

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Computational Models of Development: A Symposium

Permalink

<https://escholarship.org/uc/item/228948qt>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 18(0)

Authors

Plunkett, Kim

Shultz, Thomas R.

Publication Date

1996

Peer reviewed

Computational Models of Development: A Symposium

Kim Plunkett

Department of Experimental Psychology
Oxford University
South Parks Road, Oxford, OX1 3UD
plunkett@psy.ox.ac.uk

Thomas R. Shultz

LNSC, Department of Psychology, McGill University
1205 Penfield Avenue
Montréal, Québec, Canada H3A 1B1
shultz@psych.mcgill.ca

Introduction

Any theory of development must specify its position on three crucial assumptions:

1. The nature of the start state of the organism.
2. The effective learning environment that the organism occupies.
3. The learning procedure that enables the organism to alter its start state via interactions with the learning environment.

Once committed to these assumptions, the developmentalist can make predictions about the trajectory of learning that the organism must follow in order to achieve mature levels of performance. These predictions can be compared to the actual behaviours that the organism produces en route to masterful behaviour. A certain amount of slack is permitted in the fit of the predictions to actual behaviour depending on the manner in which the original assumptions have been stipulated—some theories will permit a greater degree of individual variation than others. In principle, the outcome of the comparison of the predicted developmental trajectory to the organism's actual developmental trajectory will lead to an acceptance or rejection of the theory. In the case of rejection, one or more of the original assumptions may require modification and the process of prediction and comparison reiterated until a successful match can be achieved.

This idealised process of theory building and evaluation is rarely achieved for complex organisms like *homo sapiens* because of the difficulties involved in stipulating the content of the three assumptions outlined above: We still know very little about the genetic code and neural architectures that specify the start state of human development. The nature of the learning environment for young humans may seem self-evident. However, ethology has taught us to be wary of equating environmental input with perceptual uptake. The sensitivities of young humans to features of the environment may vary considerably even from their mature caregivers. Some success has been achieved in discovering neural mechanisms of learning, such as long term potentiation in the mammalian brain. However, it is unlikely that this Hebbian-style learning is the only learning mechanism available to humans. Ignorance as to what gets learning off the ground in humans undermines the constraints necessary for theory building and evaluation.

When there are so many under-specified, free parameters in the theory, the reasons for failure to predict a developmental trajectory become unfathomable.

There are, nevertheless, grounds for optimism for theory building in developmental cognitive science. First, significant progress is being made in decoding the human genome and neuro-physiologists are making dramatic headway in identifying the neural circuitry and constraints on neural plasticity in developing organisms. These advances are likely to have a major impact on developmental psychology over the next decade. Second, developmental psychologists themselves are making enormous headway in understanding the sensory and perceptual sensitivities of human newborns. We know a lot more about the infant's effective learning environment than we did just one decade ago. Thirdly, the advent of computer modelling in developmental psychology permits the rapid evaluation of theories so that a wide range of configurations of start-states, effective learning environments and learning procedures can be explored. In this symposium, we focus on the third of these advances—the application of computational modelling to theory building in developmental cognitive science. We evaluate and compare some of the lessons we have learnt from the study of two particular classes of models—symbolic and neural network simulations of development.

The Role of Modelling

In order to build a computer model of a developmental process, you need a precisely formulated theory—precise enough to specify in a computer programme. This demand for precision forces the theory to be coherent. It may be the wrong theory, but at least it will be internally consistent. If it is not, it won't run properly on the computer. The modelling endeavour forces an elaboration of the fundamental principles underlying the developmental process. The model cannot encompass the whole child, so the modeller must decide which parts of the child are crucial for the problem of interest. This is part of the theory building process itself, but a part that is enforced by the decision to implement the theory on a computer.

Once these parts of the model building process are complete, the next step is to determine whether the programme produces the kind of behaviour produced by developing children. This is not always a straightforward part of the process. Children vary in their behaviour so one must decide whether any observed deviation from the

model's performance constitutes a refutation of the model. It is best if the model can vary in its performance too and these variations are caused by factors that might be related to factors underlying variation in children. The theory embodied in the model must be able to make novel predictions that can be tested empirically. In other words, it is not enough for the model to mimic children's behaviour. It must be able to generate testable hypotheses.

Connectionist and Symbolic Modelling

All the papers in this symposium present working models of some aspect of the child's linguistic or cognitive development from a connectionist or symbolic perspective. Both classes of model make their own characteristic assumptions about the nature of the start state, learning algorithms and effective learning environments. All the models assume a set of *patterns* to be processed by some computational architecture. The fundamental difference between the two classes of model is the nature of the processing and inference procedures employed. For connectionist models, patterns are manipulated by a set of adaptable weighted connections that are sensitive to the *content* of the forms encountered. Learning amounts to a process of statistical inference. In symbolic models, patterns are manipulated by a set of *rules* that may be impervious to certain details of the patterns. If a pattern is identified as a certain *type*, then it is handled according to the rules specified for that type, irrespective of any idiosyncrasies concerning the particular token of the type. Learning is often handled by the construction of decision trees that classify the patterns in terms of their symbolic attributes.

Nativism and Empiricism

A common point of contention between connectionist and symbolic models of development concerns the amount of structure that is required to get learning off the ground. From the symbolic perspective, it is sometimes held that stimuli (or training patterns) are too impoverished to account for the richness of human cognition—Plato's well-known "Shadows in the cave" problem. Learning is often seen as a process of triggering innate knowledge bases, where the trigger operates in a domain-specific fashion. Many connectionists maintain that symbolic theorists have underestimated the structural complexity of information available in the environment and the capacity of domain-general learning algorithms to induce domain-specific knowledge bases from this complexity.

The input sensitivity of connectionist systems makes them obvious tools for exploring empiricist accounts of linguistic and cognitive development. However, it should be emphasized that connectionism and empiricism are not necessary bedfellows. In fact, connectionism offers a tool for examining the trade-off between the role of the input and the role of pre-adapted structures and processes in development. Although the representations formed by connectionist systems are indeed highly sensitive to input parameters, it is the architectures and learning algorithms of connectionist systems themselves that afford this sensitivity. Because connectionist architectures differ in terms of their network structures and learning algorithms, they also differ in the

manner in which they respond to the same inputs. The potential variety of network architectures thereby contributes to the range of hypotheses that might describe the initial state of the learning device before it is exposed to any input. Consequently, connectionism can also be used to explore interactionist and nativist accounts of linguistic and cognitive development.

Symbolic models also offer considerable flexibility in exploring a range of developmental hypotheses. Although symbolic models necessarily postulate a set of symbolic primitives, the granularity of these atomic building blocks can vary can vary dramatically from one theory to the next. The commitment to the software-hardware distinction that commonly goes hand-in-hand with the symbolic approach, does not preclude the exploration of minimalist programmes of *innate structures and processes*.

Biological Plausibility and Levels of Explanation

Connectionist models clearly have a neurological appeal. But are connectionist systems constrained to providing models at the neurological level of explanation? Many researchers working within the classical symbolic approach to cognitive science argue that functional accounts of the cognitive level must be couched in terms of discrete, categorical, symbol processing systems. Furthermore, they argue that current connectionist models do not behave in the necessary symbolic fashion. According to this argument, connectionist models will not be able to provide explanations and descriptions at the cognitive level. It is conceded, however, that connectionist models, appropriately hard-wired, may be able to implement the foundations of a cognitive system in much the same way that the hardware of a computer provides the necessary working environment for symbolic programmes. Indeed, it is widely acknowledged that something like a connectionist system must provide the neurological foundations for the apparent symbolic mind. On this view, a symbol processing machine sits on top of a connectionist implementation of the neurological system. It makes sense to talk about a two-level system where the symbolic machine operates according to its own autonomous set of principles.

Many developmental connectionists resist this relegation of their explanations to the implementational level. One of the primary motivations for building connectionist models of cognitive processes is that symbolic approaches seem to lack certain characteristics that are needed at precisely the cognitive level of functioning (parallel processing, graceful degradation, learning).

A compromise solution, in which connectionist mechanisms and symbol manipulating devices work side-by-side in an harmonious cognitive system, is currently fashionable—so-called hybrid systems. Do these offer the way forward or are they merely a confusion in levels of description?

The present symposium features modelling approaches that are either purely symbolic or purely connectionist.