Developmental Science 6:4 (2003), pp 430–433

Development: it's about time

Jeff Elman

Department of Cognitive Science, University of California, San Diego, USA

Abstract

Traditionally, research has been defined as being about *development if it focuses on behaviors that* occur *during development, without reference either to what precedes or follows, or – more importantly – the mechanisms that drive change. This perspective has been challenged by two new approaches: dynamical systems, and connectionism, both of which have as core goals the explanation of change. Interestingly, these approaches have also been seen by many as being at odds with one another, for reasons that are in part valid, and in part superficial. Recent developments in both approaches suggest that a convergence is occurring, although much remains to be done before a true unification can take place.*

Introduction

Development: it's about time. Well, other things too. But what is surprising is that, with a few notable exceptions, very little work in developmental research seems to focus on what is most striking about development, namely, that it involves change over time.

This tendency to study development as if it were a succession of snap-shots – frozen in time and fixed – is more than odd. It biases the theoretical perspective, favoring theories in which the origins of behavior are of minimal interest. In the extreme, this view gives rise to recent papers with titles such as 'Language acquisition in the absence of experience' (Crain, 1991; acquisition being seen as little more than a triggering of innate potential), or statements such as 'initial knowledge may emerge through maturation or be triggered by experience, but learning and processing do not appear to shape it' (Spelke, 1994, p. 439). This maturational metaphor is increasingly invoked in a way that appears to minimize interest in understanding the mechanisms of change, presumably on some belief that maturation itself merely reflects the unfolding of structures and knowledge already internally present – a belief that is eerily reminiscent of preformationism, and squarely at odds with what is now known about the biology of development. There are no homunculi hiding the sperm (or, as some preformationists believe, humunculae in the ova). Development is an emergent process in which the biology is better seen as a constraint on interactions with the environment, and not as a process by which innate knowledge is revealed over time.

Thus, against this theoretical backdrop, it is clear that the amount of agreement between connectionist and dynamical systems approaches to development is overwhelming in comparison to their differences. Both approaches take change as the phenomenon of interest. Both agree that understanding the genesis of behavior is the best route to understanding the behavior itself. And there are additional points of agreement that stand in stark contrast to what I've described as the static view of development. Many of these points are usefully elaborated in the papers by Thelen and Bates, Munakata and McClelland, and Spencer and Schöner (this issue), but let me list what I see as the major items:

The origin of complexity

Rather than seeing behavior as intrinsically fractionated and arising from independent modules, behavior is seen as an emergent consequence of interactions across multiple domains. These are by no means insights peculiar to the connectionist or dynamical systems frameworks. Thelen and Bates cite D'Arcy Thompson and Piaget, and might easily have added others (e.g. Wolff, Waddington, or even Aristotle, to whom we owe the term 'epigenesis'). But although this view is non-controversially and widely accepted among developmental biologists, within behavioral and cognitive development it has been primarily

Address for correspondence: Jeffrey Elman, University of California, San Diego, Department of Cognitive Science, 9500 Gilman Drive La Jolla, CA 92093-0515, USA; e-mail: jelman@ucsd.edu

connectionist and dynamical systems advocates who have argued for this position.

Modularity

Behavior may appear to be modular and domain-specific, but such modularity and domain-specificity are better understood as functional outcomes rather than characteristics of the underlying mechanisms, and often soft and graded in nature. There are many aspects of language that appear relatively modular. For instance, there are systematic and significant differences in the way the regular form of the English past tense is processed, compared with irregulars (e.g. reaction times to naming, magnitude of brain activity as revealed by fMRI, etc.). But connectionist models have shown that such differences can arise because of intrinsic stimulus differences between the classes of verbs themselves, both of which are learned by the same learning mechanism, and not because they are processed by two different language modules. Or consider the finding that certain regions of the cortex are most active in skilled chess players at different points in the game (Nichelli, Grafman, Pietrini, Always, Carton & Miletich, 1994). No one would argue that there has evolved an innate 'checkmate' module. Rather, these brain areas have intrinsic capabilities that lend themselves to being recruited – as a result of learned expertise – to serve the specific needs of chess.

Nonlinearity

The processing mechanisms that underlie behavior are fundamentally nonlinear. This means that development itself will frequently have phase-like characteristics, that there may be periods of extreme sensitivity to input ('critical periods'), that change may be nonlinear and even nonmonotonic, and that the effects of damage or disease will be strikingly different at different points in time. It also means that behaviors may at times seem categorical and binary (deceptively, giving the appearance of symbolic processing) and at other times, graded and exquisitely sensitive to small differences in input or experience.

Origins of control

Control of processing is distributed over a network of processing elements. This does not mean that behavior is fragmented or incoherent. It means rather that what the organism does at any point in time reflects the dynamic interplay between competing demands and processes. Control is a consequence, not a cause.

The role of time

Most importantly, time figures as a central dimension in both dynamical systems and connectionist approaches. Time is of course one of the theoretical primitives in dynamical analyses, so its role is inevitable in a way that it is not for connectionist models. But connectionists have recognized, from the very first (and still perhaps most widely cited) neural network model on development – the past tense model of Rumelhart and McClelland (1986) – that explaining why a behavior changes over time can be key to understanding the behavior itself.

These areas of agreement are central to both the connectionist and dynamical systems approaches. Thus it might seem a bit perplexing to some that the two approaches are often viewed as in competition, if not in outright opposition. And indeed, there are differences that we should not gloss over. Many of these differences are due to historical accident, having to do with the particular focus of individual researchers and to sorts of theoretical issues that occupied their attention. But the differences are real and raise the question about whether a change in agenda in either or both of these approaches might be warranted.

Dynamical systems approaches, for reasons that are laid out in detail in both the Thelen and Bates and the Spencer and Schöner papers, have historically focused on motor behavior. They have also eschewed any sort of mentalism or interest in internal representations (as Thelen & Bates put it, 'there are no inferences about unobservable mental structures that exist outside of the behavior-in-context'). These strike me as the two most significant areas in which the dynamical systems approach differs from most connectionist work.

The focus on motor behavior has been enormously fruitful, both as a domain for elucidating principles of dynamical systems in a biological context, and also because it has led to an appreciation of how important an understanding of the body is for an understanding of the mind. This latter point is not at all self-evident, although – in large part due to the efforts of the dynamical systems folks – it is increasingly acknowledged in mainstream cognitive science (e.g. Clark, 1997; Damasio, 1994).

It would be unfair and inaccurate to say that connectionists have not been at all interested in the body. But the interest has been focused and limited to the brain. Indeed, Munakata and McClelland argue that a strength of connectionist models of development has been the focus on the underlying neural mechanisms of behavior. Still, it seems fair to say that, with this one important exception, connectionists have by and large not turned to the body as a source of explanation for understanding cognitive phenomena.

This strikes me as unfortunate, although I hasten to say that it is by no means a fatal deficiency, since disembodiment has never been a plank in the connectionist platform. And there are many ways in which taking the body into account – and in particular, changes during development in both the sensory and motor systems – might interact with learning. For example, the initial immaturity and wide spacing of photoreceptors in the infant retina, as well as limitations on the accommodative system, significantly limit what the infant sees. The specific effect is to filter out high spatial frequency information, and to make objects that are close to the infant most salient. It is not unreasonable, as Turkewitz and Kenney (1982) have proposed, that such limitations may facilitate learning about size constancy. That is, a limited depth of field allows the infant to learn about the absolute size of objects before having to deal with apparent size transformations that result when objects are viewed at a distance. Similar arguments can be made to explain why kittens' eyes are closed for the first week of life: postponing visual input allows the kitten to organize important behaviors such as homing around tactile and olfactory inputs, which might be simpler to deal with than visual stimuli. Once these behaviors have been developed, they provide a converging source of information that makes it easier for the kitten to deal with the more complex visual world. These examples are of course hypothetical, but it seems to me reasonable to believe that the fact that the body is changing so dramatically at precisely the point in time when the most learning is going on will interact with how that learning proceeds.

The dynamicists' suspicion of mental representations I view as in the long run less healthy. There is certainly much to be said for being wary of invoking invisible entities. And we must never forget that constructs such as attention, consciousness or belief are hypotheses that have been advanced to explain data. These constructs should not be confused with the data themselves. Braitenberg's (1984) *Vehicles*, Brooks' (1991) 'Intelligence without representation', and the skepticism voiced by many dynamical systems researchers are thus healthy correctives to a mentalism in danger of running amok. Furthermore, work such as that by Thelen, Smith and their colleagues on the physical basis for the A-not-B error has provided an impressive alternative to the more traditional 'cognitive' explanation, and generated a number of non-obvious predictions – many with older children – that have been verified experimentally. That said, one would not want to preclude any role at all, ever, for mental representations in behavior (i.e. precisely those 'inferences about unobservable mental structures' that Thelen and Bates refer to). The work described by Spencer and Schöner reflects an important and significant step in developing the dynamic systems approach in ways that (in their words) 'bridge the representation gap in the dynamic systems approach'. This work is encouraging, although obviously much remains to be done.

The strengths and weaknesses of the connectionist approach have in many ways been complementary to the strengths and weaknesses of the dynamical systems approach. First, where the body (apart from the brain) has been neglected, mental representations have not. Indeed, the title of one of the most influential papers on connectionist learning was 'Learning internal representations by error propagation' (Rumelhart, Hinton & Williams, 1986). Within the realm of developmental research, understanding the ontogenesis and character of internal representations has been a major focus, with significant pay-offs. For example, the concept of 'graded representations' discussed by Munakata and McClelland (see also the distinction between 'latent' and 'active' representations), has been an important contribution to the field. Among other things, this notion provides a very rich view about why and when knowledge–action dissociations might arise.

Second, learning – which goes to the heart of the connectionist enterprise – is clearly something that is as yet not well elucidated in dynamical systems approaches. It is not that there is no role for input and experience in dynamical systems models; it is simply that the mechanisms by which these models change as a function of experience is rarely specified, and there is nothing in that literature that approaches either the empirical or theoretical work on learning that one finds in the connectionist literature. If understanding how the developing body interacts with learning represents a fertile ground to be cultivated by connectionists, then connecting dynamical systems with models of learning and adaptation represents a crucial next step for dynamical systems researchers.

But this may be the wrong way to phrase things. It implies that these two communities should continue along separate lines with separate identities, engaged in a hopefully friendly competition for theoretical limelight. The fact is that the historical and theoretical differences between the approaches have at this point been outgrown. The dynamical field approach is an example of a productive rapprochement. Another example is work in which dynamical systems analysis has been applied to recurrent neural networks in order to better understand, in a more formal way, the computational properties of such networks (e.g. Blair & Pollack, 1997; Boden, Wiles, Tonkes & Blair, 1999; Rodriguez, 1999; Rodriguez & Elman, 1999; Rodriguez, Wiles & Elman, 1999; Siegelmann, 1999; Wiles & Elman, 1995). This work has been motivated by the desire to 'deconstruct the black box': recurrent networks *look* different than traditional discrete autonoma, but *are* they truly different at a deep level? The exciting result seems to be that yes, there are important differences and these differences are consistent – in a direction that favors the networks – with some aspects of human cognition. Such networks have the ability to capture hierarchical constituent structure (e.g. of the sort required to process sentences that contain relative clauses and other embedded material), but unlike traditional symbolic systems, there is an interaction between the contents of what is being stored (e.g. the particular words) and the structure. This leads to the interesting result that although sentences such as 'The rat the cat the dog chased hunted ran away' and 'Do you believe the report the stuff they put in coke causes cancer is true' are both structurally very complex, humans – and these networks – find the latter much easier to process than the former (e.g. Weckerly & Elman, 1992). The dynamical systems analysis now gives us a formal understanding of why this should be.

These are beginnings. What one would hope to see now is a growing community of developmentalists who are conversant with the tools and conceptual vocabulary of both approaches, and the emergence of a more unified and undoubtedly richer framework for studying development. I believe we are moving in that direction. Although we still have a way to go, the ideas presented in this special issue point toward a future that seems quite promising. It's about time.

References

Blair, A.D., & Pollack, J.B. (1997). Analysis of dynamical recognizers. *Neural Computation*, **9**, 1127–1142.

Boden, M., Wiles, J., Tonkes, B., & Blair, A. (1999). Learning to predict a context-free language: analysis of dynamics in recurrent hidden units. *Proceedings of the International Conference on Artificial Neural Networks*, 359–364.

Braitenberg, V. (1984). *Vehicles: Experiments in synthetic psychology*. Cambridge, MA: MIT Press.

Brooks, R.A. (1991). Intelligence without representation. *Artificial Intelligence Journal*, **47**, 139–159.

Clark, A. (1997). *Being there: Putting brain, body and world together again*. Cambridge, MA: MIT Press.

Crain, S. (1991). Language acquisition in the absence of experience. *Behavioral and Brain Sciences*, **14**, 597–611.

Damasio, A.R. (1994). *Descartes' error: Emotion, reason, and the human brain*. New York: Avon Books.

Nichelli, P., Grafman, J., Pietrini, P., Always, D., Carton, J.C., & Miletich, R. (1994). Brain activity in chess playing. *Nature*, **369**, 191.

Rodriguez, P. (1999). Mathematical foundations of recurrent neural networks in language processing. PhD dissertation, Department of Cognitive Science, University of California, San Diego.

Rodriguez, P., & Elman, J. (1999). Watching the transients: viewing a simple recurrent network as a limited counter. *Behaviormetrika*, **26**, 51–74.

Rodriguez, P., Wiles, J., & Elman, J. (1999). A recurrent neural network that learns to count. *Connection Science*, **11**, 5–40.

Rumelhart, D.E., Hinton, G.E., & Williams, R.J. (1986). Learning internal representations by error propagation. In D.E. Rumelhart & J.L. McClelland (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition, Volume I* (pp. 318–362). Cambridge, MA: MIT Press.

Rumelhart, D.E., & McClelland, J.L. (1986). On learning the past tenses of English verbs. In J.L. McClelland & D.E. Rumelhart (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition, Volume II* (pp. 216–271). Cambridge, MA: MIT Press.

Siegelmann, H.T. (1995). Computation beyond the turing limit. *Science*, **268**, 545–548.

Spelke, E.S. (1994). Initial knowledge: six suggestions. *Cognition*, **50**, 431–445.

Turkewitz, G., & Kenny, P. (1982). Limitations on input as a basis for neural organization and perceptual development: a preliminary theoretical statement. *Developmental Psychobiology*, **15**, 357–368.

Weckerly, J., & Elman, J.L. (1992). A PDP approach to processing center-embedded sentences. *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society* (pp. 482–487). Hillsdale, NJ: Erlbaum.

Wiles, J., & Elman, J.L. (1995). Learning to count without a counter: a case study of dynamics and activation landscapes in recurrent networks. *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society* (pp. 442–447). Hillsdale, NJ: Erlbaum.