Journal Title: The Behavioral and brain sciences

Volume: 14
Issue: 1
Month/Year: 1991
Pages: 101-102

Article Author: Miller,

Article Title: Let evolution take care of its own

Imprint: google IU-Link

Notes:

Skyler Place (ssplace)
993 Woodbridge Drive
Bloomington, IN 47408

We close by noting that the purposes of scientific theory are to organize knowledge, generate falsifiable hypotheses, and predict new phenomena. Although such ideas as optimality have played a critical role in theory construction, they confuse the process of theory evaluation. What matters in evaluation is how well quantitative models generate an understanding of processes and outcomes. Within this perspective, mathematical equations can potentially predict behavior such as probability matching (e.g., Heyman 1988) and environmental sampling (e.g., Stephens 1987), not whether errors determine optimality. Questions about the optimality of such behavior are not falsifiable, are outside the model's domain, and only promote loose speculation and circular argument. We hope that next-generation models will be sufficiently productive/predictive that researchers will focus on developmental issues to the exclusion of definitional controversies.

NOTE
*Correspondence should be sent to M. L. Howe.

Limits to stochastic dynamic programming

Ruth H. Mace[a] and William J. Sutherland[b]
[a]School of Development Studies, University of East Anglia, Norwich NR4 7TJ, England and [b]School of Biological Sciences, University of East Anglia, Norwich NR4 7TJ, England

Clark provides a useful description of the technique of dynamic programming and outlines its main advantages. His major worry concerns the complexity of models. He warns (sect. 10) of the temptation to include excessive detail and how this can lead to the sort of pointless heavyweight exercises that capsized and sank systems ecology. We think it is important to understand the real problem of producing complex models.

In his conclusion, Clark points out that the standard simple optimisation models of behavioral ecology have failed to stand up to quantitative tests, although they have provided qualitative insights. The real success of such simple models as the optimal diet model, the marginal value theorem and the ideal free distribution was that they changed the perspective of a cohort of ecologists so that they studied animals in a different way. This led directly to the recognition of the importance of components such as misidentification (Hughes 1979), kleptoparasitism (Thompson 1983), variation in prey quality (Durrell & Goss-Custard 1984), and individual differences in predator quality (Sutherland & Parker 1985).

Dynamic models will probably fail to stand up to quantitative tests also (albeit in different ways), but for the same reason it will not matter. If the technique can contribute to interest in new sorts of problems then it will have made a real contribution. It has already started to do so, as the lack of a quantitative framework in which to incorporate stochasticity, the time dimension and competing demands for an animal's attention definitely contributed to behavioural ecologists ignoring the importance of these factors in the past. Now, the importance of tradeoffs, central to all sorts of decision-making, is being widely examined both through modelling (e.g., Mace & Houston 1989) and purely empirical work (e.g., Cuthill & Guilford 1989).

Dynamic programming involves an interaction between nature, computers, and human brains. Nature is complex and computers are becoming increasingly capable of describing such complexity. The "curse of dimensionality" is as much a problem for human understanding as it is for the power of the computer. Computers have no problem handling four or five dimensions – the weak link is the human brain. It is possible to create models with more than one state variable and several behavioural options (and solve them numerically). But, in our experience, as the complexity of the model exceeds one state variable or two behavioural options it can become increasingly hard to make sense of the output.

Clark's abstract states that "limitations arise because nature's complexity always exceeds our modelling capacity," but this is not the real problem. The major challenge is to abstract the complexities of nature in a way that will capture the imagination of its students.

Models are just prostheses for our brains

Manfred Milinski
Zoologisches Institut, Abteilung Verhaltensökologie, University of Bern, Wohlenstrasse 50 a, CH-3032 Hinterkappelen, Switzerland

Without using algebra we can only make qualitative hypotheses in our heads and express them in words. A mathematical model is (only) a hypothesis formulated quantitatively and expressed in numbers. There is nothing more to a model that deserves either condemnation or glorification. The model is a tool to formulate hypotheses for research when the natural phenomenon under study is too complex to be handled by the limited channel capacity of our brain. In this respect, the dynamic modeling technique is no different from other well accepted models. It allows us to put more realistic complexity into our hypotheses, but at the same time it often robs us of the excuse that it is impossible to predict a behavior quantitatively because of its complex conditions.

As an empiricist I am as happy to have this new and more powerful instrument as I am to have a new and more powerful word processing program. However, in both cases the value of the results produced with the new tool depends very much on what I am using it for. The laborious part of the job is concealed in sentences like the following (Clark's abstract, emphasis mine): "The models use *biologically meaningful* parameters and variables, and lead to testable predictions." I hope that editors keep this in mind when they soon receive vast numbers of manuscripts starting with "using dynamic programming techniques we have demonstrated that . . ."

Let evolution take care of its own

Geoffrey F. Miller[a] and Peter M. Todd[b]
Psychology Department, Jordan Hall, Building 420, Stanford University, Stanford, CA 94305
Electronic mail: [a]geoffrey@psych.stanford.edu and [b]todd@psych.stanford.edu

Clark's title is somewhat misleading. Rather than modelling *behavioral adaptations* per se – specific psychological mechanisms capable of generating adaptive behavior – Clark's dynamic programming method computes specific behavioral sequences designed to optimize fitness given precalculated *adaptive pressures*. Indeed, adaptive pressures are precisely what Equation 9, the central expression of Clark's method, represents.

Yet Clark's title does point in the right direction: Behavioral adaptations *are* what we want to model. Characterizing evolved psychological mechanisms themselves is crucial to understanding behavioral responses to the adaptive pressures that emerge in complex environments, for it is at the level of mechanisms, not individual behaviors, that analyzable regularities most clearly appear. As Cosmides & Tooby (1987) argue, psychological adaptations must serve as the intermediary between adaptive pressures and behavioral strategies; one cannot take the shortcut of finding behavioral strategies directly given adaptive pressures.

The phenomenon of protean (adaptively unpredictable) behavior illustrates this levels-of-analysis problem. Simple proteanism occurs when a rabbit flees a fox by "randomly" darting back and forth (Driver & Humphries 1988). If the rabbit had internalized the sort of look-up table for escape behavior suggested by the dynamic programming method, always choosing the "optimal" escape route in its attempts to maximize Equation 9, the very predictability of this behavior would render it unfit. Foxes would evolve predictive counterstrategies. Suppose instead that rabbits have not simply evolved a set of behavioral strategies per se (as suggested by the dynamic programming method), but a more abstract, flexible behavior-generating mental mechanism that allows them to behave unpredictably in certain circumstances. Although this mechanism may violate dynamic programming optimization, perhaps causing some rabbits to perform suboptimally in the short term (e.g., zigging when they "should have" zagged), this mechanism may nonetheless increase the average fitness across the subpopulation of those rabbits possessing it. Although Houston & McNamara (1988) allude to the possibility of dynamic programming selecting probability distributions across behaviors (which would yield a kind of proteanism), the proper level of analysis here is that of the complex protean psychological mechanisms themselves. These mechanisms are the true behavioral adaptations, but ones that dynamic programming seems incapable of revealing.

More seriously, dynamic programming seems unable to adequately model the optimization of *inclusive* fitness (Hamilton 1964), rather than just *individual* fitness. With inclusive fitness, there is no specifiable final time T beyond which a behavior's effects will not propagate; because the effects of an organism's behaviors may continue long after its death, affecting its kin and offspring for many generations, there is no reasonable endpoint for assessing ultimate fitness. Thus our models of behavioral adaptations must consider fitness effects of interactions *between* individuals, both within and across generations, not just within an individual's own life-time. Dynamic programming may be sufficiently powerful in principle to represent the interaction contingencies of social behavior by breaking them down into adaptive pressures impinging on organisms considered individually. But if one tries to imagine exactly how this would work with collaborative or competitive behaviors as complex as coalitional aggression or social exchange, dynamic programming seems less than entirely efficient.

Modeling interactions with other individuals in the environment leads naturally to modeling interactions with the environment itself. This step would free us from specifying quantitative adaptive pressures impinging on the individuals. Rather, the adaptive pressures molding the evolution of behavioral mechanisms could emerge from the dynamics of the modeled environment and the fitness function defined over it.

Finally, modeling actual reproduction and inheritance directly seems simpler than representing adaptive pressures in terms of expected future reproduction or some other abstract fitness construct. The reason creatures operate in accordance with inclusive fitness is that by aiding their relatives they are aiding the spread of their own genes – their relatives are likely to have copies of their own genotypic specifications of phenotypic mechanisms. Modeling inclusive fitness without actually modeling the spread and recombination of genes just misses the point. These considerations lead us to wish for a method of modeling the evolutionary spread of successful psychological mechanisms from one generation to the next, in response to adaptive pressures emerging from a specified environment with which, and within which, individuals interact.

Clark himself suggests that "An intriguing possibility is to use the computer to emulate the evolutionary process in searching for optimal or ESS strategies via a process of natural selection, but to my knowledge this has not yet been attempted." (sect. 10, para. 7) In fact, the entire field of *genetic algorithms* (Goldberg 1989; Holland 1975; Schaffer 1989) and much of *artificial life* research (Langton 1989) rely on computer instantiations of evolutionary dynamics to produce adaptive solutions to specified problems – often solutions in the form of neural or psychological mechanisms underlying behavioral strategies.

Our research, for example, uses genetic algorithms to simulate the evolution of neural networks that control the behavior of simple organisms in simple virtual environments (Miller & Todd 1990; Miller et al. 1989; Todd & Miller, in press). Ackley (1990) has produced a more complex and suggestive model of the evolution of adaptively behaving creatures using a similar approach. In these models, adaptive pressures are not explicitly represented, but emerge from the dynamics of the environment and the interactive behavior of the simulated organisms. In all such methods, the evolutionary process itself *is* the search for optimal behavioral strategies. Although no global optimum is guaranteed to exist or to be findable in finite time, genetic algorithms have generally proven superior to any other search method for very large, complex search spaces with many local optima (Goldberg 1989).

We sympathize with the desires of Houston & McNamara (1988a) and Clark to develop computational tools for analyzing the adaptive functions of behaviors, but we are pessimistic about the ability of any simulation method to represent directly the manifold adaptive pressures that emerge from even moderately complex ecosystems. Rather, we believe that adaptive pressures can be best understood indirectly, by setting up environments, simulating an evolutionary process to produce psychological and behavioral adaptations to those environments, and comparing the resulting adaptations and behaviors to those observed in real organisms. Dynamic programming represents an attempt to understand the results of evolution without simulating evolution. But we believe that evolution can take care of its own. Simulating evolution via genetic algorithms can automatically register the differential selection of genes and gene complexes through the phenomenon Holland (1975) calls *intrinsic parallelism,* and can include the effects of kin selection and inclusive fitness.

Furthermore, through the application of our genetic algorithm to the evolution of behavioral-producing neural networks our models of adaptive psychological mechanisms can incorporate the two main advantages of the dynamic programming approach: first, the use of *evolved,* not prespecified, internal state variables in the generation of behavior (via recurrent patterns of network activation – see Elman 1988), and second, the production of ongoing dynamic behavioral sequences (Jordan 1986). Moreover, our method includes the further biologically relevant characteristics of a powerful set of learning mechanisms (Rumelhart & McClelland 1986) and the ability of networks to generalize adaptively to novel environmental situations (a crucial adaptive capability – see Shepard 1987), obviating the need for an exhaustive dynamic programming search of state-space.

If you want to model what comes out of the process of evolution (behavioral adaptations) in terms of what goes in (adaptive pressures) then why not model the process itself? The growing number of researchers using genetic algorithms answer "why not, indeed?" Genetic algorithms are transparently analogous to natural selection, applying concrete environmental and social effects to genotypically coded populations of organisms which evolve forward in time, thus performing computationally efficient searches for adaptive responses to emergent adaptive pressures. As such, they are an intuitively appealing, understandable, and tractable approach to modeling behavioral adaptations. Respect for the complexity of natural behavior demands respect for the adaptive process, natural selection, which produced that complexity. And instantiating that process in our models is the highest respect we can offer.