General Article

THE REWARDS AND HAZARDS OF COMPUTER SIMULATIONS

By Stephan Lewandowsky

Although computer simulations and other modeling tools have assumed a pivotal role in cognitive psychology, their utility continues to be questioned by some psychologists. This article presents several examples that illustrate both rewards and potential hazards associated with the simulation approach. Simulations can provide the formal framework necessary to disambiguate new ideas, they can explore the implications of complex models, and they can predict seemingly counterintuitive findings or uncover hidden relationships. At the same time, care must be taken to avoid pitfalls that may arise when computer code inadvertently differs from the intended specifications of a theory, or when predictions derive not from fundamental properties of a theory but from pragmatic choices made by the modeler.

The simulation of human behavior by computers has become a burgeoning enterprise: As early as 1971, Starbuck and Dutton compiled a bibliography encompassing 1,921 articles concerned with computer simulations. That number appears particularly staggering because it dates from a period in which only a handful of psychologists were conversant with simulation techniques. During the two succeeding decades, the number of psychologists conducting simulations has increased considerably, and computer simulations are now nearly as common as experiments in some parts of the literature. However, this trend has not been without opposition, and demands have recently been made for the complete elimination of memory models or theoretical constructs of any kind (Watkins, 1990). That position, in turn, has not gone unchallenged and has been countered by Hintzman (1991), who illustrated the utility of models by pointing mainly to known limitations of human reasoning, and how those could be circumvented by the proper use of models.

Although stimulated by recent controversy, this article was not written to continue the debate between modelers and behavioral empiricists. Instead, the article illustrates the rewards offered by computer simulations, as well as their potential hazards, by bringing together a number of relevant cases that both sides of the debate might find informative.

Address correspondence to Stephan Lewandowsky, Department of Psychology, University of Oklahoma, Norman, OK 73019-0535; e-mail: lewan@uokmax.ecn.uoknor.edu.

REWARDS AND PROGRESS

Agreement on what constitutes progress in cognitive psychology has remained about as elusive as agreement about the proper role of models and theories. How should one report the contributions of simulations when judgments about the tool's achievements are as diverse as judgments concerning the tool itself? A criterion for progress put forward recently (Lewandowsky & Hockley, 1991) may serve well here. It considers the extent to which theory and data have become interrelated, that is, the extent to which theory generates data gathering and the extent to which data, in turn, constrain theory building. For several areas of inquiry, a historical segregation between theory and data has been overcome and has been replaced by the desired constant interaction. By that criterion, which assigns neither data nor theory the pivotal role that some scholars would hesitate to accept, a variety of cases in which simulations played a constructive role can be cited.

Precise Specification of Models

On the surface, it may appear useful to differentiate between two classes of models: those simple enough to be understood and tested without formal work and those sufficiently complex to require computer simulations or other modeling exercises. Thus, the argument might go, one can safely ignore simulations if one only sticks to the former class of models. This reasoning ignores the problem that even deceptively simple models can benefit from the rigor of simulations—although often that benefit takes the form of a simulation revealing the logical incoherence of a verbal model; the benefit, then, is to science and not the model's creator.

An illustration of these interesting, but usually nonpublic, elimination and refinement processes is provided in a dissertation by McDonald (1980) concerning word recognition. Models of the word recognition process must account for a large body of literature (reviewed by Carr & Pollatsek, 1985) and, in particular, the two consistent findings that a prior related item speeds processing (semantic priming) and that more familiar words are recognized more quickly than less familiar words (the word frequency effect). Early versions of a lexical instance model (Morton, 1969) proposed that each word is instantiated by its own unique detector, or "logogen," and visual feature analyzers feed information extracted from the stimulus to all such detectors in parallel. According to this model, evidence is accumulated by a detector until its activation exceeds a threshold and the corresponding word is recognized. The logogen model accounts for priming by postulating that semantic characteristics of previous words are fed back into the detector array, thus raising the activation of related items above their resting levels. Frequency effects are implemented in a related fashion, by postulating that familiar words have lower thresholds than their low-frequency counterparts.

McDonald (1980) first addressed one of the major difficulties of the logogen model. Its prediction that stimulus degradation should affect both priming and word frequency in the same way had appeared to be disconfirmed when only priming—but not word frequency—was found to interact with stimulus degradation (Becker & Killion, 1977). However, the impact of these behavioral data on the logogen model was contingent upon acceptance of an additive factors interpretation that is known to be problematical (Luce, 1986, pp. 473–491).[1] McDonald's initial simulations of the logogen model therefore examined the effect of changes in the feature sampling rate (a simulation mechanism analogous to the presentation of a degraded stimulus) on priming and frequency effects. The observed interactivity in both cases confirmed the troubling implications of Becker and Killion's results, but did so without relying on the problematic additive factors interpretation.

Having confirmed these difficulties, McDonald turned to further simulation-driven development of an alternative verification model that, at a verbal level, had been specified earlier (Becker, Schvaneveldt, & Gomez, 1973). In the verification model, the output from logogen-like detectors produced a set of candidate items that was then edited by a serial verification process. The model was said to handle the observed additivity between word frequency and stimulus degradation because word frequency affected the order in which candidates were considered for verification but not their activation level. Priming, in contrast, was handled as in the logogen model, by feeding back information from previous items into the array of detectors, thus boosting the activation of related words while also producing the observed interaction with stimulus degradation. During simulations of this initial version, McDonald encountered a difficulty concerning the relative size of the priming and frequency effects. When the magnitude of the priming activation boost was adjusted to give reasonable amounts of facilitation, only very small frequency effects were obtained. When the boost was further increased to obtain reasonable frequency effects, benefits of priming were extinguished. Thus, although the model had appeared, at a verbal level, to account for the data that propelled its initial development, the simulations revealed otherwise.

A revision of the model restored the correct relative magnitude of priming and word frequency effects by reversing the details of the verification process: The early, verbal, version had postulated that as sensory evidence accumulated, candidates were added to the verification set. When all evidence had been extracted, the set reached its maximum size and was verified, by bringing to bear top-down expectations, in order of word frequency. The revised model reversed the growth process of the verification set into a shrinking process. That is, upon stimulus presentation, the initial verification set included all items in the lexicon, and as sensory evidence accumulated, implausible candidates were eliminated. In conjunction with several other modifications (e.g., the presence of two search sets, one based on semantic relatedness and the other on featural evidence), this version simulated both priming and frequency effects successfully.

What is to be learned from McDonald's model development process? Perhaps most important, it shows that simulation of a verbal model can reveal previously hidden insufficiencies, and that simulations allow for experimentation and modification until known empirical benchmarks are accommodated. Moreover, in the context of the criterion for progress given earlier, one must note the later empirical success of the verification model, when its predictions were found to correlate with human performance for some 900 individual words and pseudowords (Paap, Newsome, McDonald, & Schvaneveldt, 1982). It is no small feat to predict responses at the item level, and it appears doubtful that a verbal model would ever be able to do so.

Exploration of Novel Ideas or Complex Models

In a related fashion, simulations can also be used to explore attractive ideas or complex models. Simulations can be of value in this way either because a seemingly attractive idea might otherwise be too unconstrained to support predictions and tests or because a complex

1. In the additive factors approach, experimental variables that are found to interact are assumed to affect the same stage of processing, whereas those that are additive affect different stages. A major difficulty with this interpretation is that it must assume strict seriality among processing stages (see Luce, 1986, for further discussion).

model may resist analytic exploration even though it is fully specified mathematically.

Consider a case in which a simple but seemingly untenable idea was revitalized on the basis of simulation results. How would one represent and recall a serial list of items A, B, C, and D? Perhaps the simplest way is by *chaining*, that is, by pair-wise associations between list items, which are then used for serial retrieval. An obvious problem with this scheme is that a break in the chain, such as the failure to retrieve C, makes all subsequent items inaccessible for lack of proper cue. The pervasive recency effect in serial recall, in conjunction with other damaging data, therefore led to the rapid demise of the notion of chaining, which remained dormant until a series of simulations by Lewandowsky and Murdock (1989) shed new light on the issue. Their distributed memory model (TODAM; Murdock, 1983) was based on the chaining notion but, unlike previous incarnations of that idea, specified how an overt break in the chain (e.g., failure to recall C) could nonetheless engender recency. In this model, cuing of memory would always retrieve an approximation to the target response (e.g., cuing with B would retrieve the approximation C′), and even if that approximation was too "blurry" to support an overt response (recall of C), it could still be used as a cue for further items, thus providing the opportunity for the correct recall of later list items.

An example of the exploration of complex models can be found in a recent simulation of a neural network that confirmed the existence of an alternative to the discrete "lexicon" (a repository of all known word forms) embodied in most word recognition models. Seidenberg and McClelland (1989) showed that word recognition and pronunciation—for example, reading of *pint* or *lint*—could take place in the absence of a traditional lexicon of discrete, localized representations. In Seidenberg and McClelland's model, information about words was instead distributed across many different units ("weights" in neural network parlance) in memory, such that *pint* and *lint* would be represented by the same ensemble of weights. Because responses were determined by the interaction between these weights and the overall feature pattern of an item, very different pronunciations could be learned even for words with a high degree of surface similarity. Although the mathematics of the network are well understood (Rumelhart, Hinton, & Williams, 1986), it was a simulation that explored the notion of distributed representation to the point where it became a satisfactory alternative to the firmly entrenched lexicon models.

Seidenberg and McClelland's (1989) word recognition model also illustrates that simulations can be subject to public scrutiny and need not be the private enterprise they are sometimes assumed to be. Besner, Twilley, McCann, and Seergobin (1990) analyzed and extended the original simulations and concluded that the network was unable to generalize properly. That is, although the network had learned to pronounce and recognize words from a large corpus, including some nonwords, Besner et al. showed that the simulations failed to produce better-than-chance lexical decision performance on new nonwords and failed to produce various other effects shown by humans (e.g., the pseudohomophone effect). Besner et al.'s article, in turn, was followed by a rejoinder by Seidenberg and McClelland (1990); in this context, it is not the full details of the debate that are of interest, but the fact that it took place, thus underscoring the public nature of simulations. And lest one think that this type of debate is confined to esoteric, purely theoretical, exchanges between modelers, it should be noted that part of the response to Seidenberg and McClelland was empirical (Besner, 1990). The close interrelation between simulations and behavioral experimentation in this example fits well within the criterion for progress mentioned earlier.

Serendipity and Simulations

A widespread opinion among critics is that theories or simulations somehow stand in the way of serendipitous discovery: "The discoveries of penicillin, X-rays, and America have apparently failed to alert students of memory to the possibility of serendipitous findings within their own field" (Watkins, 1990, p. 333). Similarly, Lockhart (1991) challenged modelers to produce a list of important phenomena or conceptual advances that owe their existence to the powers of formal models. It takes little effort to begin compilation of that list.

Consider an experiment in which subjects learn to associate pairs of unrelated words, such as *grass* and *city*. It appears to be a matter of common sense to expect that these associations, like all other memories, will be subject to forgetting once learning ceases. It thus appeared rather curious that modeling of some classic paired-associate learning data with a distributed memory model required an unreasonably high value of a retention parameter (Murdock, 1989). That is, in contrast to commonsense expectations, the model had to presume little or no forgetting in order to describe the acquisition of long lists of paired associates. Several experiments have since established that, in a continuous recognition paradigm, there is only limited forgetting of associative information, and that performance appears to reach asymptote after a few intervening pairs (Hockley, 1991; Murdock & Hockley, 1989). Two points can be made about these data: First, they contradict commonsense expectations. Second, they would not have been collected had it not been for a model parameter assuming a peculiar value.

In a related vein, models can sometimes bridge the gap between seemingly unrelated sets of observations by specifying a single underlying process, much as desired by Estes (1975, p. 271). A particularly compelling example involves the relationship between recognition and classification, reported by Nosofsky (1991). In his experiments, subjects were trained to classify schematic faces into two categories, using a small set of exemplars. During a subsequent transfer phase, subjects were tested on a larger set of faces, including both old and new stimuli. For each stimulus, subjects had to decide to which of the study categories it belonged (classification) and whether it had been presented during training (recognition). Figure 1 shows the empirical relationship between recognition and classification obtained in Nosofsky's Experiment 1A.

The low correlation between those two tasks ($r = .36$) is apparent from the figure, and might give rise to the speculation that different, possibly independent, processes underlie recognition and classification. However, it turns out that a common model, based on analysis of the similarity between the test item and all instances stored in memory, can simultaneously account for both recognition and classification. This is shown in the two panels of Figure 2, which relate the performance predicted by Nosofsky's model to the observed data for both tasks.

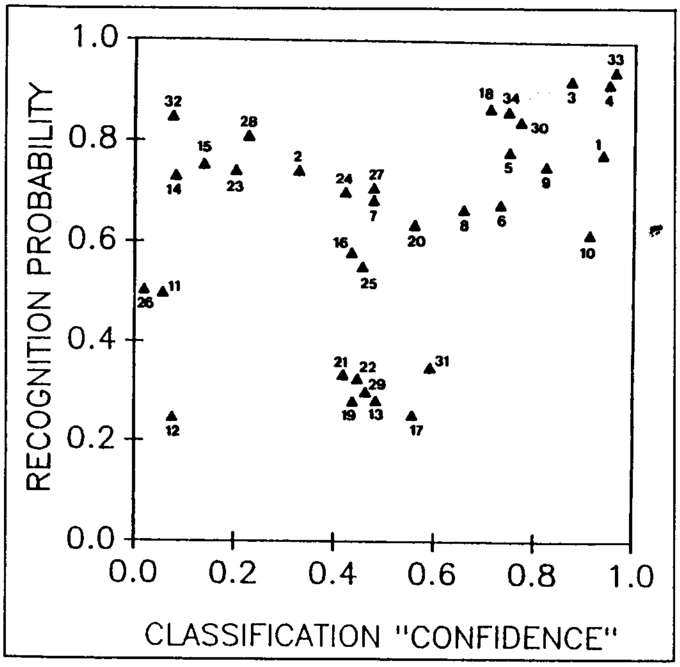

Fig. 1. Observed recognition probabilities as a function of observed classification confidences for the same set of stimuli (from Nosofsky's, 1991, Experiment 1A). Note the low apparent correlation between the two tasks. Reprinted by permission.

Nosofsky's (1991) conclusions are noteworthy: "The initial scatterplot . . . revealed little relation between classification and recognition performance. At that limited level of analysis, one might have concluded that there was little in common between the fundamental processes of classification and recognition. Under the guidance of the formal model, however, a unified account of these processes is achieved . . ." (p. 9).[2]

This success is not an isolated occurrence: Metcalfe (1993) showed that feeling-of-knowing judgments, release from proactive inhibition, and Korsakoff's syndrome can all be described by a common process embodied in her composite holographic memory model. Interestingly, that success derived from a property of the model initially deemed to be a serious problem. Metcalfe discovered, by way of simulation, that the storage of a large number of items increased the variance[3] of the common memory trace so rapidly—especially if the items were related to each other—that the system soon reached a saturation point. Exploration of possible solutions led to the development of a monitoring and control mechanism, whereby the weight given to each new item depended on its anticipated contribution to an increase in overall variance. The monitoring part of that mechanism, which assessed the novelty of a to-be-encoded event, turned out to provide a natural account of some of the intricacies of people's feeling-of-knowing judgments (e.g., that items producing errors of commission are given higher familiarity ratings than those leading to omissions, that priming of cues increases familiarity ratings whereas priming of targets does not). Further simulations revealed that the same monitoring and control mechanism enabled the model to show buildup and release from proactive inhibition, and that the model behaved much like a Korsakoff's patient when the mechanism was disabled. Thus, by using simulations to explore the implications of a potentially troublesome property of her model, and by proposing a solution on the basis of further simulations, Metcalfe (1993) was able to provide a common theoretical underpinning for feeling-of-knowing judgments, proactive inhibition, and Korsakoff's syndrome. It is unclear how an account relating phenomena as diverse as these

2. It must be noted that these two examples (Murdock & Hockley, 1989; Nosofsky, 1991) involved application of the models not by simulation but by analytic fit. It is important to bear in mind that simulations differ from analytic models mainly in the way in which predictions are generated, and that both tools can be applied to similar problems in similar ways. A more extensive comparison of these two approaches and an exploration of related issues can be found in Gregg and Simon (1967). Complete introductions to simulation techniques are provided by Lehmann (1977) and Dutton and Starbuck (1971).

3. Variance here refers to the sum of the squared activation values of all elements, or weights, in the memory trace.

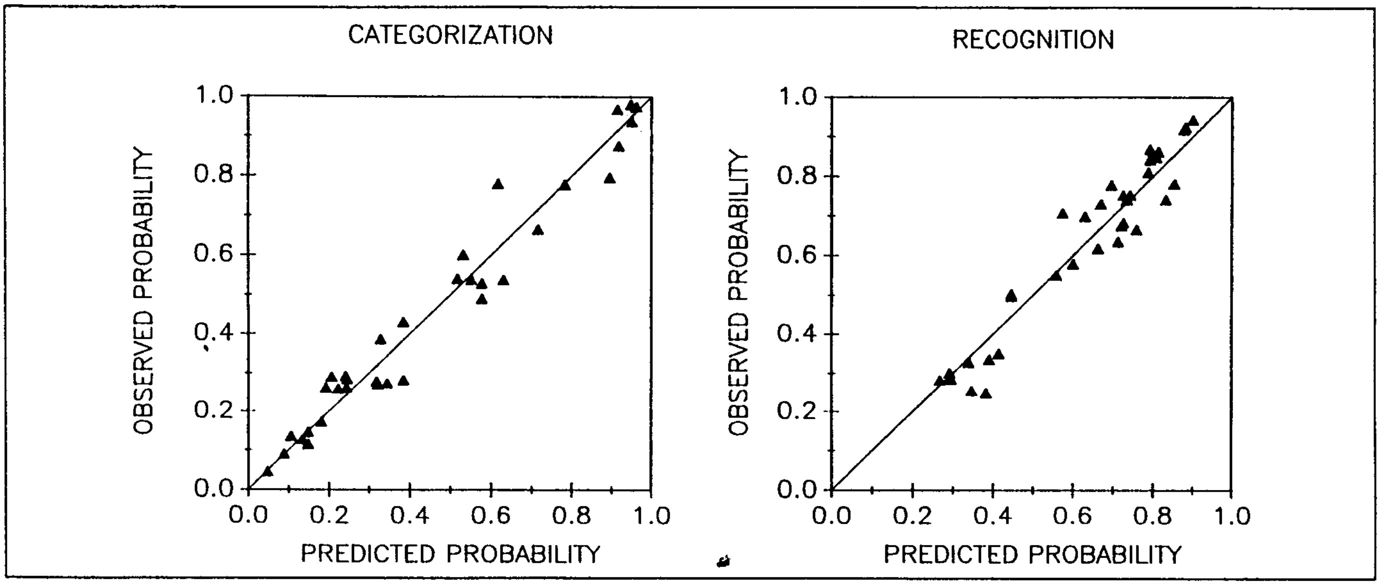

Fig. 2. Observed and predicted classification probabilities (left panel) and observed and predicted recognition probabilities (right panel). Perfect prediction would correspond to all points falling on the diagonal; the model approaches that closely, with more than 90% of the variance accounted for in both panels. From Nosofsky's (1991) Experiment 1A. Reprinted by permission.

three could be provided by experimentation and verbal theorizing alone.

HAZARDS AND SETBACKS

Having considered several cases in which simulations offered new insights into cognitive processes or engendered the collection of new behavioral data, we must now turn attention to potential pitfalls awaiting the simulation modeler. Three principal classes of problems can be identified, concerning the relation between computer program and theory, the match between real-world situations and simulated environments, and the nature of simulation results.

Not-So-Irrelevant Specifications

Implementation of an existing theory in a simulation—as described earlier with the verification model—is a step-by-step process that often involves pragmatic decisions to bridge the gap between the loose level of verbal theorizing and the tight level of description required for a program. These pragmatic decisions, in turn, may lead to inadvertent discrepancies between theory and simulation (Frijda, 1971; Reitman, 1965, p. 25). This *irrelevant-specification problem* is potentially serious because the simulation results no longer speak to properties of the theory.

The nature of the problem renders it difficult to find relevant published examples, so I discuss my own work as an illustration here. Consider the deblurring stage of retrieval that is common to most neural networks and distributed memory models. Typically, the output from the model is compared with a set of possible responses and the best match (i.e., smallest discrepancy between output and response) is chosen (Hinton & Shallice, 1991; Lewandowsky & Murdock, 1989; Seidenberg & McClelland, 1989). Not only is there some uncertainty concerning how that match should be measured (Goebel & Lewandowsky, 1991), but the comparison process also introduces a particularly pernicious problem that is usually resolved by the programmer in an ad hoc manner: What is the "best" match if two of the possible responses are identical? In the program, the comparison process must be serial, and a pointer is updated whenever the best match is improved by the next item. This must involve a conditional statement with either a $<$ or a $\leq$ comparison operator. If the former is used, the model will respond with the first of the two identical items to be compared. If the latter is used, the response will be the item that happens to be compared second. This matters only in the unique situation when (a) a serial list is to be recalled by probing from one item to the next, (b) that list contains a repeated item, and (c) the model's output is compared only to the remaining list items. These conditions are met when the Lewandowsky and Murdock (1989) chaining model is presented with a list such as A B C D C E. With the $<$ operator, the model can recall all items on the list, similarly to what human subjects would do. With the $\leq$ operator, the model recalls A B C E but can never produce the intervening D and C. The theory (TODAM; Murdock, 1983) is mute on the choice of operator, stating only that the "best match" should be obtained.

It has been argued that this type of irrelevant specification presents less of an issue in contemporary unified theories, whose larger scope comes closer to specifying all the details of the simulation, and which from the moment of their inception are implemented in a simulation program (Newell, 1990, p. 23). Hence, simulations of unified theories are not written to implement a model; rather, the model is the simulation (see also Feigenbaum, 1963). However, even then, irrelevant specifications can be avoided only if the model is communicated and discussed exclusively by referring to the programming code. Any verbal description, no matter how faithful and precise, likely opens the door to irrelevant specifications. At best, one can follow Frijda's (1971) advice to "describe the relevant processes unambiguously by naming the subroutines concerned, by stating their precise input and output conditions, the conditions of their activation, and the transformations they achieve" (p. 612).[4]

Overparameterization

A related difficulty arises when the simulation is inadvertently provided with "friendly" input conditions. Results may then derive from these extraneous conditions as opposed to fundamental properties of the theory.

To illustrate, consider Rumelhart and McClelland's (1986) neural network that simulated children's acquisition of past tense. That acquisition process consists of two distinct phases: During the first phase, children use a variety of correct past tense forms for both regular (*walk → walked*) and irregular (*go → went*) verbs. During the second stage, children tend to overregularize, leading to forms such as *go → goed* or *eat → eated*. It is only later that children reacquire the ability to produce the correct past tense for both regular and irregular verbs. The top panel of Figure 3 shows Rumelhart and McClelland's simulation results, which resemble the linguistic processes observed in children.

That simulation was thoroughly scrutinized by Pinker and Prince (1988). Their critique consisted of a conceptual component, not relevant here, and an analysis of the input conditions for Rumelhart and McClelland's simulations, which is reproduced in the bottom panel of Figure 3. The panel shows the proportion of input items consist-

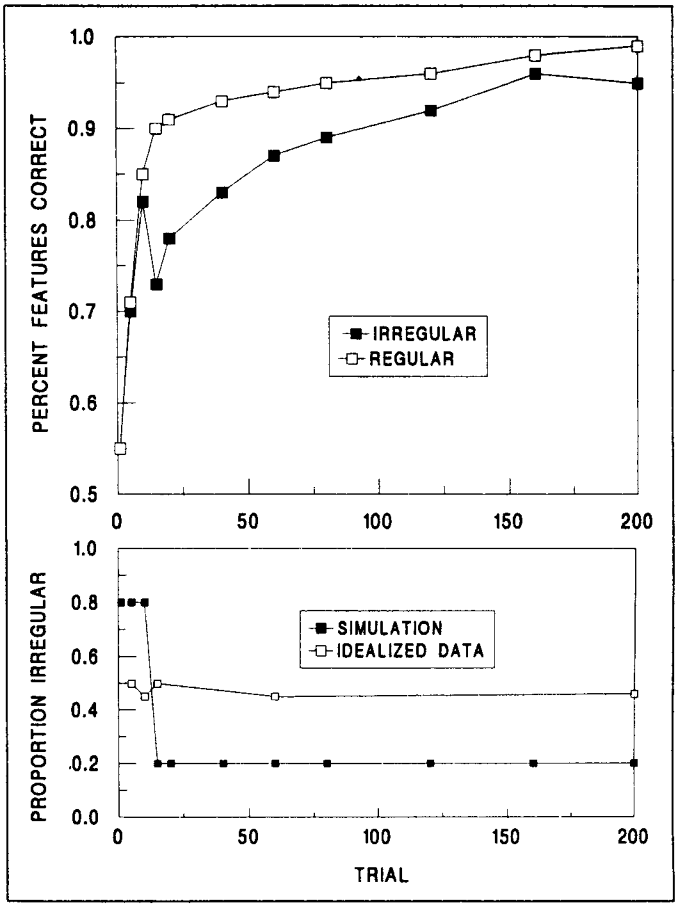

Fig. 3. Analysis of Rumelhart and McClelland's (1986) network simulating children's acquisition of past tense. The top panel shows simulated acquisition of past tense for regular and irregular verbs (adapted from Rumelhart & McClelland, 1986). The bottom panel shows proportion of irregular verbs in the input corpus (simulation) together with proportion actually found in children's vocabulary (idealized data). Data represent average of the values reported by Pinker and Prince (1988, Table 1).

ing of irregular verbs. For the first 10 simulation trials, 80% of the input was irregular, compared with 20% on the remaining trials. Comparison of the panels shows that the simulated drop in performance for irregular verbs coincided with that change in input conditions. Pinker and Prince concluded that "the model's shift from correct to overregularized forms does not emerge from any endogenous process; it is driven directly by shifts in the environment" (p. 138). Moreover, these simulated environmental shifts did not correspond to the actual composition of children's vocabulary, which tends to be evenly divided into regular and irregular verbs.

One must hasten to add that Rumelhart and McClelland (1986) chose their input conditions to resemble the presumed learning environment experienced by children: The initial set of stimuli contained the highest frequency verbs in English (which happen to be mostly irregular),

4. It has been argued that, even with greatest care, final identification of the proper model is impossible because "of the poverty of the data relative to the complexity of the implementation level theories" (Anderson, 1990, p. 21). An opposing view holds that the existing data base suffices to constrain theory building (Newell, 1990).

and the subsequent expanded set consisted mainly of medium-frequency items (fewer of which are irregular). Hence, the environmental shift was an accidental—but ultimately critical—by-product of otherwise plausible assumptions about children's linguistic environment.

The Nature of Simulation Results

What kind of results can be expected from a simulation? Because a computer program is, by definition, entirely deterministic, can we ever be truly surprised by an unexpected outcome, as might happen in behavioral experiments? Early opinions denied this possibility, preferring the view that one can get out of a simulation only what the programmer has put in (Reitman, 1965, p. 15). The counterargument, that simulations are capable of producing truly counterintuitive results (Simon, 1969, p. 15), has been supported by recent simulations that have generated novel insights in nontrivial ways. Consider, for example, Seidenberg and McClelland's (1989) network model, in which it had been entirely unclear, prior to conducting the simulation, whether the proper pronunciation patterns could be learned.

Although potentially valuable, the possibility of unexpected results opens the door for *Bonini's paradox,* which arises when the simulation turns out to be no easier to understand than the real-world processes it was supposed to illustrate (Dutton & Starbuck, 1971). It has been argued that this problem is endemic in current neural networks, and that they therefore cannot provide true explanations for the processes they purport to describe. (McCloskey, 1991, examines the issue in greater depth than is possible here.)

It appears, then, that simulations can be criticized either for not producing unforeseen results or, if they do, for being inscrutable. In response, one must point to the opportunity for experimentation afforded by successful simulation models (Dutton & Starbuck, 1971). McCloskey (1991) offered an analogy to animal models of human cognition: Neither a simulation nor an animal model in itself is an explanation, but both can point to explanations by allowing experimental manipulations not possible with human subjects.

Consider, as an illustration, Hinton and Shallice's (1991) "lesion" experiments involving an attractor network. Briefly, their net was trained to map a set of orthographic representations into semantic features, so that presentation of a spelling pattern would activate the correct "word" at a semantic output level. Subsequent "lesioning" variously involved the removal of units, the contamination of weights with noise, or the setting to zero of a set of connections. Among the most interesting findings was the persistent co-occurrence of visual (*cat* read as *mat*) and semantic (*peach* → *apricot*) errors upon lesioning of virtually any part of the same semantic pathway. Hinton and Shallice identified this result to reflect a general "breakdown characteristic of a network containing attractors when lesioned in various places" (p. 89). That generality, in turn, elegantly explained why this mix of errors (visual vs. semantic) remains remarkably constant even though other performance deficits show great variability across patients. Further experimentation established that the attractor net predicted several other intricacies of the error pattern that had previously escaped explanation by the standard dual-route models (e.g., that of Marshall & Newcombe, 1973).

CONCLUSIONS

This article cites examples that support the following conclusions. (1) Computer simulations have contributed to cognitive psychology by providing precise formulations for new ideas and verbal models and by exploring complex theories. (2) Simulations are subject to replication, inspection, and scrutiny in ways similar to traditional experiments. (3) Simulations can uncover relations between seemingly disparate findings. (4) Simulations allow for experimentation that may not be possible with human subjects.

Yet several hazards and limitations must be recognized: (1) Inadvertent discrepancies between a theory and its implementation in a simulation may exist. (2) Care must be taken to avoid "friendly" input that gives rise to the desired result. (3) Simulation results cannot always substitute for a conceptual level of explanation.

Acknowledgments—This article was made possible by the continued criticism and encouragement provided by J. Vivien Bainbridge, John T. Behrens, Francis T. Durso, Charles F. Gettys, Scott D. Gronlund, William E. Hockley, and Bennet B. Murdock. The author wishes to thank James E. McDonald for making available his dissertation and Robert M. Nosofsky for providing figures from one of his articles. I also thank two anonymous reviewers for their detailed and constructive comments on an earlier version.

REFERENCES

Anderson, J.R. (1990). *The adaptive character of thought.* Hillsdale, NJ: Erlbaum.

Becker, C.A., & Killion, T.H. (1977). Interaction of visual and cognitive effects in word recognition. *Journal of Experimental Psychology: Human Perception and Performance, 3,* 389–401.

Becker, C.A., Schvaneveldt, R.W., & Gomez, L. (1973, November). *Semantic, graphemic, and phonetic factors in word recognition.* Paper presented at the annual meeting of the Psychonomic Society, St. Louis.

Besner, D. (1990, November). *Word recognition and connectionism: More words, definitely a lexicon.* Paper presented at the annual meeting of the Psychonomic Society, New Orleans.

Besner, D., Twilley, L., McCann, R.S., & Seergobin, K. (1990). On the association between connectionism and data: Are a few words necessary? *Psychological Review, 97,* 432–446.

Carr, T.H., & Pollatsek, A. (1985). Recognizing printed words: A look at current models. *Reading Research: Advances in Theory and Practice, 5,* 1–82.

Dutton, J.M., & Starbuck, W.H. (Eds.). (1971). *Computer simulation of human behavior.* New York: Wiley.

Estes, W.K. (1975). Some targets for mathematical psychology. *Journal of Mathematical Psychology, 12*, 263–282.
Feigenbaum, E.A. (1963). Simulation of verbal learning behavior. In E.A. Feigenbaum & J. Feldman (Eds.), *Computers and thought* (pp. 297–309). New York: McGraw-Hill.
Frijda, N.H. (1971). Problems of computer simulation. In J.M. Dutton & W.H. Starbuck (Eds.), *Computer simulation of human behavior* (pp. 610–618). New York: Wiley.
Goebel, R.P., & Lewandowsky, S. (1991). Retrieval measures in distributed memory models. In W.E. Hockley & S. Lewandowsky (Eds.), *Relating theory and data: Essays on human memory in honor of Bennet B. Murdock* (pp. 509–527). Hillsdale, NJ: Erlbaum.
Gregg, L.W., & Simon, H.A. (1967). Process models and stochastic theories of simple concept formation. *Journal of Mathematical Psychology, 4*, 246–276.
Hinton, G.E., & Shallice, T. (1991). Lesioning an attractor network: Investigations of acquired dyslexia. *Psychological Review, 98*, 74–95.
Hintzman, D.L. (1991). Why are formal models useful in psychology? In W.E. Hockley & S. Lewandowsky (Eds.), *Relating theory and data: Essays on human memory in honor of Bennet B. Murdock* (pp. 39–56). Hillsdale, NJ: Erlbaum.
Hockley, W.E. (1991). Recognition memory for item and associative information: A comparison of forgetting rates. In W.E. Hockley & S. Lewandowsky (Eds.), *Relating theory and data: Essays on human memory in honor of Bennet B. Murdock* (pp. 227–248). Hillsdale, NJ: Erlbaum.
Lehmann, R.S. (1977). *Computer simulation and modeling: An introduction.* Hillsdale, NJ: Erlbaum.
Lewandowsky, S., & Hockley, W.E. (1991). Relating theory and data: Towards an integration. In W.E. Hockley & S. Lewandowsky (Eds.), *Relating theory and data: Essays on human memory in honor of Bennet B. Murdock* (pp. 3–19). Hillsdale, NJ: Erlbaum.
Lewandowsky, S., & Murdock, B.B., Jr. (1989). Memory for serial order. *Psychological Review, 96*, 25–57.
Lockhart, R.S. (1991). Varieties of theoretical perspectives. In W.E. Hockley & S. Lewandowsky (Eds.), *Relating theory and data: Essays on human memory in honor of Bennet B. Murdock* (pp. 101–107). Hillsdale, NJ: Erlbaum.
Luce, R.D. (1986). *Response times.* New York: Oxford University Press.
Marshall, J.C., & Newcombe, F. (1973). Patterns of paralexia: A psycholinguistic approach. *Journal of Psycholinguistic Research, 4*, 169–176.
McCloskey, M. (1991). Networks and theories: The place of connectionism in cognitive science. *Psychological Science, 2*, 387–395.
McDonald, J.E. (1980). *An information processing analysis of word recognition.* Unpublished doctoral dissertation, New Mexico State University, Las Cruces.
Metcalfe, J. (1993). Novelty monitoring, metacognition, and control in a composite holographic associative recall model: Implications for Korsakoff amnesia. *Psychological Review, 100*, 3–22.
Morton, J. (1969). Interaction of information in word recognition. *Psychological Review, 76*, 165–178.
Murdock, B.B., Jr. (1983). A distributed memory for serial-order information. *Psychological Review, 90*, 316–338.
Murdock, B.B., Jr. (1989). Learning in a distributed memory model. In C. Izawa (Ed.), *Current issues in cognitive processes: The Tulane Floweree Symposium on Cognition* (pp. 69–106). Hillsdale, NJ: Erlbaum.
Murdock, B.B., & Hockley, W.E. (1989). Short-term memory for associations. In G.H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 24, pp. 71–108). San Diego: Academic Press.
Newell, A. (1990). *Unified theories of cognition.* Cambridge, MA: Harvard University Press.
Nosofsky, R.M. (1991). Tests of an exemplar model for relating perceptual classification and recognition memory. *Journal of Experimental Psychology: Human Perception and Performance, 17*, 3–27.
Paap, K.R., Newsome, S.L., McDonald, J.E., & Schvaneveldt, R.W. (1982). An activation-verification model for letter and word recognition: The word superiority effect. *Psychological Review, 89*, 573–594.
Pinker, S., & Prince, A. (1988). On language and connectionism: Analysis of a parallel distributed processing model of language acquisition. *Cognition, 28*, 73–193.
Reitman, W.B. (1965). *Cognition and thought.* New York: Wiley.
Rumelhart, D.E., Hinton, G.E., & Williams, R.J. (1986). Learning internal representations by error propagation. In D.E. Rumelhart, J.L. McClelland, & the PDP Research Group (Eds.), *Parallel distributed processing* (Vol. 1, pp. 318–362). Cambridge, MA: MIT Press.
Rumelhart, D.E., & McClelland, J.L. (1986). On learning the past tense of English verbs. In J.L. McClelland, D.E. Rumelhart, & the PDP Research Group (Eds.), *Parallel distributed processing* (Vol. 2, pp. 216–271). Cambridge, MA: MIT Press.
Seidenberg, M.S., & McClelland, J.L. (1989). A distributed, developmental model of word recognition and naming. *Psychological Review, 96*, 523–568.
Seidenberg, M.S., & McClelland, J.L. (1990). More words but still no lexicon: Reply to Besner et al. (1990). *Psychological Review, 97*, 447–452.
Simon, H.A. (1969). *The sciences of the artificial.* Cambridge, MA: MIT Press.
Starbuck, W.H., & Dutton, J.M. (1971). The history of simulation models. In J.M. Dutton & W.H. Starbuck (Eds.), *Computer simulation of human behavior* (pp. 9–102). New York: Wiley.
Watkins, M.J. (1990). Mediationism and the obfuscation of memory. *American Psychologist, 45*, 328–335.