logically plausible and realizable explanations, connectionism is not the only tool for the study of language development. Even those who share Elman *et al.*'s vision of what developmental psychology should become must look beyond connectionism.

REFERENCES

Brent, M. R. (1996). Advances in the computational study of language acquisition. *Cognition* **61**, 1–38.

Hahn, U., Nakisa, R. C. & Plunkett, K. (1997). The dual-route model of the English past tense: another case where defaults don't help. Proceedings of the 1997 GALA Conference on Language Acquisition. University of Edinburgh.

Nakisa, R. C. & Hahn, U. (1996). Where defaults don't help: the case of the German plural system. Proceedings of the 18th Annual Meeting of the Cognitive Science Society. Mahwah, NJ: Erlbaum.

Plunkett, K. & Marchman, V. (1993). From rote learning to system building: acquiring verb morphology in children and connectionist nets. *Cognition* **48**, 21–69.

Rethinking learning: comments on *Rethinking innateness**

VIRGINIA VALIAN

Hunter College and CUNY Graduate Center

In his review of *Rethinking innateness* (Elman, Bates, Johnson, Karmiloff-Smith, Parisi & Plunkett, 1996; henceforth, RI), Rispoli (this volume) comments favourably on the dynamical change models presented in RI's Chapter 4. I think a more critical stance is warranted. In particular, I will argue that dynamical change models cannot in principle make reference to mental representation, that the models are stipulative, and that they fail as descriptions of behaviour. (For more extensive discussion, see Valian, in press.) The relation between dynamical change models and connectionist networks is not spelled out in RI, so it is not clear which of the criticisms that I direct at dynamical change models also hold for connectionist models.

Absence of concepts. The emphasis of RI's dynamical change models is continuity, a different kind of continuity from that proposed by nativists, which is continuity of concepts and theoretical vocabulary; nativism says nothing directly about mechanism. In RI, continuity refers to continuity of

[*] This research was supported in part by a grant from the National Institutes of Mental Health (MH 55353) and in part by a grant from The City University of New York PSC-CUNY Research Award Program. I thank Gary Marcus, Jerrold J. Katz, Fred Katz, Thomas Bever, Martin Chodorow, and Mary C. Potter for discussion and comments. I am also grateful to the authors of RI for critical comments of an earlier draft.

mechanism. New mechanisms do not come in; old mechanisms do not die out. Concepts, if they exist at all, can emerge *de novo*.

But are there any concepts? The general form of the dynamical change models in RI is (E). (More complicated models are discussed, but they do not change the basic points discussed below.)

$$\text{(E)} \quad \mathrm{d}y/\mathrm{d}t = by + c,$$

$\mathrm{d}y/\mathrm{d}t$ refers to the change in a variable y over time t. y is always a performance measure, such as size of productive vocabulary in words or percent correct on a grammaticality task. Thus, while y is a behavioural consequence of knowledge, it is not itself knowledge. b and c are constants that represent mechanisms, such as learning ability and learning efficiency. The equations model behaviour by referring to mental mechanisms, but not to mental content or knowledge

Since we can interpret the terms of (E) very broadly, it might seem that y could stand for knowledge. If so, (E) could model the acquisition of knowledge. But for y to stand for knowledge, it would have to be reducible to values on a single dimension, like number of words in one's productive vocabulary. Lexical knowledge, however, is a congeries of different types of knowledge – phonological, syntactic, semantic, pragmatic. Even within word meaning there are different aspects of knowledge, ranging from the semantic domain the word is part of to specific aspects that differentiate one word from another. Operational definitions of lexical knowledge might be represented on a single dimension, but lexical knowledge itself cannot be.

Thus, equations like (E) are restricted to modelling behaviour and are at best only very indirectly related to the underlying abstract concepts to which the behaviour is related. Some dynamical change theorists have accepted the implications of that, saying, for instance, that 'knowledge…is not a thing, but a continuous process; not a structure, but an action, embedded in, and derived from, a history of actions' (Thelen & Smith, 1994, p. 248). Dynamical change models of development, by their nature, are uninformative about mental structure. This is a serious problem, since things like the vocabulary spurt are of interest only because of what they might tell us about the organization of mental concepts and how learning takes place.

Since RI does not develop the connection between dynamical change models and connectionist nets, its position on knowledge is hard to ascertain. When it presents a connectionist model late in chapter 4, for example, it likens stages in Karmiloff-Smith's (1992) model of cognitive development – in which knowledge undergoes representational redescription – to stages in a connectionist model that detects the difference between odd and even numbers (p. 231). Here, then, they seem to intend to model knowledge and changes in knowledge. But I don't think the analogy holds up. In Karmiloff-Smith's model the child's underlying internal representation undergoes qualitative change. In the recurrent network the behaviour simply looks as if there is a changing internal representation underneath it. Indeed, the whole point of the exercise is to demonstrate that nothing intrinsic changes except the interaction among mechanisms. As RI puts it, 'The question is, what sort of mechanisms might be responsible for what seem to be qualitatively different sorts of knowledge, and how can we move from one phase to the next?' (p. 231).

The example reflects a general inconsistency in the book's stance toward knowledge. On the one hand, it claims connectionist nets model the development of knowledge and show how new concepts, such as syntactic categories, can emerge. On the other hand, it claims that connectionist nets show that behaviour can look as if there are emergent concepts when in fact there is only learning of distributional regularities among contexts that correspond to entities that we label as syntactic categories. (See e.g. the discussion in Chapter 2.)

Stipulation. To return to dynamical change models like (E): in vocabulary development, the two mechanisms *b* and *c* can be thought of as ability to learn new words and learning efficiency. (Actually, the constants are interpreted differently in different places; I've picked what seems the most reasonable interpretation.) But calling *c* learning ability is unmotivated. RI has no theory of vocabulary development to which the constants in (E) are related. RI is not formalizing an existing theory with (E). There is no theory. The constants in (E) are merely the numbers that are needed to generate a curve. The numbers are named *ad hoc*. *c* is called 'ability to learn', but could be called anything.

Even if we assume that 'ability to learn' is the right name, we face the problem of stipulation. (E) simply stipulates that there is a single ability that remains constant, shedding no light on what that ability is. (E) treats as primitives the processes to be explained and described thus begging the developmental question.

The virtue of existing hypotheses for the vocabulary spurt (like the 'naming insight') is that they are explicit and specific enough to be tested and shown to be incorrect. The only explicit and specific components of RI's change models are the dependent variables that need to be accounted for in the first place. We could not test whether learning ability is a constant because we have no idea what learning ability is or how to measure it. That problem is hardly RI's alone. But RI treats learning ability as if we could take for granted that it is a univocal constant process.

A hypothetical example shows the difficulty. Someone who goes into a coma and then comes out of it really does experience a different rate of learning: there is a real step function down, followed by flat learning, followed by a step function up. RI could model the data as a dynamical change function by smoothing the steps out into a *U*. In this example we know that the true story includes a second process. In vocabulary development, second language learning, and birdsong learning, we don't know what the true story is. Exactly what is at issue is the identity of the processes and their time courses. RI says it's possible that the processes are continuous. True enough. It's also possible that they aren't. It's possible that there are five different processes that enter and exit at different times. A possibility is not enough to lead us to prefer one model over another.

Curve-fitting. In principle, the choice of appropriate constants and values of constants will allow one to model the time course of almost any phenomenon. But the family of curves RI uses to fit data that resemble a step function smoothes out the function. There is no longer any 'step', as there would be in the coma example and as graphs like Figs. 4.7 and 4.9 show. On what grounds, then, would one prefer a curve that is so much more removed from the data? RI's answer, for both second language learning and birdsong learning, is that the dynamical change curve is a better empirical fit: it accounts for more of the variance in the phenomenon.

Does it? Take the data for second language learning (Johnson & Newport, 1989). Up until about age 16, second language learners' scores on a grammaticality judgement task decline with age of learning onset in an orderly, linear-looking way. In contrast, between about 16 and 40, no correlation is found; score variance also increases enormously. By eye, there is a downward sloping line and then an almost horizontal line. There appears to be something like a sensitive period for language learning, ending around age 16 (Johnson & Newport, 1989).

How much of the variance in language learning between birth and age 40 is accounted for by those two lines? RI says 39 %, a figure obtained by averaging the variance accounted for in each of the two lines. The dynamical change curve accounts for 63 % of the variance. But visual inspection of the different curves shows that RI's curve doesn't account for the variance in the first part of the data nearly as well as Johnson & Newport's (1989) line. And it does no better at accounting for the variance in the second part. How, then, could it account for more of the variance? It couldn't. The calculation of total variance allows the RI curve to profit from the correlation in the initial segment; the calculation of average variance prevents Johnson & Newport's two curves from profiting from its high initial correlation. What should be done is to total the squared deviation of each point from its respective line. In that case, the Johnson & Newport pair of lines would account for more of the variance. RI also reproduces data on birdsong learning (Marler & Peters, 1988). The empirical data show percentage of correct song learning flat at about 60 % for the first 75 days or so of bird life. At that point there is a dramatic drop in correct learning, with the percentage varying between about 20 % and 5 % but having a downward slope. As with the second language learning data, there is also an increase in variance.

The curve generated by the dynamical change model wipes out the step function, thus missing the empirical phenomenon altogether. The birdsong data are the closest thing in nature that we have to the hypothetical example of a person entering a coma. Yes, we can model it as dynamical change and claim that the step function is a mirage. But why would we want to do that? Don't we instead want to find out whether what looks like a change is indeed a change that results in inability to learn?

I am not claiming that there is a sensitive period for language learning or birdsong learning, but that the discontinuities in the data cannot be eliminated by curve-fitting. The particular equations RI uses sacrifice the phenomenon that needs explanation. By producing a curve that wipes the phenomenon out we go backwards rather than forwards, because we make it seem as if there is nothing to explain.

Simplicity. RI suggests that we should prefer dynamical change models on the grounds of simplicity: they don't add or change mechanisms, while hypotheses like the critical period hypothesis for language or birdsong learning do. But two models can only be compared in simplicity if they account for exactly the same set of phenomena. That is not the case here. I have already noted that the dynamical change models do not account for the step function but eliminate it.

Even if that were not so, there is the problem of accounting for a cluster of phenomena. As RI notes, in the second language learning data there is an increase in variance at the same point that grammaticality judgements stop correlating with age of exposure. The dynamical change model does not account for that cluster. So it cannot be meaningfully compared with a sensitive period hypothesis; it does not account for the same data. The data reported for birdsong learning also show increased variance at the end of a hypothesized sensitive period, suggesting some generality to the phenomenon.

One possibility is that learning is tightly controlled by a dedicated mechanism until the end of the critical period. In language learning that mechanism appears to be less and less robust with age, while in birdsong learning the mechanism is equally strong throughout the critical period. In both cases, variance among learners is limited because the language-learning mechanism is dominating behaviour. When the critical period ends, the dedicated mechanism plays either a very limited role or no role; other, more general, learning mechanisms come into play. Those other mechanisms are both less efficient and more variable.

Is this story on the right track? We don't know, but we can find out. In humans, for example, we would expect post- but not pre-sensitive period behaviour to correlate with measures of cognitive efficiency. There are, of course, other approaches to the same phenomena. In second language learning, the mechanism of language learning may stay the same but accessibility of linguistic concepts may decline. That too is something we can test. By looking at the entirety of a phenomenon we can get hints about what models could explain it.

Because dynamical change equations model only one behaviour at a time, they are particularly unsuited for describing domains like language development where phenomena are linked (see Marcus, in press, for a similar point). Thus, such models oversimplify and mischaracterize the phenomena to be explained.

Conclusion. Dynamical change models of language acquisition tell us nothing about the contents of the mind, do not fit the data, and are not simpler than other models.

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.

Johnson, J. S. & Newport, E. (1989). Critical period effects in second language learning: the influence of maturational state on the acquisition of English as a second language. *Cognitive Psychology* **21**, 60–99.

Karmiloff-Smith, A. (1992). *Beyond modularity: a developmental perspective on cognitive science*. Cambridge, MA: MIT Press/Bradford Books.

Marcus, G. (in press). Can connectionism save constructivism? *Cognition*.

Marler, P. & Peters, S. (1988). Sensitive periods for song acquisition from tape recordings and live tutors in the swamp sparrow *Melospiza georgiana*. *Ethology* **77**, 76–84.

Thelen, E. & Smith, L. (1994). *A dynamic systems approach to the development of cognition and action*. Cambridge, MA: MIT Press.

Valian, V. (in press). *Input and innateness: controversies in language acquisition*. Cambridge, MA: MIT Press.

Understanding the modelling endeavour

KIM PLUNKETT

Oxford University

JEFFREY L. ELMAN AND ELIZABETH BATES

University of California, San Diego

We enjoyed reading Matthew Rispoli's review of *Rethinking innateness* (henceforth RI-Elman *et al.* 1996). First, we respond to what we see as the major issues that he raises. Finally, we respond to the commentaries contributed by other authors in this issue.

Rispoli has two major concerns. First, he fears that RI will do the field a great disservice, increasing polarization and sowing discontent. To some extent this is an 'eye of the beholder' criticism that is difficult to counter: if Rispoli feels more polarized upon reading this book, who are we to argue?