rationality illuminate some crucial problem in the theory of rational choice. With special attention to the role of models and simulations, he proposes some low-level models that combine the two. Danielson concludes with a survey of the normative significance of unifying evolutionary game theory and rationality research, plus some speculation about human rationality's evolution.

As a whole, the *Handbook* provides an engaging and accessible survey of current philosophical studies of rationality, and a useful up-to-date roadmap of state-of-the-art thinking on this enduring topic. Of course, the coverage of the *Handbook* might have ... ctive ration ... rves conce ... one might ... , or

Schlesinger, M. (2006).
Review of the book Computational Developmental Psychology.
Philosophical Psychology, *19*, 557-561.

postmodernity in Part 2. Overall, this is an authoritative sourcebook. Whereas each chapter can be read on its own and the references within each essay direct the reader to further writings in that topic area, the selection and arrangement of the topics are thoughtful and well-designed. Each author is undoubtedly among the best qualified candidates to undertake his or her commissioned chapter, and most essays are consistently of high quality.

JING ZHU
Department of Social Sciences
Graduate School of the Chinese Academy of Sciences
Beijing, 100049, China
Email: jingzhu@gscas.ac.cn

Computational Developmental Psychology

THOMAS R. SHULTZ
Cambridge, MA: MIT Press, 2003
322 pages, ISBN: 026219483x (hbk); $38.00

Since the publication of Rumelhart and McClelland's *Parallel Distributed Processing* in 1986, critics have debated whether neural network models of learning and development are—at best—video-game-like simulations of real organisms and their nervous systems, or—at worst—clever, hand-built demonstrations that bear only a superficial resemblance to the real-world behaviors that they are intended to represent. In response to this subtle but pervasive skepticism, Thomas Shultz has delivered a comprehensive text that skillfully illustrates the merging of computer models and developmental psychology into a new beast, part of an emerging, interdisciplinary field called "developmental science."

Shultz's *Computational Developmental Psychology* is organized to achieve three goals. First, it provides an intuitive (but also rigorous) overview of the mathematical foundations for artificial neural network models. At the beginning of chapter 2, Shultz notes that "I present all the essentials here (one-stop shopping) with sufficient background to enable understanding of key ideas by basically all readers, not just those with an extensive background in neural networks and mathematics" (p. 25; the formal mathematical details and derivations are systematically addressed in the appendix). This is certainly an accurate assessment, but I would also caution that a careful reading of the chapter, while not an absolute requirement, greatly enhances understanding of the models that Shultz presents in later chapters. Skimming, on the other hand, might be comparable to the experience of watching a French movie (without subtitles) after taking a year or two of high school French classes. You get the gist, but just barely.

Chapter 2 spans an impressive range of fundamental topics, including (a) the concept of linear separability, (b) network function in mathematical terms (i.e., propagation of activity between neural units, hereafter called "units"), (c) network architectures and topologies, and (d) learning algorithms. It is at this point that Shultz highlights a major theme of the book: although back-propagation-of-error (hereafter "backprop") is one of the most common methods for tuning the connection weights of a network (thereby improving its behavior), it also has a number of significant flaws. In particular, backprop nets learn slowly, have static topologies, are vulnerable to catastrophic interference (i.e., new learning wipes out old knowledge), scale up poorly to larger problems, and worst of all, are not biologically plausible.

Fortunately, each of these weaknesses is successfully addressed by an alternative approach, called "cascade-correlation." Properly speaking, cascade-correlation is a fusion of two distinct innovations, one network-architectural and the other learning-algorithmic. As Shultz carefully explains, the architectural twist is that cascade-correlation nets "grow" or expand, analogous to the processes of synaptogenesis and neurogenesis. As they recruit new internal units, cascade-correlation nets increase their ability to represent important features in their (putative sensory) input. The algorithmic twist, meanwhile, is that Shultz and his colleagues also employ a novel learning algorithm called "quick-prop," which learns faster than backprop.

Two nice touches are left for the end of chapter 2. First, Shultz offers the reader a brief summary of other related connectionist designs, including simple recurrent, encoder, auto-associator, and feature-mapping networks. Second, in a bit of deliberate irony, Shultz uses a traditional symbolic model to classify 18 different subsymbolic or connectionist models. The result is a decision tree that groups models as a function of common objectives and constraints (e.g., whether the model is *supervised*, i.e., trained with explicit feedback). The maneuver is a little like inviting your atheist friend over for dinner, and then asking them to pray. Shultz' humor is likely to be appreciated by readers who know the history of "classical" and "modern" AI. In any case, the decision tree (Table 2.2, p. 72) is a valuable aid to novice model-builders.

As a result of the focus on cascade-correlation, there are other notable computational approaches to modeling development that are either excluded or only briefly mentioned by Shultz, including adaptive resonance theory, reinforcement learning, genetic algorithms, and dynamic field theory (e.g., Schlesinger & Parisi, 2004). While none of these are conventional connectionist models, strictly speaking, they still share a common mathematical framework and many theoretical tenets with the connectionist approach (e.g., that knowledge representations are graded and distributed). In chapters 3–5, however, it becomes apparent that Shultz' decision to highlight the virtues of cascade-correlation models is guided by a compelling blend of careful reasoning and broad empirical support. As Shultz seems to implicitly argue, other modeling approaches are no less important or interesting, but simply less successful in capturing three fundamental aspects of development: knowledge representation, developmental transitions, and stages of development.

As Shultz works towards achieving his second goal—to systematically compare the performance of selected models on well-known developmental phenomena—several unique features of the book are revealed. First, the discussion of knowledge representation (i.e., encoding, storage, and retrieval) offers balanced coverage of two prevailing views. On the one hand, the rule-based view characterizes knowledge as symbolic, proposition-like bits of information that are stored in something analogous to a catalogue or lookup table. Connectionism, on the other hand, characterizes knowledge as distributed and graded, rather than symbolic, and represented by the strength of connections between a network of units that are always active to one degree or another.

Shultz surveys a wide array of developmental phenomena in chapter 3, helping the novice model-builder to appreciate both the conceptual and methodological differences that arise from implementing symbolic versus connectionist models. A number of well-known topics are covered, including both logico-mathematical cognition (e.g., number conservation and the balance-scale) and language acquisition (e.g., semantics and morphology). Developmental researchers who specialize in these areas will be more-than-pleasantly surprised to find that Shultz is far from a "bull in the china shop." Indeed, he goes well beyond a superficial overview of each topic to offer careful, detailed descriptions that are meticulous and thorough. For example, reviews of children's reasoning on Piaget's number conservation and balance-scale tasks include discussion of the (somewhat esoteric) set-size and torque-difference effects, respectively.

The contrast between rule-based and connectionist models also allows Shultz to exploit a second unique feature of the book: the "bakeoff." As the name suggests, a bakeoff is a head-to-head comparison of two models that are designed to solve the same task. Borrowed from the field of machine learning (e.g., the performance benchmark), the bakeoff is probably an unfamiliar research tool for most experimental psychologists. Nevertheless, there are two crucial lessons that the reader should not miss at this point (Shultz explicitly notes the first; he either implies, or I infer, the second). First, on a practical level, the bakeoff creates a level playing field, on which two alternative accounts can objectively and fairly compete.

Second, on a more symbolic level, the bakeoff also illustrates the need for computational models to be precise and specific in their formulation, in contrast to "traditional" verbal theories, which are more general, and sometimes vague, ambiguous, or difficult to formalize and test.

As Shultz is careful to note, the purpose of the bakeoff is not simply to figure out whether it is the connectionist or rule-based model that learns faster, or better than the other, but rather to perform a much more in-depth *comparative analysis* of the two models. In other words, a comparative analysis highlights the shape of the developmental trajectory, and includes questions such as: Is the process of learning in the model gradual or discontinuous? If discontinuous, are the changes analogous to stages of development in human children? If the model, like children, learns through a series of stages, when do those stage transitions occur, and what learning experiences precede or predict the emergence of a new stage?

Chapter 4, a *tour-de-force* in just under 50 pages, focuses on the issue of developmental transitions, i.e., the "motor" that drives the developmental process forward. To illustrate the topic, Shultz returns to the development of children's reasoning on the balance-scale and conservation tasks, this time using the bakeoff between symbolic and connectionist models to highlight how and when new behaviors emerge. For example, on the balance-scale task, young children watch as weights are placed on each side of a balance-scale (sometimes at different distances), and then are asked to predict which side (if either) will tip. An intriguing developmental pattern is that young children initially focus exclusively on the magnitude of the weights on each side. With experience, though, they begin to attend to, and incorporate into their reasoning, both the magnitudes and distances of the weights (from the center of the scale). Perhaps not surprisingly, Shultz showcases the virtues of cascade-correlation as the model best-suited for explaining (and sometimes predicting) these changes in behavior and thought.

Shultz not only addresses a number of empirical topics in chapter 4, but several fundamental conceptual issues are also systematically explored. First, both "classical" (i.e., Piagetian) and modern theoretical approaches to developmental transitions are reviewed. Second, Shultz refutes (or at least responds to) "Fodor's paradox," part of an analytical, nativist assault on connectionist models which argues that learning and development are a reorganization of prior (read: innate) knowledge, not an acquisition of something new (Fodor & Pylyshyn, 1988). The chapter concludes with two topics that also link psychology and philosophy via epistemology: the relation between learning and development (a distinction developmentalists constantly bicker about), and the relation between evolution and development (i.e., nativism, empiricism, and constructivism). While the latter issue is inevitably a minor theme for Shultz, it should be noted that other authors—in particular, Elman et al. (1996)—have tackled this question in greater detail.

In chapter 5, Shultz combines two lines of attack: the first half of the chapter presents an in-depth analysis of discontinuity in development (i.e., the emergence of qualitatively new behaviors), while the second half emphasizes the now-familiar blend of behavioral and modeling studies to illustrate the phenomenon of developmental stages and how they are investigated. Shultz' coverage of *functional data analysis*—while arguably the appropriate tool for identifying qualitative changes—runs a bit jargon-heavy (e.g., velocity peaks, B-spline basis function, curvature of growth, etc.), and may only stir the souls of true math-lovers. In this case, the topic is an admittedly dry one, and so there may be no sure remedy for grabbing the attention of the more casual reader. However, the latter half of chapter 5 allows Shultz to focus on a number of concrete examples, including development on both traditional cognitive tasks (e.g., conservation, balance-scale, and seriation) and linguistic measures (e.g., pronoun use and phoneme discrimination). Shultz succeeds in finishing the chapter on a high note for developmental psychologists, as the question of developmental precursors (i.e., how prior structures influence the emergence of later ones) is interpreted and addressed in computational terms.

The final section of the book (ch. 6–7) represents the culmination of Shultz's third and most laudable goal: first, to raise and then respond to the wide array of issues that critics have raised against computational modeling in general and connectionism in particular, and second, to forecast where the field of computational developmental psychology is headed. The manner in which Shultz dispatches the critics of connectionism is especially fun to watch, as he incorporates a compelling mix of empirical data, logical analysis, and just plain common sense to rebut the skeptics. Meanwhile, the last chapter highlights numerous emerging trends, each of which is likely more than a passing fad, and will inevitably exert some influence on how computational models of development are designed and studied in the next 20 years.

References

Elman, J. L, Bates, E. A., Johnson, M. H., Karmiloff-Smith, A., Parisi, D., & Plunkett, K. (1996). *Rethinking innateness: A connectionist perspective on development*. Cambridge, MA: MIT Press.

Fodor, J. A., & Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition*, *28*, 3–71.

Schlesinger, M., & Parisi, D. (Eds.). (2004). Beyond backprop: Emerging trends in connectionist models of development. *Developmental Science*, *7*, 131–132.

Matthew Schlesinger
Psychology Department
Southern Illinois University
Carbondale, IL 62901, USA
Email: matthews@siu.edu