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]
[a]*Department of Computer Science, Memorial University of Newfoundland, St. John's, NL, A1B 3X5, Canada;* [b]*Nijmegen Institute for Cognition and Information, Radboud University Nijmegen, 6525HR Nijmegen, The Netherlands;* [c]*Department 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-*