "heat"), or those with multiple causal agents (e.g., "breaking"). However, there is little evidence in the literature supporting such differential difficulty (cf. Goswami & Brown 1989).

It is also unclear where constraints on priming come from. Objects can enter in multiple relations, each of which could be primed. In addition, a set of entities semantically associated with these objects could be primed as well. Consider the following problem: *Falcon* is to *Pigeon* as *Human* is to what (*Canary*, *Chimp*, or *Chicken*)? In addition to *Canary* (which could be primed by both *falcon* and *pigeon*), the *Falcon-Pigeon* pair can also prime several relations: (1) *Larger* (in which case *Canary* or *Chicken* could be selected); (2) *Similar Animal* (in which case *Chimp* could be selected); or (3) *Predate* (in which case *Chicken* could be selected). Given that priming does not lead to an unambiguous response, response selection becomes a non-trivial problem; however, the proposed account does not offer an obvious way of addressing this problem.

There are also important developmental questions associated with the proposed account. How do children learn that transformations are important? How do they encode, store, and retrieve relations-as-transformations in the absence of an obvious causal agent? Do relations-as-transformation undergo developmental change, or do they remain the same throughout development? And most importantly, how do people acquire the ability to perform analogy when relations cannot be construed as transformations? Overall, while the target article makes a promising proposal, it also provides further evidence of "how hard the problem of analogy is."

ACKNOWLEDGMENTS

Writing of this commentary was supported by grants from the NSF (BCS-0720135) and from the Institute of Education Sciences, U.S. Department of Education (R305H050125), to Vladimir M. Sloutsky. I thank Alex Doumas and John Opfer for their helpful comments.

Computational modeling of analogy: Destined ever to only be metaphor?¹

doi:10.1017/S0140525X08004676

Ann Speed

Cognitive and Exploratory Systems Department, Sandia National Laboratories, Albuquerque, NM 87185-1011.

aespeed@sandia.gov

http://www.sandia.gov

Abstract: The target article by Leech et al. presents a compelling computational theory of analogy-making. However, there is a key difficulty that persists in theoretical treatments of analogy-making, computational and otherwise: namely, the lack of a detailed account of the neurophysiological mechanisms that give rise to analogy behavior. My commentary explores this issue.

The target article by Leech et al. and the hypothesis about the nature of analogy on which it is based, are quite compelling and different from the majority of analogy literature. However, the authors repeat a key difficulty present in other computational modeling efforts. This commentary focuses, first, on the persuasive aspects of the article and model; then, on the difficulties associated with the computational model; and finally, presents a high-level summary of the things that I believe need to be addressed to demonstrate a computational hypothesis that is both psychologically and physiologically plausible regarding how analogy emerges from the brain.

Generally, Leech et al.'s article and model are compelling for a couple of reasons. First is the hypothesis that relational priming is at the core of analogy and the parsimonious theoretical implications for that hypothesis. This theoretical construct easily accounts for adult novice-expert differences (Novick 1988); inter-domain transfer difficulty (e.g., variations on the convergence problem); the ubiquity and effortlessness of analogical transfer in everyday life (Hofstadter 2001; Holyoak & Thagard 1995); and the variety of developmental phenomena cited in the target article. Correspondence with the knowledge accretion hypothesis also provides a convincing reason explaining why analogies are so difficult to elicit in the laboratory despite the general agreement that analogy underlies much of human cognition.

Second is the fact that the authors' computational model uses distributed representations. The debate between those who argue that cognition requires discrete symbolic representation and those who argue that physiology dictates distributed representations is an important one in the cognitive literature (Dietrich & Markman 2003; Spivy 2007), and there is really only one other model that proposes a computational implementation of analogy using distributed representations (Eliasmith & Thagard 2001; see Spivy 2007).

I agree with the authors' assertion that analogy is an emergent phenomenon. However, I would assert that it is not, at the most fundamental level, an emergent property of more basic psychological processes (e.g., relational priming), but is rather an emergent property of basic physiological processes. The authors repeatedly identify that prior models of analogy do not provide mechanisms for analogy from a developmental perspective. I also agree with this assertion, but I believe that the developmental perspective must take into account the thing that is really developing: the brain.

Specifically, the authors' implementation of their hypothesis, especially when applied to the Saddam-Hitler analogy, replicates a key deficiency in other models of analogy. This omission is the lack of an in-depth account of how human neurophysiology produces analogy (despite LISA's [i.e., the Learning and Inference with Schemas and Analogies model's] use of node oscillations to create bindings; Hummel & Holyoak 2003). Such an account must provide several details:

1. What is the nature of information representation? How are those representations physically manifested in the brain? As the authors observe, "how object attributes and relations are represented is important for modeling analogy because it constrains what else can take place" (sect. 2.2, para. 1). Although I agree with the spirit of this statement, I would argue that the knowledge representation must be firmly grounded in neurophysiological processes – that to ignore neurophysiological mechanisms as we currently understand them is to ignore important constraints on the resulting cognitive behavior.

2. How do those representations come to exist? How does the physical brain transform energetic information (e.g., photons, sound waves) into the hypothesized representation format? Using distributed representations, as in the target article, partially addresses the first issue, but it is still unclear how these distributed representations might come to exist. Nor is it clear what these representations are analogous to in the brain. This issue is not taken into detailed account in any existing computational model of analogy or higher-order cognition, as far as I am

aware. Divorcing higher-order cognition from low-level sensory processes ignores the fact that the brain is connected to physical sensors and that higher-order cognitive processes necessarily use the output of those physical sensors.

3. Once these grounded representations exist, how do they give rise to, or contribute to the “calculation” of analogy (and other higher-order cognitive processes)? An extension of the first two points, this issue has been partially addressed from a theoretical perspective (e.g., Mesulam 1998; cf. the literature on prefrontal cortex function) but has yet to be fully addressed computationally. However, there exists a significant body of literature that is coming closer to making this link (e.g., Barabas 2000; Elston 2003; Mogami & Tanaka 2006; Rougier et al. 2005; Tanaka 1992). One implication for analogy is that existing computational models are totally unable to identify structural characteristics of problems on their own – a key feature in analogy-making, and the thing that differentiates experts from novices and positive from negative transfer (Gentner & Markman 1997; Novick 1988; Ross 1987; Ross & Kennedy 1990). Even though the first model in the target article takes steps in this direction, it remains unclear whether the mechanisms by which it accomplishes analogy are those used by humans.

As computational models of cognitive processes are hypotheses about the underlying mechanisms giving rise to cognitive behaviors, it seems necessary to account for known neurophysiological processes. Of course, this approach will radically increase the complexity of the resulting models, but dynamical systems theorists would argue that the typical reductionist approach, which yields crisp, transparent, simple-to-analyze models, likely does not produce theories that accurately reflect human cognitive mechanisms.

And this point is the core of the issue that my colleagues and I believe modelers need to begin addressing – how does the brain, as a physical system, take raw sensory information and perform higher-order cognition using that information (e.g., solving Raven's Progressive Matrices)? And, how much of this physiological detail do we as a modeling community need to include in our models in order to accurately emulate what the brain/mind is really doing?

As Stephen Grossberg has said, theories of consciousness posited in the absence of links to neurophysiology can only ever be metaphors (Grossberg 1999). How long before our computational models of analogy become more than just metaphor?

NOTE

1. The author of this commentary is employed by a government agency and as such this commentary is considered a work of the U.S. government and not subject to copyright within the United States.

ACKNOWLEDGMENT

Sandia National Laboratories is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under Contract DE-AC04-94-AL85000.

Creativity or mental illness: Possible errors of relational priming in neural networks of the brain

doi:10.1017/S0140525X08004688

James E. Swain^a and John D. Swain^b

^aChild Study Center, Yale University School of Medicine, New Haven, CT 06520; ^bDepartment of Physics, Northeastern University, Boston, MA 02115.

james.swain@yale.edu

<http://myprofile.cos.com/jameseswain>

John.swain@cern.ch

<http://www.physics.neu.edu/Department/Vtwo/faculty/swain.htm>

Abstract: If connectionist computational models explain the acquisition of complex cognitive skills, errors in such models would also help explain unusual brain activity such as in creativity – as well as in mental illness, including childhood onset problems with social behaviors in autism, the inability to maintain focus in attention deficit and hyperactivity disorder (ADHD), and the lack of motivation of depression disorders.

Malfunction of artificial neural networks, such as those that include analogy priming, might cause those networks to make mistakes about the nature of reality and display analogous behaviour to human brains and minds which are not normal. This may be especially relevant to understanding creativity, as well as the development of childhood onset neuropsychiatric illnesses. The following are three such types of defect in computational models of reasoning – well known in the artificial intelligence field – which can arise in analogy priming. (For general references, see Rumelhart & McClelland [1986], Haykin [1998], Rojas and Feldman [1996], and Muller et al. [1995].)

Example 1. In the first example (Fraser 1998), a network was trained on pictures of a landscape with and without a tank present. In those with the tank, it was placed in all sorts of different locations, obscured to various degrees, shot from different angles. The landscapes without the tank were similarly photographed from various angles, and some pictures from each set were used to train the network, the data being provided to indicate whether or not the tank was present. The goal was to see if the network could “learn” the task of telling whether or not a tank was present in a picture it had not seen – a very simple form of analogical reasoning based on previous pictures but without the need to explicitly isolate the tank as an object in the scene, which hopefully would lead to the ability to recognize arbitrary objects in arbitrary scenes. The pictures not used in training were used as tests – and the results were spectacular, with essentially 100% correct discrimination between pictures with and without tanks present. How the network had done this was stored in all the weights of connections between nodes, and the general feeling was that whatever was going on, it would likely not be understandable by a human being. However, on examination, it turned out that what the network had done was effectively to sum up the brightness of each pixel in the photograph. The pictures with the tank had been taken on a different day than the pictures without, and the network had discovered a significant difference between the two sets of pictures – one set had been taken on a cloudier, and therefore darker, day than the other. The network had learned to classify the pictures by analogy, but had used the total brightness of the scene rather than the presence or absence of the tank.

This “error” might be seen as a basis for unexpected creativity in neural networks, in which new perspectives result from dropping prejudices. For example, in the above example, the input consisted of the raw data. If instead the network had received the input as a list of possible tank-like features (perhaps extracted on the basis of some more logic-based algorithm such as matching features in the scene to features on photographs of tanks), it might have counted up the number of tank-like features and their quality and made a different discrimination. Dropping the preclassification of features in the scene, in a sense, opened up the “creative realization” that the two scenes were of different brightness – something that typical humans might well (and indeed in this case did!) miss. Additionally, it also shows how selective prejudices can sharpen cognition by making some features stand out.

On the other hand, this alternate solution to the problem might explain how autistic children make social misjudgements – perhaps using unusual aspects of a social scene. For example, if the qualities of the tank were like the intricacies of facial expressions, the processing of which appears to be impaired in autism (Schultz 2005), then excluding this information would lead to some of the social errors that autistic children make in missing facial expressions.

Example 2. Networks that have too many connections between too many neurons often do not work well (Müller et al. 1995; Rojas & Feldman 1996). This is perhaps not surprising, since it essentially means that almost no weights (connections) are close to zero. Given the high apoptosis rate in the developing brain, one might wonder whether or not any mental disorders are associated with defects in this apoptotic process. Indeed, autism is associated with unusually large brain size (Courchesne et al. 2004). Perhaps future therapies for autism could be based upon restoring normal apoptotic mechanisms during infancy.

Contrasting mechanisms of neurogenesis, neural sprouting, and new synapse formation would also be important in regulating neural network performance. Abnormalities in those new connections and activity, because of either genetic or environmental issues, could lead to problems such as structural non-uniformity in computational models (Rumelhart & McClelland 1986).

Example 3. An efficient neural network must appropriately switch between flexible and stable states (Haykin 1998; Rumelhart & McClelland 1986). The stable state of a neural network might be akin to a focused state. Perhaps difficulties in reaching and maintaining stable states in children's brains manifest as the lack of focus and hyperactivity of attention deficit/hyperactivity disorder (ADHD) (American Psychological Association 2000). Perhaps understanding overactive brain circuits may also inform our understanding of abnormally active cortex in epilepsy. Alternatively, states that are too stable may appear like the psychomotor retardation of depression (Sadock & Sadock 2004). Perhaps treatments such as electroconvulsive therapy in adults are a sort of "reset," helping the brain out of a state of excessive stability. Thinking about mentally ill brains as connectionist neural networks which have an impaired ability to attain, maintain, and switch between stable states may lead to novel therapies aimed at augmenting these brain mechanisms.

Analogy does indeed lie at the heart of the acquisition of human cognition, as Leech et al. posit. Connectionist models of the neural networks in brains may help explain how the acquisition of cognitive skills in humans actually works. In addition, apparent errors in the development and maintenance of these networks, which may be modelled computationally, may mimic aspects of mental illness and lead to improved and alternative treatments. This kind of innovative approach may be especially helpful to understand and treat infants and children who are learning critical cognitive skills, yet are not necessarily able to communicate their problems clearly.

ACKNOWLEDGMENTS

James E. Swain is supported by a grant from the National Alliance for Research on Schizophrenia and Depression, the Yale Center for Risk, Resilience, and Recovery, and associates of the Yale Child Study Center. John D. Swain is partially supported by the National Science Foundation.

Computational complexity analysis can help, but first we need a theory

doi:10.1017/S0140525X0800469X

Todd Wareham,^a Iris van Rooij,^b and Moritz Müller^c

^aDepartment of Computer Science, Memorial University of Newfoundland, St. John's, NL, A1B 3X5, Canada; ^bNijmegen Institute for Cognition and Information, Radboud University Nijmegen, 6525HR Nijmegen, The Netherlands; ^cDepartment of Mathematics, Albert-Ludwigs-Universität Freiburg, 79098 Freiburg, Germany.

harold@cs.mun.ca

<http://www.cs.mun.ca/~harold>

i.vanrooij@nici.ru.nl

<http://www.nici.ru.nl/~irisvr/>

moritz.mueller@math.uni-freiburg.de

<http://home.mathematik.uni-freiburg.de/mueller/home.html>

Abstract: Leech et al. present a connectionist algorithm as a model of (the development) of analogizing, but they do not specify the algorithm's associated computational-level theory, nor its computational complexity. We argue that doing so may be essential for connectionist cognitive models to have full explanatory power and transparency, as well as for assessing their scalability to real-world input domains.

Leech et al. describe a connectionist algorithm that reproduces several known effects in the development of analogy-making. The authors claim that this algorithm models how children develop the ability to make analogies in a manner not yet captured by previous (primarily non-connectionist) models of analogy such as Structure Mapping Theory (SMT) (Gentner 1983). The current version of the algorithm does not account for (the development of) the complex analogies made by adults. Moreover, Leech et al.'s target article is silent on two issues prominent in previous work on analogy, namely: (1) a computational-level theory in the sense of Marr (1982), that is, a precise formulation of a cognitive ability as an input-output function, of which the presented connectionist model supposedly provides an algorithmic-level implementation; and (2) the computational complexity of the proposed algorithm and/or its associated computational-level theory. In this commentary, we discuss why consideration of (1) and (2) may be essential for making progress in research programs such as Leech et al.'s.

To start, we find it useful to re-cast the problem of deriving models of cognitive development (be they algorithmic- or computational-level) in terms of satisfying various constraints. The most basic of these is the empirical constraint, that is, the model must mimic/predict observed cognitive behavior. Though this is often construed only in terms of adult cognitive behavior, the model should be able to fit performance across the different stages of development (e.g., in infancy, childhood, and adulthood) and account for any apparent discontinuities between different stages (e.g., relational shift in the case of analogy). This constraint holds for any model of natural phenomena. In the case of cognitive abilities, which develop over time, Leech et al. point out the need for a developmental constraint, that is, "all proposed mechanisms [of the model] must have a developmental origin" (sect. 5.5, para. 1). That is, the model should incorporate mechanisms which allow the ability to mature consistent with the empirical constraint. Overlooked or ignored so far is a third and equally important constraint, the computational constraint; that is, a cognitive model must satisfy both the empirical constraint and the developmental constraint while operating within the computational resource-limits imposed by the human body and its environment.

Computational complexity analysis is the tool of choice for assessing whether or not a cognitive model can satisfy the computational constraint, thereby placing such analysis at the heart of cognitive modeling. This is not to say that such analysis is easy: Though well-developed measures such as worst-case asymptotic time complexity are applicable to algorithms operating on digital computational architectures, it is not obvious which measures are most appropriate for connectionist algorithms. Potential measures include the time it takes for the network to settle, the number of training-cycles required to develop a given level of performance, and the number of nodes and layers in a network required for computing a given input-output mapping. Once defined, such measures can be used in conjunction with a suitable criterion for computational tractability (see, e.g., van Rooij 2003; in press). Doing so would enable cognitive modelers such as Leech et al. to evaluate how their models' computational resource requirements scale for the larger inputs that are characteristic of real-world domains of analogizing, and to show whether or not modifications are necessary to accommodate adult-level analogizing.

Though algorithmic-level models can be evaluated against the three constraints mentioned above, there are additional benefits in specifying the computational problems that these algorithms

are supposed to be solving – that is, formulating computational-level theories. The first benefit is explanatory transparency: A computational-level theory provides a precise input-output characterization of the cognitive capacity that is to be explained, which is the primary explanandum.¹ In the absence of such a characterization, it is hard to tell if the proposed algorithmic-level theory is explaining anything at all (Cummins 2000). The second benefit is explanatory power: Computational-level theories postulate organizing principles that govern cognitive abilities, which in turn give insight into the rationale of cognitive computations. This is not obviously supported by algorithmic-level theories, especially when we are dealing with connectionist algorithms (Cummins 1995). The third benefit is analytical power: Computational-level complexity analyses can demonstrate that *no* algorithm (let alone a given algorithm) meets the computational constraint for a particular computational-level theory (see also Tsotsos 1990). Moreover, the same analyses can highlight aspects of one's theory that are responsible for excessive resource demands, and as such can guide the formulation of new theories that meet the computational constraint (see van Rooij & Wareham, in press; van Rooij et al. 2005; Wareham 1999).

Formulating computational-level theories of cognitive capacities is not easy, and seems to be particularly hard for connectionist architectures. Yet, such theories can be formulated (for an example, see Thagard 2000), and given the benefits we have identified, it may well be worth the effort. Such theories may counteract the acknowledged temptation to focus on getting connectionist algorithms to work rather than focusing on why they work (Mareschal & Thomas 2007, p. 148), and enable them to actually count as explanations (Green 2001). Such theories may also enable the exploitation of known connectionist-oriented complexity results (Bruck & Goodman 1990; Judd 1990; Parberry 1994), which, given the computational intractability of theories of analogy such as SMT (Veale & Keane 1997), may be crucial in helping approaches such as Leech et al.'s scale to adult-level performance. Finally, computational-level connectionist theories may more clearly expose relationships to non-connectionist theories. For example, reading the target article, we wonder to what extent connectionist algorithms could be trained to map analogies according to the criteria set forth by SMT, and hence to what degree these approaches are complementary rather than competing. Lacking a characterization of the problem that the connectionist network is supposed to be solving, we are so far unable to tell.

NOTE

1. Following Cummins (2000), we consider cognitive capacities to be the primary explananda of cognitive science. The effects considered by Leech et al., on the other hand, are secondary explananda in that they help constrain (via the empirical constraint) theoretical accounts of the cognitive capacity for forming analogies.

Development and evolution of cognition: One doth not fly into flying!

doi:10.1017/S0140525X08004706

Edward A. Wasserman

Department of Psychology, University of Iowa, Iowa City, IA 52242.

ed-wasserman@uiowa.edu

<http://www.psychology.uiowa.edu/Faculty/Wasserman/>

Abstract: Abstract thought, in general, and – reasoning by analogy, in particular, have been said to reside at the very summit of human cognition. Leech et al. endeavor to comprehend the development of analogous thinking in human beings. Applying Leech et al.'s general

approach to the evolution of analogical behavior in animals might also prove to be of considerable value.

He who wisheth one day to fly, must first learn standing and walking and running and climbing and dancing—one doth not fly into flying!
— Friedrich Nietzsche, *Thus Spoke Zarathustra*

Highly complex and abstract skills – such as analogous thinking – have frequently been deemed to represent the very pinnacle of human cognition (e.g., Penn et al. 2008). Leech et al. appear to accept this view, but they do not appear to be content to revel in it. Instead, they seek to understand the emergence of analogous thinking, thereby making its study a decidedly developmental matter.

Leech et al. hypothesize that analogical completion arises from simple cognitive mechanisms. Specifically, they suggest that relational priming is a basic building block for completing analogies. To the extent that their innovative account is successful, at least some key aspects of analogical reasoning may not require the hypothesization of analogy-specific mechanisms.

Leech et al. further suggest that analogical processing may best be viewed as an umbrella term that comprises different task-specific concatenations of basic memory and control processes. Analogy-specific mechanisms may very well exist, but other possibilities should be entertained first because of their greater parsimony and plausibility for infants and young children. Finally, any shifts in children's strategies of task mastery are believed to be the result of children acquiring greater and richer relational knowledge. Leech et al. thus stress the interaction between learning mechanisms and environmental experiences in determining children's developmental trajectory.

Leech et al.'s approach suggests how an advanced cognitive competence – such as analogy formation and performance – can be grounded in more elementary processes, and it promises to provide a fuller picture of the mechanisms underlying the transition from simple to more complex reasoning. If there is at least a seed of truth to Ernst Haeckel's recapitulation theory that “ontogeny recapitulates phylogeny,” then applying Leech et al.'s approach to the evolution of analogical behavior might prove to be valuable as well.

For instance, different species of animals appear to be more or less successful in solving a wide range of relational discrimination problems (Wasserman & Zentall 2006). One of the most intensely studied of such relational discrimination problems is same-different discrimination learning (Cook & Wasserman 2006; Delius 1994). Here, the behavioral evidence suggests that pigeons, baboons, chimpanzees, and humans all can discriminate *first-order* same-different relations; they can reliably report whether two or more stimuli are identical ($A=A$ or $B=B$) or nonidentical ($A \neq B$).

An even more advanced form of same-different discrimination involves *higher-order* relations between *first-order* relations. Task mastery here requires organisms to discriminate groups of two or more stimuli that involve the *same higher-order* relations ($[A=A]=[B=B]$ or $[A \neq B]=[C \neq D]$; both groups of stimuli are the same or both groups of stimuli are different) from groups of two or more stimuli that entail *different higher-order* relations ($[A=A] \neq [C \neq D]$; one group of stimuli is the same and the other group of stimuli is different). Such *higher-order* relations may share important similarities with human analogical reasoning (Thompson & Oden 2000).

Can only human beings discriminate such *higher-order* relations and exhibit analogical reasoning? Perhaps not. Premack (1983) and Thompson and Oden (2000) have suggested that both humans and apes can appreciate *higher-order* stimulus relations. Comparative study thus becomes critical in deciding among these and other rival hypotheses and in elucidating the evolutionary origins of analogical thinking.

Such comparative study is already under way. Cook and Wasserman (2007) and Fagot et al. (2001) have reported that pigeons and baboons, respectively, can discriminate *second-*

order same-different relations in a matching-to-sample task; that discrimination also transfers to novel stimuli. So, the hypothesis that only humans and apes can learn about *higher-order* same-different relations seems to be too restricted (Premack 1983; Thompson & Oden 2000).

Note that the pigeons and baboons in these two studies were not experimentally naive; each animal had some form of *first-order* same-different discrimination training prior to *second-order* relational matching-to-sample training. Perhaps previously learning to attend to simpler *first-order* relations is a prerequisite to pigeons' and baboons' subsequent success in discriminating more difficult *higher-order* relations.

Also note that, in both *first-order* and *second-order* relational discriminations, pigeons and baboons appeared to rely heavily on the variability of the items in a sample array (Wasserman et al. 2004). Young and Wasserman (1997) hypothesized that variability or entropy discrimination lies at the root of same-different discrimination (for a mechanistic theory of entropy discrimination, see Young et al. 2007). Two or more identical items always involve *zero* entropy; two or more nonidentical items always involve *nonzero* entropy, with entropy increasing as a direct function of the number of nonidentical items. Not only do pigeons (Young & Wasserman 1997; Young et al. 1997) and baboons (Wasserman et al. 2001a; 2001b) show strong sensitivity to the variability of the items in an array, but so also do college students – in both choice accuracy and choice reaction time (Castro et al. 2006; Young & Wasserman 2001; 2002).

Hence, we see that the interpretive framework proposed by Leech et al. may be effectively applied to both the development and evolution of relational discrimination behavior. Assuming that complex cognitive processes depend on and progressively evolve from more basic perceptual and cognitive processes – the notion of anagenesis (Gottlieb 1984) – we can make the following predictions. First, some species might only be sensitive to the variability of arrays of stimuli. Second, other, more select species may also be able to categorize stimulus arrays into those comprising identical or nonidentical items. Third, other, even more select species may be able to discriminate identical from nonidentical relations. Finally, we might expect to see this task ordering preserved in the development of relational responding in individual animals and to find this ordering to strongly depend on the animal's environmental experiences.

There's no "flying into flying" according to this proposal for the development and evolution of relational responding!

Authors' Response

Growing cognition from recycled parts

doi:10.1017/S0140525X08004718

Robert Leech, Denis Mareschal, and Richard P. Cooper
School of Psychology, Birkbeck, University of London, London WC1E 7HX,
United Kingdom.

r.leech@bbk.ac.uk

<http://www.alphalab.psyc.bbk.ac.uk/RobertLeech.html>

d.mareschal@bbk.ac.uk

<http://www.bbk.ac.uk/psyc/staff/academic/dmareschal>

r.cooper@bbk.ac.uk

<http://www.bbk.ac.uk/psyc/staff/academic/rcooper>

Abstract: In this response, we reiterate the importance of development (both ontogenetic and phylogenetic) in the understanding

of a complex cognitive skill – analogical reasoning. Four key questions structure the response: Does relational priming exist, and is it sufficient for analogy? What do we mean by relations as transformations? Could all or any relations be represented as transformations? And what about the challenge of more complex analogies? In addressing these questions we bring together a number of supportive commentaries, strengthening our emergentist case for analogy (in particular with insights from comparative psychology), and review new supportive evidence. We also rebut those commentaries that ignore development at their peril. Along the way, we revisit the main assumptions underlying the analogy as relational priming (ARP) account of analogy, clarifying and elaborating as necessary.

R1. Introduction

Analogy as relational priming (henceforth, ARP) views the analogical abilities of young children as developing out of priming – a simple, ubiquitous, cognitive mechanism. Central to the theory is the importance of development as a fundamental theoretical constraint for any account of analogical reasoning (and by extension, other high-level cognitive processes). It is our view that development must constrain all theories of analogy (even those just focusing on adult competence) because, to paraphrase Stephen Jay Gould quoting Francis Crick, “you can't identify the value in something until you know how it is made” (Gould 1992, p. 137). Consequently, we are happy to note that many commentators were highly sympathetic to this goal of our work (Cheshire, Ball, & Lewis [Cheshire et al.]; Goswami; Markman & Laux; Morrison & Cho; Schwering & Kühnberger; Wasserman). Indeed, several commentators provided interesting additional evidence supporting central tenets of our account (e.g., Estes & Jones; Raffray, Pickering, & Branigan [Raffray et al.]), and the converging support provided by the comparative angle of Greene and Wasserman on the evolution of high-level reasoning and relational knowledge is of particular interest. However, as made clear by other commentaries, much research into how we view analogical reasoning remains polarized (e.g., French, Holyoak & Hummel, Markman & Laux).

In this response, we begin by reconsidering why we believe development is so central to understanding analogical reasoning and why ignoring it (as some commentators do) inevitably distorts any theoretical account of analogy. We also reiterate our understanding of analogy as an emergent property, recycling existing cognitive processes to build up complex cognitive skills. We then consider the two main assumptions underlying the current implementation of the ARP account: relational priming and representing relations as transformations between semantic states, all the time paying particular attention to development. We then turn to issues arising from more difficult analogies that, we agree, require mechanisms or processes beyond those countenanced by ARP. We close by discussing other models of analogical completion and some meta-theoretical issues raised by some commentators.

R2. One doth not fly into flying

A key principle of the target article is the centrality of ontology to our account of analogy. While some commentators appear to have overlooked this issue (most notably

Holyoak & Hummel), others clearly appreciate the importance of an ontogenetic (and phylogenetic) perspective on cognition and the consequences of such a perspective for theories of analogy. **Wasserman** is perhaps the most articulate of the commentators in this respect, when he notes that, quoting Nietzsche (1891), “one doth not fly into flying.”

In our view, there are precursors to the kind of analogical reasoning and problem solving that is typically examined in experimental work with adult humans, just as there are precursors to flying. This is demonstrated by the developmental literature reviewed in the target article, which provides strong evidence that even very young children are able to apply analogical principles to solve simple proportional analogies, and that this ability improves with age as children move from basing their responses on surface similarity to relational similarity. As an account of the development of early analogical ability, ARP therefore provides more than just an account of a set of behavioral findings. It also provides an *explanation* of why behavior changes as children acquire greater knowledge of the world, and more generally of how analogical abilities relate to other cognitive mechanisms (i.e., priming).

Two points here require emphasis. First, in presenting this argument we do not mean to imply that ARP is intended as a precursor of a greater or more complex theory that at some point during adolescence replaces or supplants it. Rather, our view is that the developmental findings on which ARP is based provide fundamental insights into basic cognitive mechanisms that are preserved and refined as successive mechanisms (e.g., relating to inhibitory control and working memory) come into play.

The second point to note concerns the nature of explanation. There are many constraints that theorists might adopt in theory development. As an empirical science, psychological theories are generally constrained by or evaluated against empirical findings. But it is generally agreed that for a theory to have an explanatory role it must do more than merely match behavioral findings. It should not, for example, posit special purpose mechanisms or entities that have no support beyond the range of situations for which they were introduced. Equally, theoretical explanations are generally preferred if they are continuous with theories in related areas, accounting for behavior in one area using existing, well-supported concepts from other areas. ARP offers just such a theory of the analogical abilities of young children.¹

At the same time, if a theory cannot reach a particular level of cognitive complexity proposed as an explanation of adult competence, then this is not an explanation of the adult competence. Building a bridge on the basis of principles that cannot be implemented would never happen in engineering. Our approach therefore is extending (or in the case of Piaget et al. 1977, re-extending) to analogical reasoning the developmental principles of decades of work in other areas of developmental biology and psychology (see Mareschal et al. 2007).

It is important to note that this issue is independent of the nature versus nurture debate that rages though developmental psychology. Regardless of whether one believes that the essential components of cognition are somehow encoded throughout the genome, acquired through interaction with the environment, or some combination of the

two, the weight of proof remains on researchers to show how the adult state comes into being (Marcus 2004).

R3. Could analogy boil down to priming?

R3.1. Empirical evidence of relational priming

Several commentators critique ARP by questioning the existence of relational priming (e.g., **Morisson & Cho**), citing, in particular, the negative findings of Spellman et al. (2001). However, as argued by **Estes & Jones**, **Greene**, and **Raffray et al.**, there is in fact ample evidence of relational semantic priming in a broad range of domains (including memory retrieval, noun-noun comprehension, and syntactic text processing). Much of this evidence has appeared since Spellman et al.'s (2001) article (for a review, see Estes & Jones 2006). Furthermore, as Estes & Jones make clear in their commentary, the ARP model captures many aspects of the relational priming literature, while at the same time accounting for many aspects of the development of analogy. In addition, Raffray et al. point out recent evidence of a “lexical boost” in relational priming. This is consistent with the ARP account and further supports the concept of relational priming.

Despite this growing body of empirical findings, it is true that no studies have yet shown relational priming in children using verbal stimuli. **Goswami** suggests that behavioral evidence of relational priming in children is not critical, and that it would in fact be sufficient to provide evidence at the neural level of repetition suppression (Dehaene et al. 2001). We agree, but recent work by Duscherer et al. (in press) suggests that it is plausible to expect behavioral evidence. Duscherer et al. found, using a visual priming paradigm, that the prior perception of an action pantomime facilitates the recognition of the corresponding image of a tool associated with the action, but not of an image of a tool unassociated with the action, in 5- to 12-year-olds. This provides initial evidence that components of a causal relation can be primed in young children. Moreover, it is reasonable to expect verbal relational priming in younger age groups, since semantic priming effects are normally larger in younger, less fluent participants (Chapman et al. 1994).

Regardless of these empirical findings, there is currently something of a debate over what exactly relational priming consists of; however, we are agnostic with regards to many of the details of relational priming. Although we implement a specific type of relational priming in ARP (pattern completion in a dynamic recurrent network), more broadly, any kind of semantic priming would carry out the same effect. Indeed, when making an analogy, all sorts of different priming mechanisms could be at play – semantic priming, lexical priming in addition to “pure” relational priming – all or any of which would lead to the completion of the analogy through mechanisms that do not involve special architectures or processes such as explicit structural alignment.

R3.2. An associationist straitjacket?

Given the evidence for relational priming, is the mechanism sufficient for reasoning by analogy? **Holyoak & Hummel** argue that “the target article illustrates an all-

too-common approach to connectionist modeling of those cognitive processes most central to human intelligence: suck the essence out, then force-fit what's left into an associationist straitjacket." If by "sucking the essence out" they mean de-mystifying (i.e., explaining cognitive processes in terms of simpler processes such as memory that we have a better handle on), then yes, we absolutely agree with them. The aim of psychological science is to explain phenomena in terms of well-defined causal mechanisms whose effects are reliably measurable (see Smith [2000] for an extended discussion of this point). Priming in a semantic memory is an example of such a phenomenon.

Unfortunately, some of **Holyoak & Hummel's** criticisms seem to stem from a failure to grasp key elements of the model's functioning, leading them to characterize it as a "glorified paired associates task." So, to reiterate in case there is any confusion, the ARP model consists of a network learning about temporally contiguous states of the world. It is designed to learn about the causal semantics of the world and *not* to draw analogies. However, analogical completion falls naturally out of its dynamic pattern completion properties. To assess analogical completion, the *a* and *b* terms of an analogy (e.g., *bread* and *cut bread*) are first presented to the network. These values are clamped until the network has settled. This results in an internal state of the network that is consistent with its previous experience of *bread* and *cut bread*. At this point, the *a* and *b* inputs are unclamped. The network is then presented with the *c* term of the analogy (e.g., *lemon*) and allowed to settle again. Because it was in a state consistent with *cut* and it now receives a new input *lemon*, it settles into the output state of *cut lemon*. Note that (contrary to Holyoak & Hummel's claim) "knife" is *never* presented and thus *never* clamped during the formation of the analogy. It follows that the network's ability to complete this analogy is not simply an example of "glorified paired associates" in which the presence of *bread* paired with *knife* leads to *cut bread*. The model's ability to complete analogies comes instead from the recurrent network's dynamic pattern completion abilities. Different prior events (i.e., different *a:b* probes) will situate the network in different locations within a continuous state space. Its response to the current input (e.g., *lemon*) is a function of that new input and its location within that state space.

R3.3. Systematicity

A second concern over the sufficiency of relational priming for reasoning by analogy concerns *systematicity*. **Holyoak & Hummel**, **Schwering & Kühnberger**, and **Markman & Laux** all flag the phenomenon of systematicity as a selection constraint in analogical reasoning that, these commentators claim, the ARP model fails to capture. We fully agree that a bias towards systems of relations is a desirable feature of any model of analogical reasoning; but, furthermore, we also believe that a bias towards systematicity is a natural feature of many artificial neural networks with hidden layers and nonlinear activation rules, that arises without the need to use more complex structural representations. The point (made in the target article) is that systems of relations will vary coherently across multiple domains reflecting an underlying statistical structure, whereas superfluous or irrelevant relations will

often be restricted to individual domains. As such, the hidden units in a network will extract this underlying statistical structure across multiple domains. Relations that frequently co-occur will be represented closer together in the trained network's internal semantic space. Subsequent complex, multi-relation analogies using these relations will key into this proximity between representations of relations constituting a system, ensuring that the network has a bias towards picking systematic relations. We refer the reader to similar phenomena involving semantic cognition (Rogers & McClelland 2004) and metaphor (Thomas & Mareschal 2001).

R3.4. Analogical inference

Markman & Laux make the valid point that one of the main functions of analogical reasoning is to draw useful inferences that extend our knowledge. They also point out that ARP at present does not model how people draw these inferences. Although we have to date not focused on analogical inference, the possibility of drawing novel inferences is consistent with the ARP account. Under ARP we envisage analogical inferences as proceeding in a somewhat different way from more traditional mapping accounts (e.g., the Structure Mapping Engine) that explicitly choose features and relations to map from a base to a target domain. Instead, we suggest that analogical inferences might best be understood as novel generalizations governed by the distributional information about which input features and relations co-vary across the base and target domains. (The Metaphor as Pattern Completion account provides a template as to how this might proceed; see Thomas & Mareschal 2001.)

R4. Are relations simply transformations in a semantic space?

R4.1. Reversible relations

Several commentators (**Costello, Petrov, Sloutsky**) point out that although some relations (particularly simple causal relations such as *cutting*) could usefully be viewed as additive transformations in a semantic space, this cannot be the case for many of the relations habitually used in analogies, particularly adult analogies. In fact, Costello clearly demonstrates the impossibility that a bidirectional relation (e.g., the relation *kicks* that can be reversed from *John kicks Mary* to *Mary kicks John*) could be represented as an additive transformation. However, this misapprehension arises not from an inherent limitation of our account but from an overly narrow reading of the notion of relations as transformations. To clarify, we may distinguish between a *specific* and a *general* theory of relations as transformation. The *specific* theory is that there is a one-to-one (or close to one-to-one) correspondence between a given relation and a given transformation and, consequently, that relations are additive transformations in semantic space. In our first simulation, we followed this approach because it makes sense in the context of simple causal relations of the kind easily used in analogies by young children. For a relation such as *cutting* it is plausible that *cutting* transforms everyday objects in somewhat similar ways. Furthermore, these simple causal relations tend to be

unidirectional in that we see apples being transformed into cut apples but not the other way around. However, as **Goswami** points out, we need not be wedded to this specific theory of relations as additive transformations – instead what is important to the ARP account is a much more general theory of relations as transformations.

The *general* account of relations as transformations is that relations are transformations from one state to another in semantic/conceptual space, but there is no need for equivalence between a given relation and a specific additive transformation. This allows the network to represent relations that are reversible. Indeed, the second simulation considering the Gulf War/World War II analogies (sect. 4 of the target article) is implemented using this more general form of relation as transformations. What is important is that a given exemplar of a relation is a transformation between states (i.e., a mapping from one state to another) and the network learns different mappings (i.e., transformations) for different exemplars of the same relations. This approach, as we point out in the target article, is not unique to us but similar to what is envisaged in other transformational accounts of cognitive abilities (e.g., Rogers & McClelland 2004).

In general, therefore, we contend that relations should be understood as transformations in semantic space (a conceptual system), but, importantly, that a given relation is not necessarily represented by a specific single transformation. Instead, a given relation covers a set of learned transformations that are triggered by the specific interaction of objects with the available context.

Within the model, relations are represented implicitly in the matrix of connection weights, so there is nothing in the model per se that corresponds directly to the relation *cutting*. These implicit representations lead **Dietrich** to question whether the ARP model could account for explicit (i.e., conscious) awareness of relational knowledge. We have two responses to this point: first, (as made clear by the animal literature referred to by **Greene** and **Wasserman**; see also Oden et al. [2001] for a discussion of analogy in nonhuman species) there is no necessity for the relations underlying analogies to be explicitly known. Humans and some animal species appear to be capable of solving analogies using implicit or close-to-implicit knowledge and without much conscious awareness. Second, we accept that more complex, adult-like analogies will make considerable use of explicit strategies (see sect. R5 of this response). Note finally that by associating verbal labels with relations, for example, by augmenting the context layer of the network (see sect. 3.3.5 of the target article), it is easy to see how the implicit representations of relations could be linked to explicit (or at least verbalizable) representations in a manner that makes the knowledge consciously accessible.

As the previous point makes clear, we are not committed to any particular type of context representation driving analogies, such as causal agents that correspond to physical objects in the world (see **Sloutsky**). We introduced causal agents in the first simulations because we were modeling a specific task involving causal agents. However, as pointed out in the second simulation on more complex analogies, the context could take the form of many different types of mental representation, for example, something as concrete as a knife but also something more abstract such as heat or a verbal label.

R4.2. Multiple possible relations

A second objection to relations as transformations, raised most directly by **Akman** but also underlying **Bouissac's** uneasiness with ARP and central to one of **French's** concerns, is that multiple relations may hold between two objects. Consider Akman's example: *puppy:dog::kitten:?*. Is the relation of concern *offspring*, or is it any one of a number of other possibilities (*is_younger_than*, *is_cuter_than*)? Bouissac argues that the multiplicity of relations that might hold between the given terms in a proportional analogy undermines any possibility that there might be any single "correct" answer. His specific concern is with the socializing process that leads our society to prioritize one relation over another (and hence to suggest that there is just one right answer). In fact, this is wholly consistent with the ARP approach: when multiple relations hold between the *a* and *b* terms of a proportional analogy, the one that the model predicts will be applied is the one that is favored by the model's training history and current context. In other words, the model does not predict that there is a single "correct" answer to an analogy. As Bouissac notes, correctness is a function of the child's/system's knowledge, which in turn is a function of culture.

French uses the *puppy:dog* semi-analogy to make a related point. His argument is that the relation between the *a* and *b* terms that is appropriate within a proportional analogy depends on the *c* term. If the *c* term is *kitten* then *offspring* is the appropriate relation, but if the *c* term is *watch* then the appropriate relation is something else, maybe *small_version_of*. In fact, there is substantial semantic overlap between *offspring* and *small_version_of*. We therefore see no difficulty in accounting for this analogy within ARP, provided of course that the system is exposed to exemplar watches and clocks. A more critical example would be a case where multiple very different relations hold between the *a* and *b* terms of the analogy and when, given *c*₁, the appropriate response is *d*₁, but when, given *c*₂, an alternate response *d*₂ is appropriate. It remains to be demonstrated that young children can respond appropriately in such a situation. ARP would predict that, prior to the mastery of executive control, they could not. That is, within ARP the successful completion of such analogies is likely to require the kind of additional machinery for inhibiting responses included in our model of the Gulf War/World War II analogy, as described in the target article and further discussed in section R5 here.

R4.3. Abstract relations

A final issue concerning relations as transformations that several commentators brought up (**Dietrich**, **Petrov**, **Phillips**, **Schwering** & **Kühnberger**) is how the model can handle abstract relations. Again, in part this issue arises because of the confusion between the specific and general theories of relations as transformations. The general account of relations as transformations allows for abstract relations to be represented. For example, "pen *left_of* pencil," although not involving a meaningful transformation in the world, could be implemented as a mapping between the representation of pen and the representation of pencil. However, we acknowledge that there is still the important question of how relations

implemented implicitly as transformations can be used abstractly and flexibly in novel contexts. For example, in the situation “pen *left_of* pencil” this is a temporary configuration that we may never have seen before (or at least never analogized about before). As Phillips makes clear, we have not had the chance to be exposed to hundreds of epochs of learning about the relative location of pens and pencils; and even if we had, it would not be any use if we then saw a pencil to the left of a pen. This is a legitimate concern, and again we have two responses to this issue. First, we maintain that some of the complex and flexible usage of relations observed in the world relies on complex cognitive control strategies involving working memory and inhibition. Second, and perhaps more important, we contend, however, that there really are no examples of completely abstract and novel analogical or relational reasoning. Although we may never have seen a given pen and pencil in a specific and completely temporary configuration (or even two novel objects, e.g., greebles), we have seen and thought about a vast array of other objects in similar types of configurations.

R4.4. Empirical evidence for relations as transformations

As discussed in the target article, Thomas and Mareschal (1997) and Hahn et al. (2003) have both stressed the importance of transformations for judging similarity and, by extension, the importance of transformations for analogical reasoning. In addition, as discussed in section 5.1.2 of the target article, Leech et al. (2007) directly tested a strong prediction of the current implementation of causal relations as transformation. They found that the size of a transformation instantiating a relation between two states of the world (e.g., shrinks) was a performance factor affecting analogical reasoning, in both adults and adolescents, using a range of different analogical reasoning tasks. Such similarity effects fall naturally out of the implicit distributed nature of relations as transformation and are difficult to interpret in terms of explicit propositional representations of relations.

R5. ARP and the bigger picture

R5.1. The challenge of complex analogies

Researchers of adult analogy often make use of very complex analogies requiring far more than just the mechanisms that they propose in order to explain analogy. Take, for example, the analogy described by French in the beginning of his commentary relating baseball to academia. It is (on the surface at least) far more complex than $a:b::c:d$. But what has increased its complexity? Such analogies require additional components, such as advanced semantic knowledge of baseball and the world of academia, as well as sentence and text processing skills, working memory, and pragmatic conversational skills. Are these intended to be central to analogy? It seems entirely plausible that children's failure on these more involved analogies is at least in part due to shortcomings in world knowledge, working memory (see sect. R5.2), sentence processing, and pragmatic skills, all of which develop substantially throughout childhood.

We use proportional analogies in presenting and exploring ARP because they embody what is essential to analogies (Sternberg 1977a). Could French's example be recast as a proportional analogy? Does “Ty Cobb:Home runs::Academic:Journal articles” make sense? Absolutely, but only if the relevant background knowledge is primed (i.e., if the reader is located in the appropriate part of their semantic space) – as French does, by providing the reader with the background context so as to highlight the intended relation between Ty Cobb and home runs that French wishes us to transfer to the domain of academia.

R5.2. Working memory and cognitive control

We agree with Dumas & Richland that there is more to solving multirelational analogies than relational knowledge. As Morrison & Cho point out, there are clear constraints from working memory and changes in inhibitory control on children's developing abilities to use and comprehend increasingly complex analogies. Again, this is entirely consistent with the role of executive control that we posit for interpreting complex multi-relational analogies. We chose, however, to focus on the kind of analogies that very young children (3 to 4 years old and younger) seem competent at, and we produced an emergentist account of how these types of analogy could emerge. The next, and by no means easy, step is to investigate the interplay of the emergence of analogy based on relational priming (with a central focus on knowledge accretion) with the other emerging abilities of working memory and executive or cognitive control (e.g., active and selective maintenance and updating of working memory) that presumably underlie very complex analogical reasoning. This research agenda will have to address some of the issues regarding working memory that Morrison & Cho note (i.e., integration of multiple relations and the suppression of distracting stimuli). However, from our perspective it seems likely that analogy is parasitic on the rest of cognition. So, as Petrov points out, quoting French (2002, p. 204), “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.” The implication of this statement is that to model the development of analogical reasoning, we must model the development of a huge raft of interrelated cognitive skills (semantic cognition, working memory, inhibition, number, mathematical and spatial reasoning skills, etc.). As discussed in section R6 of this response, these are all capabilities pre-assumed by other existing models of analogical development.

R5.3. Reversibility and role assignment

A further point raised by multiple commentators is that people are able to reason about analogies with the terms reversed; as Akman says, “one apparent objection has to do with the ‘direction’ of the relationship.” So Holyoak & Hummel ask whether ARP, trained on *apple:sliced apple* and *bread:sliced bread*, could correctly complete the analogy “*sliced apple:apple::sliced bread:?*”. Estes & Jones similarly suggest that ARP could usefully be extended to have some mechanism for resolving roles, that is, discovering, for example, whether relational terms are in the order cause–effect, or effect–cause. The strength of this

criticism, as pointed out by Akman, depends upon whether young children can process these types of analogies. If, following Akman's intuition, such analogies are harder for young children when they appear in a non-canonical form (e.g., *sliced apple:apple::sliced bread:bread*), what additional processes are occurring that make them harder? Tentatively we suggest that in the extended account of analogical reasoning (i.e., ARP embedded within additional cognitive control systems), the analogizer may learn to funnel information that is noted to be in a non-canonical order into the appropriate order for relational priming to work.

R5.4. ARP as part of a larger theory

We are very comfortable with **Markman & Laux's** comment that "ARP theory does not explain full competence, and cannot, in principle, be extended to do so without it becoming a part of a larger theory." Where we differ from Markman & Laux is that we see ARP as a potential fundamental component of the wider theory, sufficient on its own for some analogies and necessary for many other more complex analogies. As illustrated by the second simulation regarding the Gulf War and World War II (sect. 4 of the target article), we show that, in principle, ARP could be embedded within other cognitive processes to construct the kind of complex extended analogies used to test models such as the Structure Mapping Engine (SME) and Learning and Inference with Schemas and Analogies (LISA). In essence, relational priming is iteratively applied in a controlled way to construct large multi-relation analogies. It may be helpful here to make an analogy with research into language development, which until relatively recently placed most weight on studying syntax as the fundamental aspect of human language (analogous to explicit structure mapping). It is, in fact, increasingly acknowledged that syntactic development and semantic development are fundamentally inseparable (Bates & Goodman 1997), especially when a truly developmental perspective is taken. Similarly, we believe that research on analogical development (and analogy in general) will only truly get anywhere if simple processes such as relational priming that are ontogenetically and phylogenetically (**Wasserman**) plausible can be put to good use in increasingly complex ways (i.e., with executive control building up to complex structure mapping).

R5.5. Meta-cognitive processes

Providing children with feedback on their analogical performance has been shown to improve subsequent analogical performance (at least in the short term), as has asking them to explain their responses (Cheshire et al. 2005; 2007). **Cheshire et al.** argue in their commentary that such effects cannot be accounted for within ARP as it stands, and question how (or indeed whether) ARP might be extended to allow behavior to be influenced by feedback and reflection. On the one hand, we see this as a criticism of a specific connectionist implementation of ARP, rather than a criticism of ARP itself. On the other hand, we are wary about arbitrarily extending the implementation to incorporate feedback and self-explanation when, as noted in our target article, analogical abilities develop in the absence of explicit training on an analogy. To elaborate,

Cheshire et al.'s data does not undermine either the mechanism of relational priming or the representation of relations as transformations. Rather, it suggests additional learning mechanisms that may influence performance on proportional analogies. Could the learning mechanism of our implementation be extended to account for the role of feedback and self-explanation? Possibly, but it is our view that the primary effects of feedback and self-explanation lie beyond the semantic system held to be involved in relational priming processes. They may, for example, be accounted for by learned attentional processes that differentially weight features of the stimuli, resulting in different inputs to the semantic system. Alternatively, they may affect post-semantic decision processes, resulting, for example, in the rejection (and subsequent inhibition) of targets that match only on surface similarity.

R6. Other models of analogy

Several commentators raise implicit objections to ARP by presenting alternative models of one or another aspect of analogy. **Borges, d'Avila Garcez, & Lamb (Borges et al.)** suggest that we have prematurely rejected hybrid "neural-symbolic" models, arguing that such models provide the great advantage of combining "trainable neural networks with symbolic representations," whereas **Schwering & Kühnberger** question the expressive power of ARP in relation to analogies in the physical sciences involving quantified variables and logic operations.

We accept the potential advantages of hybrid systems (i.e., combining learning and symbolic representations) in some circumstances. However, we dispute the relevance of these properties to our target phenomena. Thus, in presenting their position, **Borges et al.** describe two neural-symbolic systems, Connectionist Inductive Learning and Logic Programming (CILP) and Sequential Connectionist Temporal Logic (SCTL). The latter certainly appears to be capable of performing analogical reasoning of the type addressed by ARP, but **Borges et al.** make no attempt to relate the account to empirical phenomena, let alone developmental evidence. In particular, there is no evidence that the SCTL approach might replicate any of our seven observations on the development of analogy, beyond perhaps Observation 1 (concerning knowledge accretion). Even in that case, however, it is unclear how the learning mechanism of SCTL might relate to a child's learning of relational knowledge, which, we claim, is the product of witnessing many cases of a relation in a variety of contexts, and not learning from a set of propositional clauses. Furthermore, as noted by **Schwering & Kühnberger**, our non-symbolic (or more precisely, non-atomic) representation of relations has the advantage that it allows for semantically similar relations to have similar representations. Thus, our encoding is in principle able to account for analogies involving similar, but non-identical, relations (such as "occupies" and "annexes," to use **Schwering & Kühnberger's** example). At the same time ARP cannot be easily compared with the Heuristic Driven Theory Projection (HOTP) model (see Gust et al. 2007) cited by **Schwering & Kühnberger** – another model that is concerned solely with adult-level competence.

Doumas & Richland correctly point out that we overlooked one model that accounts for how explicit relational

representations can be learned. The Discovery of Relations by Analogy (DORA) model (Dumas et al. 2008) proposes to do just this. We apologise for this oversight and, therefore, provide due coverage of this model in the present reply. At first glance, we are highly sympathetic to the aims of this work. As argued throughout our target article, answering the question of how and from where an ability emerges is essential to any understanding of that ability. Unfortunately, a more detailed look at the DORA model suggests that it fails on this vital task. In fact, the very name of the model, “Discovery of Relations by Analogy,” suggests that there is a presupposition that the mechanisms for analogical retrieval and inference are in place prior to the discovery of any form of relational representations.

So, what are the assumptions that, according to Dumas et al. (2008), underlie the development of this ability in DORA? Fortunately, the authors are exceptionally clear in what they take as necessary requirements for the model to function. Thus, we read, “Armed with a basic vocabulary of perceptual and relational invariants ... DORA discovers relations through general learning processes and develops as a result of experience” (Dumas et al. 2008, p. 30). We have nothing to quibble about concerning this proposition. Next come the assumptions: “The model assumes that memory and perception are present at the start of learning. *It also requires a capacity for comparison, mapping, and SSL [Self Supervised Learning], as well as the ability to flexibly treat a feature as either a (holistic) feature of an object or as a feature of an explicit predicate*” (Dumas et al. 2008, p. 30; emphasis added). Even putting aside further mysteries such as how the complex machinery involving multiple modules, multiple levels of representation, and multiple processes and control streams got put together in the first place, most if not all of the skills believed by many to underlie analogy are pre-assumed in the DORA model. To us, this is not an explanation of how analogical abilities emerge, but rather an explanation of how, once analogy is in place, it can be used to generate new explicitly accessible relational representations.

Dumas & Richland make the further claim that DORA can account for all of the phenomena accounted for by ARP and much more. At the moment, there is no published evidence of this. As the DORA model is not a model of analogical development (nor does it claim to be), it does not try to account for the same tasks as we do. Although it may be able to capture some of the performance characteristics of the ARP model on *a:b::c:d* analogies, we would be curious to see how easily it could capture specific developmental phenomena such as the profile of errors that children are found to make at different ages (see sect. 1.1.3 of the target article; see also Rattermann & Gentner 1998a) and the dynamic temporal markers associated with development on analogical completion tasks (sect. 1.1.4; see also Hosenfeld et al. 1997).

There is one area in which the DORA and ARP models do overlap – they both suggest that there is often a shift in children’s performance from appearing to rely on surface similarity in drawing analogical inferences to relying on relational similarity. The two accounts differ in that ARP suggests that it should be possible to find examples in which the opposite shift occurs. It turns out that this is in fact the case. Opfer and Bulloch (in press) tested 3-, 4-,

and 5-year-olds on how they would generalize new information. On offspring problems, relational information yields more accurate generalizations, whereas in prey problems, surface similarity yields more accurate generalization. Opfer and Bulloch found that with offspring problems there was an increasing reliance on relational similarity with age (as previously reported), but on prey problems there was a decreasing reliance on relational similarity and an increasing reliance on surface similarity with age. These authors suggest that the relational trend commonly observed in analogical completion may, in fact, reflect increasing sensitivity to cue validity rather than an overall preference to generalize on the basis of relational similarity over perceptual similarity. These findings are entirely consistent with what is observed in the ARP model, and in fact validate a primary prediction made in section 5.1.1 of the target article. It is difficult to see how such a trend would fall naturally out of DORA or any of the associated structural alignment models.

A precursor to DORA coming from the same group of researchers is LISA, which **Morrison & Cho** claim can model many developmental findings. In particular, Richland et al. (2006) demonstrate that by increasing an inhibition parameter, LISA is able to account for effects of decreasing sensitivity to distracter items and increasing ability with complex relations with age. As Richland and colleagues argue, the functional effect of increasing inhibition within LISA is to allow increased working memory capacity, and, as argued earlier in our response, we agree that such developmental changes are likely to promote improved performance in some analogical tasks. However, as in the case of DORA, the LISA model is not really a developmental model – it assumes that complex computational machinery is in place, and all that “develops” is better inhibitory control. This is also the case with the efforts to model development using SME that **Markman & Laux** direct us to (e.g., Gentner et al. 1995). While it is a valid hypothesis to argue that performance improvements are caused by an increase in capacity (such as working memory), this leaves unanswered the question of where the specialist machinery comes from in the first place. This position is in stark contrast to ARP, which does not posit any special machinery for analogical completion and argues that such abilities simply fall out of existing general-purpose semantic memory systems.

R7. Wider issues

A number of commentators raise miscellaneous higher-level issues concerning metaphysical assumptions and the methodological approach embodied within ARP. Our responses to these meta-issues are collected here.

R7.1. Metaphysical assumptions

Linhares is concerned with some of the metaphysical assumptions underlying ARP, namely, the representational module assumption and the associated assumption of metaphysical realism. The central issue in both cases is the nature of mental representation. Thus, Linhares suggests that “mental objects” are dynamically built, consisting of a concept with associated sets of connotations, whereas “[a]nalogies are mechanisms by which a mental

object acquires connotations from different mental objects.” In contrast, ARP (or at least its connectionist implementation) assumes that mental representations can be adequately characterised as sets of features. Consider first the abstract claims of ARP. Although Linhares expresses no opinion about relational priming per se, he agrees with our view of relations as transformations between mental representations, and with the corollary that the application of such transformations during analogy effectively results in the acquisition of source features or connotations by a target mental object. What, though, are connotations? The word itself is slightly ethereal and mysterious. Linhares, however, provides some clues: they are (a) “rules for carving waves,” (b) “relations between mental objects,” and (c) “properties of particular objects.” The first of these appears to be a restatement of the representational module assumption that Linhares denies, while the third is entirely equivalent to the featural representations that we have assumed. The difference then between Linhares’ position and ARP reduces to the claim that mental objects consist, in addition to featural representations, of an atomic concept that may stand in some relation to other atomic concepts. Therefore, we see Linhares’ position as an elaboration of our position. That elaboration may well turn out to be justified when more complex cognitive abilities are considered, but the simulations reported in our target article demonstrate that it is not necessary for the target phenomena.

It is also important to be clear about the extent to which our representational assumptions shape the ARP model’s behavior. The ARP model would produce equivalent results regardless of the precise featural representations chosen, provided that appropriate abstract regularities over the domain hold (e.g., that the act of cutting affects some but not all features of an object that may be cut, but preserves all features of an object that cannot be cut). Thus, the ARP model requires that mental objects have a core (featural) representation that is invariant within an individual, but not, as a strong form of metaphysical realism would suggest, a representation that is invariant across individuals. Such core representations are presumably acquired throughout early development as the child becomes attuned to regularities within its environment, or as Linhares might argue, acquires “rules for carving waves.” We have no concerns in accepting this, given its widespread and apparently successful use throughout cognitive modelling (e.g., Elman et al. 1996; Mareschal & Thomas 2007; Rogers & McClelland 2004; Rumelhart & McClelland 1986; Shultz 2003).

R7.2. Computational-level theory and computational complexity of ARP

Two meta-theoretical issues concerning ARP are raised by Wareham, Rooij, & Müller (Wareham et al.). First, they are concerned with our failure to specify a computational-level theory of our target domain (in the sense of Marr 1982), and second, they are concerned with our failure to discuss the computational complexity of ARP. In our view the computational-level theory – the input-output function – is given by the empirical effects underlying our seven observations. These observations relate the inputs that the developing child is exposed to (both in terms of increasing world knowledge and the

specific analogy problems provided to children in the various developmental studies cited) and the outputs that the child produces (when presented with an analogy for completion). We therefore do not see this issue as a major concern, although it is perhaps behind several other commentaries that, as discussed above, object to our equating analogy with performance on the *a:b::c:d* task employed originally by Piaget et al. (1977) and later by Goswami and Brown (1989; 1990).

The computational complexity of ARP, on the other hand, is a more difficult question and Wareham et al. are right to raise it. We agree that there is an important role for computational complexity results to play in cognitive theorizing, particularly when considering issues of scaling. As Wareham et al. point out, Structure Mapping Theory (SMT) is computationally intractable, and this is, or should be, of concern to advocates of SMT. With regard to the computational complexity of ARP, it is again relevant to stress that our connectionist model, with its specific architecture, training procedure, and algorithm, is presented as an implementation of ARP that demonstrates the theory’s plausibility, but alternative approaches to the implementation of the theoretical claims of ARP (i.e., of relational priming and relations as transformations between representations) are surely possible. The question is therefore not one of the computational complexity of contrastive Hebbian learning, but whether there exists a computationally tractable implementation of ARP. This is an open question, but the inherently parallel nature of priming gives us reason to be cautiously optimistic about the scaling behavior of the theory.

R7.3. Relation to neurophysiology

An alternative meta-theoretical position is taken by Speed, who is concerned with the neurophysiological processes underlying analogical abilities – Marr’s (1982) implementation level – and our failure to provide a theory at this level. Specifically, Speed argues that analogy is an emergent property not of more basic *psychological* processes such as relational priming, but of more basic *physiological* processes. Unfortunately, she does not suggest what these physiological processes might be. Although we accept that psychological processes must eventually be grounded in physiological ones (see Mareschal et al. 2007), we find this reductionist perspective to be of limited use. To accept it would lead one to discard the entire field of cognitive psychology in favor of neurophysiology. While neurophysiology has told us much about the functioning of the brain, it has told us relatively little of the functioning of the mind.

We could of course have speculated more about the neural mechanisms underlying ARP (and the first of our models is at least broadly compatible with what is known of neural processing mechanisms; see O’Reilly 1996), but any comments would have been entirely speculative. Worse, in our view, is that they would conflate levels of explanation. We firmly believe that the neural level is important, but our model is concerned with the cognitive level. One may wish for well-grounded models of how “the brain, as a physical system, [can] take raw sensory information and perform higher-order cognition using that information (e.g., solving Raven’s Progressive Matrices),” as Speed writes, but the field is not yet in a position to

develop such theories. In any case, to do so would in our view be greatly facilitated by first developing a cognitive-level theory of how people solve Raven's Progressive Matrices, and this is the level at which ARP is conceived. In a similar vein before considerations of neuropsychological problems (e.g., autism or depression) can be mapped onto connectionist models (see the commentary by **Swain & Swain**) we must have a good understanding of the relations between the components of the model (e.g., hidden units, connections) and the brain (see also, Mareschal & Thomas 2007).

R8. Conclusion

Where, then, does this leave us? Young children, infants as young as 12 months old even, show signs of analogical ability (as do some chimpanzees). ARP accounts for such abilities, and for many of the improvements in analogizing seen in the first decade of the child's life. We accept, though, that we need to get to the other end – to adult-level competence. As argued throughout, our view is of analogy as an umbrella for a variety of concepts and processes – a view also explicitly endorsed by numerous commentators (**Borges et al., Sloutsky, Wasserman**) and denied by none. We accept that adult-level analogical competence involves processes beyond relational priming. We suspect, though (given what we know of the evolution and development of our cognitive abilities), that it does not involve an “analogy module,” as rival models seem to claim, but instead that it is firmly grounded in recycling old parts.

ACKNOWLEDGMENTS

We are grateful to European Commission grant NEST-029088 (ANALOGY) and the Economic and Social Research Council, United Kingdom.

NOTE

1. There are of course other requirements for a cognitive scientific explanation. Neural plausibility, for example, is one. This is addressed with respect to ARP in section R7.3.

References

Letters “a” and “r” appearing before authors' initials refer to target article and response, respectively.

- Aizawa, K. (2002) *The systematicity arguments, studies in brain and mind*. Kluwer Academic. [SP]
- Akman, V. (2007) Similar situations. In: *Context and appropriateness: Micro meets macro*, ed. A. Fetzer, pp. 31–54. John Benjamins. [VA]
- Altmann, G. T. M. (2002) Learning and development in neural networks: The importance of prior experience. *Cognition* 85:43–50. [aRL]
- Ambrose, S. E. (1998) *Citizen soldiers*. Simon & Schuster. [AL]
- American Psychological Association (2000) *Diagnostic and statistical manual of mental disorders: DSM-IV-TR*. American Psychological Association. [JES]
- Andrews, G. & Halford, G. S. (2002) A cognitive complexity metric applied to cognitive development. *Cognitive Psychology* 45:153–219. [LAAD, aRL]
- Andrews, G., Halford, G. S., Bunch, K. M., Bowden, D. & Jones, T. (2003) Theory of mind and relational complexity. *Child Development* 74(5):1476–99. [aRL]
- Arterberry, M. E. & Bornstein, M. H. (2001) Three-month-old infants' categorization of animals and vehicles based on static and dynamic attributes. *Journal of Experimental Child Psychology* 80:333–46. [UG]

- Barabas, H. (2000) Connections underlying the synthesis of cognition, memory, and emotion in primate prefrontal cortices. *Brain Research Bulletin* 52:319–30. [ASp]
- Barnden, J. A. (1994) On the connectionist implementation of analogy and working memory matching. In: *Advances in connectionist and neural computation theory, vol. 3: Analogy, metaphor and reminding*, ed. J. A. Barnden & K. J. Holyoak, pp. 327–74. Ablex. [aRL]
- Barnden, J. A. & Holyoak, K. J. (1994) *Advances in connectionist and neural computation theory. Analogy, reminding, and metaphor*. Ablex. [aRL]
- Barsalou, L. W., Simmons, W. K., Barbey, A. K. & Wilson, C. D. (2003) Grounding conceptual knowledge in modality-specific systems. *Trends in Cognitive Sciences* 7:84–91. [UG]
- Barwise, J. & Perry, J. (1983) *Situations and attitudes*. MIT Press. [VA]
- Bates, E. & Goodman, J. C. (1997) On the inseparability of grammar and the lexicon: Evidence from acquisition, aphasia and real-time processing. *Language and Cognitive Processes* 12:507–84. [rRL]
- Bates, E. & MacWhinney, B. (1989) Competition and connectionism. In: *The crosslinguistic study of sentence processing*, ed. B. MacWhinney & E. Bates. Cambridge University Press. [aRL]
- Becker, S., Moscovitch, M., Behrmann, M. & Joordens, S. (1997) Long-term semantic priming: A computational account and empirical evidence. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 23:1059–82. [aRL]
- Blanchette, I. & Dunbar, K. (2000) How analogies are generated: The roles of structural and superficial similarity. *Memory and Cognition* 28:108–24. [aRL]
- (2001) Analogy use in naturalistic settings: The influence of audience, emotion, and goals. *Memory and Cognition* 29:730–35. [aRL]
- Bowlde, B. & Gentner, D. (1997) Informativity and asymmetry in comparisons. *Cognitive Psychology* 34(3):244–86. [aRL]
- Branigan, H. P., Pickering, M. J. & Cleland, A. A. (2000) Syntactic coordination in dialogue. *Cognition* 75:B13–25. [CNR]
- Breton, A. (1924/1969) *Manifestos of surrealism*. Trans. R. Seaver & H. R. Lane. University of Michigan Press. [PB]
- Brown, A. L. (1989) Analogical learning and transfer: What develops? In: *Similarity and analogical reasoning*, ed. S. Vosniadou & A. Ortony, pp. 199–241. Cambridge University Press. [aRL]
- Brown, A. L. & Kane, M. J. (1988) Preschool children can learn to transfer: Learning to learn and learning from example. *Cognitive Psychology* 20:493–523. [aRL]
- Bruck, J. & Goodman, J. (1990) On the power of neural networks for solving hard problems. *Journal of Complexity* 6:129–35. [TW]
- Bunge, S. A., Wendelken, C., Badre, D. & Wagner, A. D. (2005) Analogical reasoning and prefrontal cortex: Evidence for separable retrieval and integration mechanisms. *Cerebral Cortex* 15:239–49. [RGM]
- Castro, L., Young, M. E. & Wasserman, E. A. (2006) Effects of number of items and visual display variability on same-different discrimination behavior. *Memory and Cognition* 34:1689–1703. [EAW]
- Cer, D. & O'Reilly, R. (2006) Neural mechanisms of binding in the hippocampus and neocortex: Insights from computational models. In: *Handbook of binding and memory: Perspectives from cognitive neuroscience*, ed. H. Zimmer, A. Mecklinger & U. Lindenberger, pp. 193–220. Oxford University Press. [AAP]
- Chalmers, D. (1990) Syntactic transformations on distributed representations. *Connection Science* 2(1):53–62. [RMF]
- Chalmers, D. J., French, R. M. & Hofstadter, D. R. (1992) High-level perception, representation, and analogy: A critique of artificial intelligence methodology. *Journal of Experimental and Theoretical Artificial Intelligence* 4:185–211. [RMF, AL]
- Chapman, L. J., Chapman, J. P., Curran, T. E. & Miller, M. B. (1994) Do children and the elderly show heightened semantic priming? *Developmental Review* 14:159–85. [aRL]
- Chen, Z., Sanchez, R. P. & Campbell, T. (1997) From beyond to within their grasp: The rudiments of analogical problem solving in 10- and 13-month-olds. *Developmental Psychology* 33:790–801. [aRL, VMS]
- Cheshire, A., Ball, L. J. & Lewis, C. N. (2005) Self-explanation, feedback and the development of analogical reasoning skills: Microgenetic evidence for a metacognitive processing account. In: *Proceedings of the Twenty-Second Annual Conference of the Cognitive Science Society*, ed. B. G. Bara, L. Barsalou & M. Bucciarelli, pp. 435–41. Erlbaum. [AC, rRL]
- Cheshire, A., Muldoon, K., Francis, B., Lewis, C. N. & Ball, L. J. (2007) Modelling change: New opportunities in the analysis of microgenetic data. *Infant and Child Development* 16:119–34. [AC, rRL]
- Chittka, L. & Osorio, D. (2007) Cognitive dimensions of predator responses to imperfect mimicry. *PLoS Biology* 5(12):e339. Available at: www.plosbiology.org. [PB]
- Cho, S., Holyoak, K. J. & Cannon, T. D. (2007a) Analogical reasoning in working memory: Resources shared among relational integration, interference resolution, and maintenance. *Memory and Cognition* 35:1445–55. [RGM]

- Cho, S., Moody, T. D., Cannon, T. D., Poldrack, R. A., Knowlton, B. J. & Holyoak, K. J. (2007b) Neural substrates for resolving interference from semantic association in verbal analogical reasoning. Paper presented at the Annual Meeting of the Society for Neuroscience, San Diego, California, November 3–7, 2007. [RGM]
- Cho, S., Moody, T. D., Poldrack, R. A., Cannon, T. D., Knowlton, B. J. & Holyoak, K. J. (2007c) Neural bases of analogical reasoning: Event-related fMRI yields evidence of prefrontal substrate for relational integration and interference resolution. Paper presented at the Annual Meeting of the Society for Cognitive Neuroscience, New York, May 5–8, 2007. [RGM]
- Christoff, K., Prabhakaran, V., Dorfman, J., Zhao, Z., Kroger, J. K., Holyoak, K. J. & Gabrieli, J. D. E. (2001) Rostrolateral prefrontal cortex involvement in relational integration during reasoning. *Neuroimage* 14:1136–49. [RGM]
- Chun, M. M. & Phelps, E. A. (1999) Memory deficits for implicit contextual information in amnesic subjects with hippocampal damage. *Nature Neuroscience* 2:844–47. [AJG]
- Cleland, A. A. & Pickering, M. J. (2003) The use of lexical and syntactic information in language production: Evidence from the priming of noun phrase structure. *Journal of Memory and Language* 49:214–30. [CNR]
- Clement, C. A. & Gentner, D. (1991) Systematicity as a selection constraint in analogical mapping. *Cognitive Science* 15:89–132. [aRL, ABM]
- Cohen, N. J., Poldrack, R. A. & Eichenbaum, H. (1997) Memory for items and memory for relations in the procedural/declarative memory framework. *Memory* 5:131–78. [AJG]
- Cook, R. G. & Wasserman, E. A. (2006) Relational learning in pigeons. In: *Comparative cognition: Experimental explorations of animal intelligence*, ed. E. A. Wasserman & T. R. Zentall, pp. 307–24. Oxford University Press. [EAW]
- (2007) Learning and transfer of relational matching-to-sample by pigeons. *Psychonomic Bulletin and Review* 14:1107–14. [EAW]
- Courchesne, E., Redcay, E. & Kennedy, D. P. (2004) The autistic brain: Birth through adulthood. *Current Opinion in Neurology* 17(4):489–96. [JES]
- Crisafi, M. A. & Brown, A. L. (1986) Analogical transfer in very young children: Combining two separately learned solutions to reach a goal. *Child Development* 57:953–68. [aRL]
- Cummins, R. (1995) Connectionism and the rationale constraint on cognitive explanation. *Philosophical Perspectives: AI, Connectionism and Philosophical Psychology* 9:105–25. [TW]
- (2000) “How does it work?” vs. “What are the laws?”: Two conceptions of psychological explanation. In: *Explanation and cognition*, ed. F. Keil & R. Wilson, pp. 117–44. MIT Press. [TW]
- Davelaar, E. J., Goshen-Gottstein, Y., Ashkenazi, A., Haarmann, H. J. & Usher, M. (2005) The demise of short-term memory revisited: Empirical and computational investigations of recency effects. *Psychological Review* 112:3–42. [aRL]
- Davidson, D. (1967) The logical form of action sentences. In: *The logic of decision and action*, ed. N. Rescher, pp. 81–95. University of Pittsburgh Press. [aRL]
- D’Avila Garcez, A., Broda, K. & Gabbay, D. (2002) *Neural-symbolic learning systems: Foundations and applications*. Springer-Verlag. [RVB]
- Davis, H. (1992) Transitive inference in rats (*Rattus norvegicus*). *Journal of Comparative Psychology* 106:342–49. [AAP]
- Dehaene, S., Naccache, L., Cohen, L., LeBihan, D., Mangin, J. F., Poline, J. B. & Rivière, D. (2001) Cerebral mechanisms of word masking and unconscious repetition priming. *Nature Neuroscience* 4:752–58. [UG, rRL]
- Delius, J. D. (1994) Comparative cognition of identity. In: *International perspectives on psychological science, vol. 1: Leading themes*, ed. P. Bertelson & P. Eelen, pp. 25–40. Erlbaum. [EAW]
- Dempster, F. N. & Brainerd, C. J. (1995) *Interference and inhibition in cognition*. Academic Press. [aRL]
- Devlin, K. (1991) *Logic and information*. Cambridge University Press. [VA]
- Diamond, A. (2002) Normal development of prefrontal cortex from birth to young adulthood: Cognitive functions, anatomy, and biochemistry. In: *Principles of frontal lobe function*, ed. D. T. Stuss & R. T. Knight, pp. 466–503. Oxford University Press. [RGM]
- Dietrich, E. (2000) Analogy and conceptual change, or you can’t step into the same mind twice. In: *Cognitive dynamics: Conceptual change in humans and machines*, ed. E. Dietrich & A. Markman, pp. 265–94. Erlbaum. [ED]
- Dietrich, E. & Markman, A. B. (2003) Discrete thoughts: Why cognition must use discrete representations. *Mind and Language* 18:95–119. [Asp]
- Dietrich, E., Markman, A. B. & Winkley, M. (2003) The prepared mind: The role of representational change in chance discovery. In: *Chance discovery by machines*, ed. Y. Ohsawa & P. McBurney, pp. 208–30. Springer-Verlag. [ED]
- Doumas, L. A. A. & Hummel, J. E. (2005a) A symbolic-connectionist model of relation discovery. In: *Proceedings of the Twenty-third Annual Conference of the Cognitive Science Society*, ed. B. G. Bara, L. Barsalou & M. Bucciarelli, pp. 606–11. LEA. [LAAD]
- (2005b) Modeling human mental representations: What works, what doesn’t, and why. In: *The Cambridge handbook of thinking and reasoning*, ed. K. J. Holyoak & R. G. Morrison, pp. 73–94. Cambridge University Press. [KJH]
- (2007) A computational account of the development of the generalization of shape information. In: *Proceedings of the Twenty-fifth Annual Conference of the Cognitive Science Society*, ed. D. McNamara & G. Trafton, pp. 221–26. LEA. [LAAD]
- Doumas, L. A. A., Hummel, J. E. & Sandhofer, C. M. (2008) A theory of the discovery and prediction of relational concepts. *Psychological Review* 115:1–43. [LAAD, KJH, rRL, RGM]
- Duscherer, K., Mounoud, P., Moy, G. & Perraudin, S. (in press) The influence of action perception on object recognition: A developmental study. *Developmental Science*. [rRL]
- Efron, B. & Tibshirani, R. J. (1998) *Introduction to the Bootstrap*. Chapman & Hall. [aRL]
- Elasmith, C. & Thagard, P. (2001) Integrating structure and meaning: A distributed model of analogical mapping. *Cognitive Science* 25:245–86. [aRL, Asp]
- Ellenbogen, J. M., Hu, P. T., Payne, J. D., Titone, D. & Walker, M. P. (2007) Human relational memory requires time and sleep. *Proceedings of the National Academy of Science USA* 104:7723–28. [AJG]
- Elman, J., Bates, E., Johnson, M. H., Karmiloff-Smith, A., Parisi, D. & Plunkett, K. (1996) *Rethinking innateness: A connectionist perspective on development*. MIT Press. [aRL]
- Elman, J. L. (1990) Finding structure in time. *Cognitive Science* 14(2):179–211. [RVB]
- (2005) Connectionist models of cognitive development: Where next? *Trends in Cognitive Sciences* 9(3):111–17. [UG]
- Elston, G. N. (2003) Cortex, cognition and the cell: New insights into the pyramidal neuron and prefrontal function. *Cerebral Cortex* 13:1124–38. [ASp]
- Estes, Z. (2003) Attributive and relational processes in nominal combination. *Journal of Memory and Language* 48:304–19. [ZE, aRL]
- Estes, Z. & Jones, L. L. (2006) Priming via relational similarity: A copper horse is faster when seen through a glass eye. *Journal of Memory and Language* 55(1):89–101. [VA, ZE, aRL]
- Fagot, J., Wasserman, E. A. & Young, M. E. (2001) Discriminating the relation between relations: The role of entropy in abstract conceptualization by baboons (*Papio papio*) and humans (*Homo sapiens*). *Journal of Experimental Psychology: Animal Behavior Processes* 27:316–28. [EAW]
- Falkenhainer, B., Forbus, K. D. & Gentner, D. (1989) The structure-mapping engine: Algorithm and examples. *Artificial Intelligence* 41(1):1–63. [VA, aRL, ABM]
- Fauconnier, G. & Turner, M. (2002) *The way we think: Conceptual blending and the mind’s hidden complexities*. Basic Books. [AL]
- Fenker, D. B., Waldmann, M. R. & Holyoak, K. J. (2005) Accessing causal relations in semantic memory. *Memory & Cognition* 33:1036–46. [ZE]
- Fodor, J. A. & Pylyshyn, Z. W. (1988) Connectionism and cognitive architecture: A critical analysis. *Cognition* 28:3–71. [SP]
- Forbus, K. D., Gentner, D. & Law, K. (1994) MAC/FAC: A model of similarity-based retrieval. *Cognitive Science* 19:141–205. [aRL]
- Frank, M., O’Reilly, R. & Curran, T. (2006) When memory fails, intuition reigns: Midazolam enhances implicit inference in humans. *Psychological Science* 17(8):700–707. [AAP]
- Frank, M., Rudy, J., Levy, W. & O’Reilly, R. (2005) When logic fails: Implicit transitive inference in humans. *Memory and Cognition* 33(4):742–50. [AAP]
- Frank, M., Rudy, J. & O’Reilly, R. (2003) Transitivity, flexibility, conjunctive representations, and the hippocampus: II. A computational analysis. *Hippocampus* 13:299–312. [AAP]
- Fraser, N. (1998) Neural network follies. Available at: <http://neil.fraser.name/writing/tank/>. [JES]
- Frawley, W. (1997) *Vygotsky and cognitive science: Language and the unification of the social and computational mind*. Harvard University Press. [AC]
- French, R. M. (1995) *The subtlety of sameness*. Bradford Books/MIT Press. [aRL, AAP]
- (2000) When coffee cups are like old elephants, or why representation modules don’t make sense. In: *Understanding representation in the cognitive sciences: Does representation need reality?* ed. A. Riegler, M. Peschl & A. von Stein, pp. 93–100. Springer. [AL]
- (2002) The computational modeling of analogy-making. *Trends in Cognitive Sciences* 6:200–205. [aRL, AAP, VMS]
- Freudenthal, D., Pine, J. M. & Gobet, F. (2006) Modelling the development of children’s use of optional infinitives in Dutch and English using MOSAIC. *Cognitive Science* 30:277–310. [UG]
- Gagné, C. L. (2001) Relation and lexical priming during the interpretation of noun–noun combinations. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 27:236–54. [ZE, aRL]
- (2002) Lexical and relational influences on the processing of novel compounds. *Brain and Language* 81:723–35. [aRL]
- Gagné, C. L. & Shoben, E. J. (1997) The influence of thematic relations on the comprehension of modifier–noun combinations. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 23:71–87. [aRL]

- Gagné, C. L., Spalding, T. L. & Ji, H. (2005) Re-examining evidence for the use of independent relational representations during conceptual combination. *Journal of Memory and Language* 53:445–55. [ZE, aRL]
- Garrod, S. C. & Anderson, A. (1987) Saying what you mean in dialogue: A study in conceptual and semantic co-ordination. *Cognition* 27:181–218. [CNR]
- Gelman, S. A. & Coley, J. D. (1990) The importance of knowing a dodo is a bird: Categories and inferences in 2-year-old children. *Developmental Psychology* 26:796–804. [UG]
- Gentner, D. (1983) Structure-mapping: A theoretical framework for analogy. *Cognitive Science* 7:155–70. [ED, RMF, aRL, AAP, TW]
- (1988) Metaphor as structure mapping: The relational shift. *Child Development* 59:47–59. [aRL]
- (1989) The mechanisms of analogical learning. In: *Similarity and analogical reasoning*, ed. S. Vosniadou & A. Ortony, pp. 199–241. Cambridge University Press. [aRL]
- Gentner, D. & Markman, A. B. (1997) Structure mapping in analogy and similarity. *American Psychologist* 52:45–56. [ASp]
- Gentner, D., Rattermann, M. J., Markman, A. B. & Kotovsky, L. (1995) Two forces in the development of relational similarity. In: *Developing cognitive competence: New approaches to process modeling*, ed. T. J. Simon & G. S. Halford, pp. 263–313. LEA. [arRL, ABM]
- Gentner, D. & Toupin, C. (1986) Systematicity and surface similarity in the development of analogy. *Cognitive Science* 10:277–300. [aRL, ABM]
- Gerrig, R. J. & Murphy, G. L. (1992) Contextual influences on the comprehension of complex concepts. *Language and Cognitive Processes* 7:205–30. [aRL]
- Gick, M. L. & Holyoak, K. J. (1983) Schema induction and analogical transfer. *Cognitive Psychology* 15:1–38. [aRL]
- Glimcher, P. W. (2003) *Decisions, uncertainty, and the brain: The science of neuroeconomics*. MIT Press. [PB]
- Gold, J. I. & Shadlen, M. N. (2007) The neural basis of decision making. *Annual Review of Neuroscience* 30:535–74. [PB]
- Goswami, U. (1991) Analogical reasoning: What develops? A review of research and theory. *Child Development* 62:1–22. [aRL]
- (1992) *Analogical reasoning in children*. Erlbaum. [aRL]
- (1995) Transitive relational mappings in three- and four-year-olds: The analogy of Goldilocks and the Three Bears. *Child Development* 66(3):877–92. [aRL]
- (1996) Analogical reasoning and cognitive development. *Advances in Child Development and Behavior* 26:91–138. [aRL]
- (1998) Is relational complexity a useful metric for cognitive development? *Behavioral and Brain Sciences* 21:838–39. [aRL]
- (2001) Analogical reasoning in children. In: *The analogical mind: Perspectives from cognitive science*, ed. K. Holyoak, D. Gentner, & B. Kokinov, pp. 437–70. MIT Press. [aRL, AAP, VMS]
- (2008) *Cognitive development: The learning brain*. Psychology Press. [UG]
- Goswami, U. & Brown, A. L. (1989) Melting chocolate and melting snowmen: Analogical reasoning and causal relations. *Cognition* 35:69–95. [FJC, arRL, ABM, AAP, CNR, VMS]
- (1990) Higher-order structure and relational reasoning: Contrasting analogical and thematic relations. *Cognition* 36:207–26. [arRL, AAP, VMS]
- Goswami, U., Leivers, H., Pressley, S. & Wheelwright, S. (1998) Causal reasoning about pairs of relations and analogical reasoning in young children. *British Journal of Developmental Psychology* 16:553–69. [LAAD, FJC]
- Goswami, U. & Pauen, S. (2005) The effects of a “family” analogy on class inclusion reasoning by young children. *Swiss Journal of Psychology* 64:115–24. [VMS]
- Gottlieb, C. (1984) Evolutionary trends and evolutionary origins: Relevance to theory in comparative psychology. *Psychological Review* 91:448–56. [EAW]
- Gould, S. J. (1992) *Bully for brontosaurus*. Penguin. [rRL]
- Greco, C., Hayne, H. & Rovee-Collier, C. (1990) Roles of function, reminding and variability in categorization by 3-month-old infants. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 16:617–33. [UG]
- Green, A., Fugelsang, J. & Dunbar, K. N. (2007) Automatic activation of categorical and abstract analogical relations in analogical reasoning. *Memory and Cognition* 34:1414–21. [aRL]
- Green, C. (2001) Scientific models, connectionist networks, and cognitive science. *Theory and Psychology* 11:97–117. [TW]
- Greene, A. J. (2007) Human hippocampal-dependent tasks: Is awareness necessary or sufficient? *Hippocampus* 17:429–33. [AJG]
- Greene, A. J., Gross, W. L., Elsinger, C. L. & Rao, S. M. (2006) An fMRI analysis of the human hippocampus: Inference, context and task awareness. *Journal of Cognitive Neuroscience* 18:1156–73. [AJG]
- (2007) Hippocampal differentiation without recognition: An fMRI analysis of the contextual cueing task. *Learning and Memory* 14:548–53. [AJG]
- Greene, A. J., Spellman, B. A., Dusek, J. A., Eichenbaum, H. B. & Levy, W. B. (2001) Relational learning with and without awareness: Transitive inference using nonverbal stimuli in humans. *Memory and Cognition* 29:893–902. [AJG]
- Gross, W. L. & Greene, A. J. (2007) Analogical inference: The role of awareness in abstract learning. *Memory* 15:838–44. [AJG]
- Grossberg, S. (1999) The link between brain learning, attention, and consciousness. *Consciousness and Cognition* 8:1–44. [ASp]
- Gust, H., Krumnack, U., Kühnberger, K.-U. & Schwering, A. (2007) Integrating analogical and inductive learning at different levels of generalization. In: *LNVD 2007: Learning from Non-Vectorial Data. Proceedings of the KI-2007 Workshop Held in Osnabrück, Germany, September 10th, 2007; Publications of the Institute of Cognitive Science*, vol. 6, ed. P. Geibel & B. J. Jain, pp. 46–57. Institute of Cognitive Science. [rRL, ASch]
- Gust, H., Kühnberger, K.-U. & Schmid, U. (2006) Metaphors and heuristic-driven theory projection (HDTP). *Theoretical Computer Science* 354(1):98–117. [ASch]
- Hahn, U., Chater, N. & Richardson, L. B. C. (2003) Similarity as transformation. *Cognition* 87:1–32. [arRL]
- Halford, G., Wilson, W., Guo, J., Gayler, R., Wiles, J. & Stewart, J. (1994) Connectionist implications for processing capacity limitations in analogies. In: *Advances in connectionist and neural computation theory*, vol. 2: *Analogical Connections*, ed. K. Holyoak & J. Barnden, pp. 363–415. Ablex. [aRL]
- Halford, G. S. (1993) *Children’s understanding: The development of mental models*. Erlbaum. [KJH, aRL]
- Halford, G. S., Bain, J. D., Maybery, M. T. & Andrews, G. (1998) Induction of relational schemas: Common processes in reasoning and complex learning. *Cognitive Psychology* 35(3):201–45. [SP]
- Halford, G. S., Wilson, W. H. & Phillips, W. (1998) Processing capacity defined by relational complexity: Implications for comparative, developmental, and cognitive psychology. *Behavioral and Brain Sciences* 21:803–31; discussion pp. 831–65. [aRL, AAP]
- Haykin, S. (1998) *Neural networks: A comprehensive foundation*. Prentice-Hall. [JES]
- Hinton, G. (1990) Mapping part-whole hierarchies into connectionist networks. *Artificial Intelligence* 46:47–75. [AAP]
- Hinton, G. & Sejnowski, T. (1986) Learning and relearning in Boltzmann machines. In: *Parallel distributed processing: Explorations in the microstructure of cognition. Vol. 1: Foundations*, ed. D. Rumelhart & J. McClelland, pp. 282–317. MIT Press. [AAP]
- Hofstadter, D. R. (1984) *The Copycat project: An experiment in nondeterminism and creative analogies*. MIT AI Memo, No. 755. [aRL, AAP]
- (2001) Epilogue: Analogy as the core of cognition. In: *The analogical mind: Perspectives from cognitive science*, ed. D. Gentner, K. J. Holyoak & B. N. Kokinov, pp. 499–538. MIT Press. [AL, ASp]
- Hofstadter, D. R. & the Fluid Analogies Research Group (1995) *Fluid concepts and creative; computer models of the fundamental mechanisms of thought analogies*. Basic Books. [AL]
- Holyoak, K., Gentner, D. & Kokinov, B. (2001) The place of analogy in cognition. In: *The analogical mind: Perspectives from cognitive science*, ed. K. Holyoak, D. Gentner & B. Kokinov, pp. 1–19. MIT Press. [aRL]
- Holyoak, K. J. & Barnden, J. A. (1994) *Advances in connectionist and neural computation theory, Analogical connections*. Ablex. [aRL]
- Holyoak, K. J. & Hummel, J. E. (2001) Toward an understanding of analogy within a biological symbol system. In: *The analogical mind: Perspectives from cognitive science*, ed. D. Gentner, K. J. Holyoak, & B. N. Kokinov, pp. 161–95. MIT Press. [aRL]
- Holyoak, K. J., Junn, E. N. & Billman, D. O. (1984) Development of analogical problem-solving skill. *Child Development* 55:2042–55. [LAAD, KJH, aRL]
- Holyoak, K. J. & Thagard, P. (1989) Analogical mapping by constraint satisfaction. *Cognitive Science* 13:295–355. [RMF, aRL]
- (1995) *Mental leaps: Analogy in creative thought*. MIT Press. [ASp]
- Hosenfeld, B., van der Maas, H. L. J. & van den Boom, D. C. (1997) Indicators of discontinuous change in the development of analogical reasoning. *Journal of Experimental Child Psychology* 64:367–95. [arRL]
- Houdé, O. (2000) Inhibition and cognitive development: Object, number, categorization, and reasoning. *Cognitive Development* 15:63–73. [aRL]
- Hummel, J. E. & Holyoak, K. J. (1997) Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review* 104(3): 427–66. [LAAD, RMF, KJH, aRL, ABM, RGM]
- (2003) A symbolic-connectionist theory of relational inference and generalization. *Psychological Review* 110:220–63. [LAAD, ZE, RGM, ASp]
- Inagaki, K. & Hatano, G. (1987) Young children’s spontaneous personifications as analogy. *Child Development* 58:1013–20. [aRL]
- Japkowicz, N. (2001) Supervised and unsupervised binary learning by feedforward neural networks. *Machine Learning* 42:97–122. [aRL]
- Judd, J. S. (1990) *Neural network design and the complexity of learning*. MIT Press. [TW]
- Karmiloff-Smith, A. (1998) Development itself is the key to understanding developmental disorders. *Trends in Cognitive Sciences* 2:389–98. [aRL]
- Keane, M., Ledgeway, T. & Duff, S. (1994) Constraints on analogical mapping: A comparison of three models. *Cognitive Science* 18:387–438. [aRL]

- Keane, M. T. & Brayshaw, M. (1988) The incremental analogy machine: A computational model of analogy. In: *Third European working session on learning*, ed. D. Sleeman, pp. 53–62. Morgan Kaufmann. [aRL]
- Kendler, T. S. (1995) *Levels of cognitive development*. Erlbaum. [SP]
- Kloos, H. & Sloutsky, V. M. (2008) What's behind different kinds of kinds: Effects of statistical density on learning and representation of categories. *Journal of Experimental Psychology: General* 137:52–72. [VMS]
- Kokinov, B. (1990) Associative memory-based reasoning: Some experimental results. In: *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society*, pp. 741–49. Erlbaum. [aRL, AAP]
- Kokinov, B. & French, R. M. (2003) Computational models of analogy-making. In: *Encyclopedia of cognitive science*, vol. 1, ed. L. Nadel, pp. 113–18. Nature Publishing Group. [aRL]
- Kokinov, B. & Petrov, A. (2001) Integration of memory and reasoning in analogy-making: The AMBR model. In: *The analogical mind: Perspectives from cognitive science*, ed. D. Gentner, K. Holyoak & B. Kokinov, pp. 59–124. MIT Press. [aRL, AAP]
- Kotovsky, L. & Gentner, D. (1996) Comparison and categorization in the development of relational similarity. *Child Development* 67:2797–822. [aRL, VMS]
- Kroger, J. K., Sabb, F. W., Fales, C. L., Bookheimer, S. Y., Cohen, M. S. & Holyoak, K. J. (2002) Recruitment of anterior dorsolateral prefrontal cortex in human reasoning: A parametric study of relational complexity. *Cerebral Cortex* 12:477–85. [RGM]
- Kutas, M. & Federmeier, K. D. (2000) Electrophysiology reveals semantic memory use in language comprehension. *Trends in Cognitive Science* 12: 63–70. [aRL]
- Lamb, L., Borges, R. V. & d'Avila Garcez, A. (2007) A connectionist cognitive model for temporal synchronisation and learning. In: *Proceedings of the Twenty-second AAAI Conference on Artificial Intelligence*, pp. 827–32. AAAI Press. [RVB]
- Larkey, L. B. & Love, B. C. (2003) CAB: Connectionist analogy builder. *Cognitive Science* 27:781–99. [aRL]
- (2007) Relations as transformations: Implications for analogical reasoning. *Quarterly Journal of Experimental Psychology* 60:897–908. [aRL]
- Leech, R., Mareschal, D. & Cooper, R. (2007) Relations as transformations: Implications for analogical reasoning. *Quarterly Journal of Experimental Psychology* 60:897–890. [aRL]
- Linhares, A. (2000) A glimpse at the metaphysics of Bongard problems. *Artificial Intelligence* 121(1–2):251–70. [AL]
- (2005) An active symbols theory of chess intuition. *Minds and Machines* 15(2):131–81. [AL]
- (submitted) Decision-making and strategic thinking through analogies. [AL]
- Linhares, A. & Brum, P. (2007) Understanding our understanding of strategic scenarios: What role do chunks play? *Cognitive Science* 31(6):989–1007. [AL]
- Loewenstein, J. & Gentner, D. (2005) Relational language and the development of relational mapping. *Cognitive Psychology* 50:315–53. [aRL, VMS]
- Marcus, G. F. (2004) *The birth of the mind: How a tiny number of genes creates the complexities of human thought*. Basic Books. [rRL]
- Mareschal, D., Johnson, M. H., Sirois, S., Spratling, M., Thomas, M. & Westermann, G. (2007) *Neuroconstructivism, vol. I: How the brain constructs cognition*. Oxford University Press. [aRL]
- Mareschal, D. & Thomas, M. S. C. (2007) Computational modeling in developmental psychology. *IEEE Transactions on Evolutionary Computation (Special Issue on Autonomous Mental Development)* 11(2):137–50. [rRL, TW]
- Markman, A. B. (1997) Constraints on analogical inference. *Cognitive Science* 21(4):373–418. [ABM]
- Markman, A. B. & Gentner, D. (1997) The effects of alignability on memory storage. *Psychological Science* 8(5):363–67. [aRL]
- Marr, D. (1982) *Vision: A computational investigation into the human representation and processing of visual information*. W. H. Freeman. [rRL, TW]
- Maybery, M., Bain, J. & Halford, G. (1986) Information processing demands of transitive inferences. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 12:600–13. [AAP]
- McClelland, J., McNaughton, B. & O'Reilly, R. (1995) Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review* 102:419–57. [AAP]
- McClelland, J. L. & Thompson, R. M. (2007) Using domain-general principles to explain children's causal reasoning abilities. *Developmental Science* 10:333–56. [aRL]
- McKoon, G. & Ratcliff, R. (1995) Conceptual combinations and relational contexts in free association and in priming in lexical decision and naming. *Psychonomic Bulletin & Review* 2:527–33. [aRL]
- Mesulam, M. M. (1998) From sensation to cognition. *Brain* 121:1013–52. [ASp]
- Mitchell, M. (1993) *Analogy-making as perception: A computer model*. Bradford Books/MIT Press. [RMF, aRL, AAP]
- Mitchell, M. & Hofstadter, D. R. (1990) The emergence of understanding in a computer model of concepts and analogy-making. *Physica D* 42:322–34. [RMF]
- Mogami, T. & Tanaka, K. (2006) Reward association affects neuronal responses to visual stimuli in macaque TE and perirhinal cortices. *Journal of Neuroscience* 26(25):6761–70. [ASp]
- Morrison, R. G. (2005) Thinking in working memory. In: *Cambridge handbook of thinking and reasoning*, ed. K. J. Holyoak & R. G. Morrison, pp. 457–73. Cambridge University Press. [RGM]
- Morrison, R. G., Doumas, L. A. A. & Richland, L. E. (2006) The development of analogical reasoning in children: A computational account. In: *Proceedings of the Twenty-eighth Annual Conference of the Cognitive Science Society*, ed. D. McNamara & G. Trafton, pp. 603–608. Erlbaum. [RGM]
- (2007) Inhibition vs. relational knowledge constraints in children's analogical reasoning: A symbolic-connectionist approach. Poster presented at the Fifth Biennial Meeting of the Cognitive Development Society, Santa Fe, New Mexico, October 26–27, 2007. [RGM]
- Morrison, R. G., Holyoak, K. J. & Truong, B. (2001) Working memory modularity in analogical reasoning. In: *Proceedings of the Twenty-fourth Annual Conference of the Cognitive Science Society*, ed. C. D. Schunn & W. Gray, pp. 663–68. Erlbaum. [RGM]
- Morrison, R. G., Krawczyk, D., Holyoak, K. J., Hummel, J. E., Chow, T., Miller, B. & Knowlton, B. J. (2004) A neurocomputational model of analogical reasoning and its breakdown in frontotemporal dementia. *Journal of Cognitive Neuroscience* 16:260–71. [RGM]
- Muldoon, K., Lewis, C. & Berridge, D. (2007) Predictors of early numeracy: Is there a place for mistakes when learning about number? *British Journal of Developmental Psychology* 25:543–58. [AC]
- Müller, B., Reinhardt, J. & Strickland, M. T. (1995) *Neural networks: An introduction*. Springer. [JES]
- Naccache, L. & Dehaene, S. (2001) The priming method: Imaging unconscious repetition priming reveals an abstract representation of number in the parietal lobes. *Cerebral Cortex* 11:966–74. [UG]
- Nietzsche, F. (1891) *Thus spake Zarathustra*. [rRL]
- Noesselt, N., Rieger, J. W., Schoenfeld, M. A., Kanowski, M., Hinrichs, H., Heinze, H.-J. & Driver, J. (2007) Audiovisual temporal correspondence modulates human multisensory superior temporal sulcus plus primary sensory cortices. *Journal of Neuroscience* 27:11431–41. [UG]
- Novick, L. R. (1988) Analogical transfer, problem similarity, and expertise. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 14:510–20. [aRL, ASp]
- Oden, D., Thompson, K. R. & Premack, D. (2001) Can an ape reason analogically? Comprehension and production of analogical problems by Sarah, a chimpanzee (*Pan troglodytes*). In: *The analogical mind: Perspectives from cognitive science*, ed. D. Gentner, K. Holyoak & B. Kokinov, pp. 471–97. MIT Press. [rRL, AAP]
- Opfer, J. & Bulloch, M. (in press) What makes relational reasoning smart? Revisiting the perceptual-to-relational shift in the development of generalization. *Developmental Science*. [rRL]
- O'Reilly, R. C. (1996) Biologically plausible error-driven learning using local activation differences. *Neural Computation* 8:895–938. [aRL]
- (2006) Biologically based computational Models of high-level cognition. *Science* 314:91–94. [aRL, AAP]
- O'Reilly, R. C. & Munakata, Y. (2000) *Computational explorations in cognitive neuroscience: Understanding the mind by simulating the brain*. MIT Press. [aRL]
- Pagonis, W. G. & Cruikshank, J. L. (1994) *Moving mountains*. Harvard Business School Press. [AL]
- Parberry, I. (1994) *Circuit complexity and neural networks*. MIT Press. [TW]
- Pauen, S. & Wilkening, F. (1997) Children's analogical reasoning about natural phenomena. *Journal of Experimental Child Psychology* 67:90–113. [aRL]
- Penn, D. C., Holyoak, K. J. & Povinelli, D. J. (2008) Darwin's mistake: Explaining the discontinuity between human and nonhuman minds. *Behavioral and Brain Sciences* 31(2):109–78. [PB, KJH, EAW]
- Phillips, S. (1999) Systematic minds, unsystematic models: Learning transfer in humans and networks. *Minds and Machines* 9(3):383–98. [SP]
- (2000) Constituent similarity and systematicity: The limits of first-order connectionism. *Connection Science* 12(1):45–63. [SP]
- (2007) Kenneth Aizawa, the systematicity arguments, studies in brain and mind. *Minds and Machines* 17(3):357–60. [SP]
- Piaget, J. (1970) *Genetic epistemology*, trans. E. Duckworth. Columbia University Press. [aRL]
- Piaget, J., Montangero, J. & Billeter, J. (1977) La formation des corrélats. In: *Recherches sur l'abstraction réfléchissante I*, ed. J. Piaget, pp. 115–29. Presses Universitaires de France. [aRL]
- Pickering, M. J. & Branigan, H. P. (1998) The representation of verbs: Evidence from syntactic persistence in written language production. *Journal of Memory and Language* 39:633–51. [CNR]

- Pickering, M. J. & Garrod, S. (2004) Towards a mechanistic theory of dialogue. *Behavioral and Brain Sciences* 27:169–226. [CNR]
- Premack, D. (1983) The codes of man and beasts. *Behavioral and Brain Sciences* 6:125–67. [EAW]
- Putnam, H. (1981) *Reason, truth, and history*. Cambridge University Press. [AL]
- Raffray, C. N., Pickering, M. J. & Branigan, H. P. (2007) Priming the interpretation of noun-noun combinations. *Journal of Memory and Language* 57:380–95. [ZE, CNR]
- Ramsay, J. O. & Silverman, B. W. (2002) *Applied functional data analysis: Methods and case studies*. Springer-Verlag. [aRL]
- Rattermann, M. J. & Gentner, D. (1998a) More evidence for a relational shift in the development of analogy: Children's performance on a causal-mapping task. *Cognitive Development* 13:453–78. [LAAD, arRL]
- (1998b) The effect of language on similarity: The use of relational labels improves young children's performance in a mapping task. In: *Advances in analogy research: Integration of theory and data from the cognitive, computational, and neural sciences*, ed. K. Holyoak, D. Gentner & B. Kokinov, pp. 274–82. New Bulgarian University. [aRL]
- Rattermann, M. J., Gentner, D. & DeLoache, J. (1990) The effects of familiar labels on young children's performance in an analogical mapping task. In: *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society*, pp. 22–29. LEA. [aRL]
- Reber, P. J., Knowlton, B. J. & Squire, L. R. (1996) Dissociable properties of memory systems: Differences in the flexibility of declarative and nondeclarative knowledge. *Behavioral Neuroscience* 110:861–71. [AJG]
- Recanati, F. (2004) *Literal meaning*. Cambridge University Press. [VA]
- Richland, L. E., Morrison, R. G. & Holyoak, K. J. (2006) Children's development of analogical reasoning: Insights from scene analogy problems. *Journal of Experimental Child Psychology* 94:249–73. [LAAD, arRL, RGM]
- Riecke, L., van Opstal, A. J., Goebel, R. & Formisano, E. (2007) Hearing illusory sounds in noise: Sensory-perceptual transformations in primary auditory cortex. *Journal of Neuroscience* 27:12684–89. [UG]
- Ripoll, T., Brude, T. & Coulon, D. (2003) Does analogical transfer involve a term-to-term alignment? *Memory and Cognition* 31:221–30. [AJG, aRL]
- Rogers, T. T. & McClelland, J. L. (2004) *Semantic cognition: A parallel distributed processing approach*. MIT Press. [aRL, AAP]
- (2005) A parallel distributed processing approach to semantic cognition: Applications to conceptual development. To appear in: *Building object categories in developmental time: Proceedings of the Carnegie Symposium on Cognition*, vol. 32, ed. D. Rakison and L. Gershkoff-Stowe, pp. 335–87. Erlbaum. [aRL]
- Rojas, R. & Feldman, J. (1996) *Neural networks: A systematic introduction*. Springer. [JES]
- Rosch, E. (1978) Principles of categorisation. In: *Cognition and categorisation*, ed. E. Rosch & B. B. Lloyd, pp. 27–48. Erlbaum. [UG]
- Ross, B. H. (1987) This is like that: The use of earlier problems and the separation of similarity effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 13:629–39. [ASp]
- Ross, B. H. & Kennedy, P. T. (1990) Generalizing from the use of earlier examples in problem solving. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 16:42–55. [ASp]
- Rougier, N. P., Noelle, D., Braver, T. S., Cohen, J. D. & O'Reilly, R. C. (2005) Prefrontal cortex and the flexibility of cognitive control: Rules without symbols. *Proceedings of the National Academy of Sciences* 102:7338–43. [aRL, AAP, ASp]
- Rumelhart, D. E. & McClelland, J. L. (1986) *Parallel distributed processing: Explorations in the microstructure of cognition foundations*, vol. 1. MIT Press/Bradford Books. [JES]
- Rutherford, E. (1911) The scattering of a and b particles by matter and the structure of the atom. *Philosophical Magazine*. Series 6, vol. 21:669–88. [ED]
- Ryan, J. D., Althoff, R. R., Whitlow, S. & Cohen, N. J. (2000) Amnesia is a deficit in relational memory. *Psychological Science* 11:454–61. [AJG]
- Sadock, B. J. & Sadock, V. A. (2004) *Kaplan and Sadock's comprehensive textbook of psychiatry*. Lippincott, Williams & Wilkins. [JES]
- Satpute, A. B., Fenker, D. B., Waldmann, M. R., Tabibnia, G., Holyoak, K. J. & Lieberman, M. D. (2005) An fMRI study of causal judgments. *European Journal of Neuroscience* 22:1233–238. [KJH]
- Schober, M. F. (1993) Spatial perspective-taking in conversation. *Cognition* 47:1–24. [CNR]
- Schultz, R. T. (2005) Developmental deficits in social perception in autism: The role of the amygdala and fusiform face area. *International Journal of Developmental Neuroscience* 23(2–3):125–41. [JES]
- Schunn, C. D. & Dunbar, K. (1996) Priming, analogy, and awareness in complex reasoning. *Memory and Cognition* 24:271–84. [AJG, aRL, AAP]
- Shastri, L. & Ajanagadde, V. (1993) From simple associations to systematic reasoning: A connectionist representation of rules, variables, and dynamic bindings using temporal synchrony. *Behavioral and Brain Sciences* 16:417–51. [aRL, RVB]
- Shultz, T. R. (2003) *Computational developmental psychology*. MIT Press. [arRL]
- Shultz, T. R. & Rivest, F. (2001) Knowledge-based cascade-correlation: Using knowledge to speed learning. *Connection Science* 13:43–72. [aRL]
- Shultz, T. R., Rivest, F., Egri, L. & Thivierge, J. P. (2006) Knowledge-based learning with KBCC. In: *Proceedings of the Fifth International Conference on Development and Learning*. [aRL]
- Siegler, R. S. (2006) Microgenetic analyses of learning. In: *Handbook of child psychology: Vol. 2, Cognition, perception, and language*, 6th edition, ed. D. Kuhn & R. S. Siegler, pp. 464–510. Wiley. [AC]
- Siegler, R. S. & Svetina, M. (2002) A microgenetic/cross-sectional study of matrix completion: Comparing short-term and long-term change. *Child Development* 73:793–809. [AC]
- Smith, B. C. (1996) *On the origin of objects*. MIT Press. [AL]
- Smith, L. B. (2000) Avoiding associations when it's behaviorism you really hate. In: *Becoming a word learner: A debate on lexical acquisition*, ed. R. M. Golinkoff, K. Hirsh-Pasek, L. Bloom, L. B. Smith, A. L. Woodward, N. Akhtar, M. Tomasello & G. Hollich. Oxford University Press. [rRL]
- Smolensky, P. & Legendre, G. (2006) *The harmonic mind: From neural computation to optimality-theoretic grammar*. MIT Press. [AAP]
- Spelke, E. S. (1994) Initial knowledge: Six suggestions. *Cognition* 50:431–45. [UG]
- Spellman, B. A. & Holyoak, K. J. (1992) If Saddam is Hitler then who is George Bush? Analogical mapping between systems of social roles. *Journal of Personality and Social Psychology* 62: 913–33. [RMF, KJH, aRL]
- Spellman, B. A., Holyoak, K. J. & Morrison, R. G. (2001) Analogical priming via semantic relations. *Memory and Cognition* 29:383–93. [ZE, KJH, arRL, RGM]
- Spivy, M. (2007) *The continuity of mind*. Oxford University Press. [ASp]
- Squire, L. R. (1992) Memory and the hippocampus: A synthesis from findings with rats, monkeys, and humans. *Psychological Review* 99:195–231. [AJG]
- St. John, M. & McClelland, J. (1990) Learning and applying contextual constraints in sentence comprehension. *Artificial Intelligence* 46:217–57. [AAP]
- Sternberg, R. J. (1977a) Component processes in analogical reasoning. *Psychological Review* 84:353–78. [arRL]
- (1977b) *Intelligence, information processing, and analogical reasoning: The componential analysis of human abilities*. Erlbaum. [RGM]
- Sternberg, R. J. & Nigro, G. (1980) Developmental patterns in the solution of verbal analogies. *Child Development* 51:27–38. [aRL]
- Sternberg, R. J. & Rifkin, B. (1979) The development of analogical reasoning processes. *Journal of Experimental Child Psychology* 27:195–232. [aRL]
- Tanaka, K. (1992) Inferotemporal cortex and higher visual functions. *Current Opinion in Neurobiology* 2:502–505. [ASp]
- Thagard, P. (2000) *Coherence in thought and action*. MIT Press. [TW]
- Thagard, P., Holyoak, K. J., Nelson, G. & Gochfeld, D. (1990) Analog retrieval by constraint satisfaction. *Artificial Intelligence* 46(3):259–310. [VA]
- Thelen, E. & Smith, L. B. (1994) *A dynamic systems approach to the development of cognition and action*. MIT Press. [aRL]
- Thomas, M. S. C. & Mareschal, D. (1997) Connectionism and psychological notions of similarity. In: *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society*, ed. M. G. Shafto & P. Langley, pp. 757–62. Erlbaum. [arRL]
- (2001) Metaphor as categorisation: A connectionist implementation. *Metaphor and Symbol* 16:5–27. [arRL]
- Thomas, M. S. C., Mareschal, D. & Hinds, A. (2001) A connectionist account of the emergence of the literal-metaphorical-anomalous distinction in young children. In: *Proceedings of the 23rd Annual Meeting of the Cognitive Science Society*, ed. J. D. Moore & K. Stenning, pp. 1042–47. LEA. [aRL]
- Thompson, R. K. R. & Oden, D. L. (2000) Categorical perception and conceptual judgments by nonhuman primates: The paleological monkey and the analogical ape. *Cognitive Science* 24:363–96. [EAW]
- Thompson, R. K. R., Oden, D. L. & Boysen, S. T. (1997) Language-naïve chimpanzees (*Pan troglodytes*) judge relations-between-relations in a conceptual matching task. *Journal of Experimental Psychology: Animal Behavior Processes* 23:31–43. [VMS]
- Towell, G. & Shavlik, J. (1994) Knowledge-based artificial neural networks. *Artificial Intelligence* 70(1–2):119–65. [RVB]
- Tsotsos, J. K. (1990) Analyzing vision at the complexity level. *Behavioral and Brain Sciences* 13(3):423–69. [TW]
- Tulving, E. & Schacter, D. L. (1990) Priming and human memory systems. *Science* 247:301–306. [aRL]
- Tunteler, E. & Resing, W. C. (2002) Spontaneous analogical transfer in 4-year-olds: A microgenetic study. *Journal of Experimental Child Psychology* 83:149–66. [aRL]
- Valiant, L. (2003) Three problems in computer science. *Journal of the ACM* 50(1):96–99. [RVB]
- van der Maas, H. & Molenaar, P. (1992) A catastrophe-theoretical approach to cognitive development. *Psychological Review* 99:395–417. [aRL]

- Van Elzakker, M., O'Reilly, R. & Rudy, J. (2003) Transitivity, flexibility, conjunctive representations, and the hippocampus: I. An empirical analysis. *Hippocampus* 13:292–98. [AAP]
- van Rooij, I. (2003) Tractable cognition: Complexity theory in cognitive psychology. Unpublished doctoral dissertation, Department of Psychology, University of Victoria, Canada. [TW]
- (in press) The tractable cognition thesis. *Cognitive Science*. [TW]
- van Rooij, I., Stege, S. & Kadlec, H. (2005) Sources of complexity in subset choice. *Journal of Mathematical Psychology* 49:160–87. [TW]
- van Rooij, I. & Wareham, T. (in press) Parameterized complexity in cognitive modeling: Foundations, applications and opportunities. *Computer Journal*. [DOI: 10.1093/comjnl/bxm034] [TW]
- Veale, T. & Keane, M. T. (1997) The competence of sub-optimal theories of structure mapping on hard analogies. In: *Proceedings of the 1997 International Joint Conference on Artificial Intelligence (IJCAI'97)*, vol. 1, pp. 232–37. Morgan Kaufmann. [TW]
- Viskontas, I. V., Morrison, R. G., Holyoak, K. J., Hummel, J. E. & Knowlton, B. J. (2004) Relational integration, inhibition, and analogical reasoning in older adults. *Psychology and Aging* 19:581–91. [RGM]
- von Baeyer, H. C. (2003) *Information: The new language of science*. Orion Books. [PB]
- von Fersen, L., Wynne, C., Delius, J. & Staddon, J. (1991) Transitive inference in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes* 17:334–41. [AAP]
- Vosniadou, S. (1989) Analogical reasoning as a mechanism in knowledge acquisition: A developmental perspective. In: *Similarity and analogical reasoning*, ed. S. Vosniadou & A. Ortony, pp. 199–241. Cambridge University Press. [aRL]
- Vosniadou, S. & Brewer, W. F. (1992) Mental models of the earth: A study of conceptual change in childhood. *Cognitive Psychology* 24(4):535–85. [ABM]
- Vygotsky, L. (1978) *Mind in society*. Harvard University Press. [UG]
- Waltz, J. A., Knowlton, B. J., Holyoak, K. J., Boone, K. B., Mishkin, F. S., de Menezes Santos, M., Thomas, C.R. & Miller, B.L. (1999) A system for relational reasoning in human prefrontal cortex. *Psychological Science* 10:119–25. [RGM]
- Waltz, J. A., Lau, A., Grewal, S. K. & Holyoak, K. J. (2000). The role of working memory in analogical mapping. *Memory and Cognition* 28:1205–12. [RGM]
- Wareham, T. (1999) Systematic parameterized complexity analysis in computational phonology. Unpublished doctoral dissertation, Department of Computer Science, University of Victoria, Canada. [TW]
- Wasserman, E. A., Fagot, J. & Young, M. E. (2001a) Same-different conceptualization by baboons (*Papio papio*): The role of entropy. *Journal of Comparative Psychology* 115:42–52. [EAW]
- Wasserman, E. A., Young, M. E. & Cook, R. G. (2004) Variability discrimination in humans and animals: Implications for adaptive action. *American Psychologist* 59:879–90. [EAW]
- Wasserman, E. A., Young, M. E. & Fagot, J. (2001b) Effects of number of items on the baboon's discrimination of same from different visual displays. *Animal Cognition* 4:163–70. [EAW]
- Wasserman, E. A. & Zentall, T. R. (2006) *Comparative cognition: Experimental explorations of animal intelligence*. Oxford University Press. [EAW]
- Wikipedia, Rutherford Scattering: http://en.wikipedia.org/wiki/Rutherford_scattering, and Gold Foil Experiment: http://en.wikipedia.org/wiki/Geiger-Marsden_experiment. [ED]
- Wisniewski, E. J. & Love, B. C. (1998) Relations versus properties in conceptual combination. *Journal of Memory and Language* 38:177–202. [aRL]
- Wright, S. B., Matlen, B. J., Baym, C. L., Ferrer, E. & Bunge, S. A. (in press) Neural correlates of fluid reasoning in children and adults. *Frontiers in Human Neuroscience*. [RGM]
- Young, M. E. & Wasserman, E. A. (1997) Entropy detection by pigeons: Response to mixed visual displays after same-different discrimination training. *Journal of Experimental Psychology: Animal Behavior Processes* 23:157–70. [EAW]
- (2001) Entropy and variability discrimination. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 27:278–93. [EAW]
- (2002) Detecting variety: What's so special about uniformity? *Journal of Experimental Psychology: General* 131:131–43. [EAW]
- Young, M. E., Wasserman, E. A. & Ellefson, M. R. (2007) A theory of variability discrimination: Finding differences. *Psychonomic Bulletin and Review* 14:805–22. [EAW]
- Young, M. E., Wasserman, E. A. & Garner, K. L. (1997) Effects of number of items on the pigeon's discrimination of same from different visual displays. *Journal of Experimental Psychology: Animal Behavior Processes* 23:491–501. [EAW]