

# Computational Cognitive Modeling, the Source of Power, and Other Related Issues

*Ron Sun and Charles X. Ling*

■ In computational cognitive modeling, we hypothesize internal mental processes of human cognitive activities and express such activities by computer programs. Such computational models often consist of many components and aspects. Claims are often made that certain aspects play a key role in modeling, but such claims are sometimes not well justified or explored. In this article, we first review some fundamental distinctions and issues in computational modeling. We then discuss, in principle, systematic ways of identifying the source of power in models.

The workshop entitled "Computational Modeling: The Source of Power," which we cochaired and organized, was held at the Thirteenth National Conference on Artificial Intelligence (AAAI-96) on 5 August 1996 in Portland, Oregon. In this well-attended workshop, 14 talks and a panel discussion were given on the nature and fundamental issues of computational cognitive modeling and on the source of power in models. Although viewpoints differed from speaker to speaker, some common themes emerged from the presentations and discussions. The difficulties that cognitive modelers face and the identification of sources of power in cognitive models were especially prominent issues in the discussions. To summarize and continue the theme of the workshop, in this article, we highlight a few important issues in cognitive modeling, including the reasons why computational models work, measures of success, and systematic ways of identifying the source of power in models.

Computational cognitive modeling is an important aspect in cognitive science because it plays a central role

in the computational understanding of the mind. A common methodology of cognitive science is to express a theory about human cognition in a computer program and compare the program's behavior with human cognitive behavior. However, many fundamental questions remain unanswered: To what extent will we match the behaviors of humans and computer models? How do we measure the success of a model? How do we identify the source of such a success? We hope that the workshop and this article will stimulate more discussions and that a better understanding will be reached in the near future.

Cognitive modeling has traditionally been underrepresented in American Association for Artificial Intelligence conferences and journals. However, its importance should not be underestimated. It could provide key insights into understanding intelligence and, thus, provide great impetus for the further development of AI. In the early days of AI, models and systems were closely tied to the study of cognition. Although this tradition has been kept at a few places, more connections between AI and other disciplines such as psychology are certainly necessary, even at the present time of specialization and technical sophistication in AI.

## Types of Cognitive Model

Cognitive models can be of two major types: One type consists of (to some extent) detailed computational process models (or AI models), and the other consists of (behavioral) mathematical models. The former seeks to capture internal computational (mechanistic) processes that generate overt cognitive behavior (for example,

Anderson [1993, 1982] and Sun [1995]). The latter works through measuring a number of behavioral parameters (such as recall rates, response time, or learning curves) and relating them in a precise way through mathematical equations (Luce 1995; Coombs, Dawes, and Tversky 1970).

Mathematical models can reveal fundamental regularities and structures in cognitive behavior in much the same way as physicists find regularities in the physical world by relating physical measures through mathematical equations (Luce 1995). If we can find suitable measures and relations of these measures that reveal fundamental (as opposed to superfluous) regularities, mathematical models are superb tools to gain insights into cognitive processes. They can also be useful in serving as a kind of abstract ideal that computational models try to match at the outcome level.

The problem with mathematical models, however, is that it is extremely difficult to find suitable measures. The lack of good behavioral measures leads to, on the one hand, the apparent lack of regularities and, on the other hand, superfluous regularities that can be misleading, which is even worse for cognitive modelers than an apparent lack of regularities. Real-world cognitive behavior is extremely complex and varied, and thus, it is difficult to find good behavioral measures that can provide deep and theoretically interesting insights.

Computational models complement mathematical models by providing detailed internal process descriptions that reveal underlying mechanisms of behavior. Computational modeling opens up the black box, although it is usually done so in a highly hypothetical way. Because in most cases, we do not have sufficient cognitive data that can lead directly and unambiguously to a computational model, numerous assumptions need to be made, and parameters need to be set. Thus, models are often underconstrained from data, even with all the methodologies of protocol analysis (Ericsson and Simon 1993) and other stylized procedures. It is also highly difficult to verify

empirically all the aspects (or even the major aspects) of a computational model because of various practical problems, such as the prohibitive amount of work necessary, the lack of a clear measure that eliminates confounding factors, and the ethical objections to performing certain kinds of experiment. It is also difficult to identify the key aspects in a model (the source of power) that lead to capturing certain cognitive phenomena.

Another shortcoming of computational models is that they often fail to account for individual differences and, thus, serve only as an "average subject," which is nonexistent and ultimately meaningless. They might even average over some important variables and thus fail to account for variance in human behavior (c.f., Walker and Catrambone [1993]). In addition, an individual model usually cannot account for all relevant data within its claimed scope if the scope is sufficiently broad; as a matter of fact, a model usually can only account for a small portion therein. Given the shortcomings of each of the two approaches, it is clear that some combination of the two approaches might be useful or even necessary (Anderson 1993).

Abstract models can serve as a compromise between mathematical and computational models. In abstract models, assumptions and details not essential to the phenomena are omitted, yet computer programs are still necessary to generate model behavior. For example, Pat Langley (1996) presented in the workshop such an abstract model for the sensory-rich flight-control domain. Later, we discuss further the role of abstract models in identifying the source of power (also see Schunn and Reder [1996]).

## Matching Human Cognitive Phenomena

Computational (AI) models, because of their indirect capturing of behavioral data, can be difficult to match against cognitive phenomena. Nevertheless, computational models can be made to correspond with cognitive processes in a variety of ways and thus shed light on cognition accordingly. There are at least the following types

of correspondence between models and cognitive processes, in an increasing order of precision:

First is behavioral outcome modeling; in this case, a computational model produces roughly the same types of behavior as human subjects do, under roughly the same conditions. For example, given a set of scenarios for decision making, a model makes roughly the same kinds of decision as a human decision maker (as, for example, in expert systems or commonsense reasoning models; Collins and Michalski [1989]; Sun [1995]), or given the same piece of text, an AI model extracts roughly the same kind of information that a human reader would (as, for example, in AI natural language-processing systems).

Second is qualitative modeling: A model produces the same qualitative behaviors that characterize human cognitive processes under a variety of circumstances. For example, the performance of human subjects improves or deteriorates when one or more control variables are changed; if a model shows the same changes, given the same manipulations, we can say that the model captures the data qualitatively (see, for example, Medin, Wattenmaker, and Michalski [1987]; Sun [1994]).

Third is quantitative modeling: A model produces exactly the same quantitative behaviors that are exhibited by human subjects, as indicated by certain quantitative performance measures. For example, we can perform point-by-point matching of the learning curve of a model to that of the human subjects, or we can match step-by-step performance of a model with the corresponding performance of humans in a variety of situations (see, for example, Anderson [1993, 1982]; Rosenbloom, Laird, and Newell [1993]).

To date, computational models have had some successes in terms of modeling a wide variety of cognitive phenomena in one of these three senses (especially the first two). These phenomena include concept learning (Medin, Wattenmaker, and Michalski 1987), skill learning (VanLehn 1995; Anderson 1982), and child development (Schmidt and Ling 1996; Shultz,

Mareschal, and Schmidt 1996a). They also include everyday commonsense reasoning, such as in logic-based models (Collins and Michalski 1989), case-based models (Riesbeck and Schank 1989), and connectionist models (Sun 1995) that combine some of the important features of the first two types (Sun 1996). Other phenomena being tackled include word-sense disambiguation in natural language processing, analogical reasoning (Thagard and Holyoak 1990; Holyoak and Thagard 1989), the acquisition of verb past-tense forms (Ling and Marinov 1993; Rumelhart and McClelland 1986), game playing (such as Go or Tic-tac-toe) (Epstein and Gelfand 1996), and expertise (expert versus nonexpert performance) (Chi et al. 1989).

## Source of Power in Computational Models

A computational model consists of many components. When a computational model is judged successful without a proper justification and detailed analysis, it might not be clear which aspects of the model are the main reasons for success. Many claims have been made in the past and are still being made at the present. Therefore, it is essential to consider ways of identifying the source of power in a computational model to account for an apparent success and make meaningful predictions. The source of power in a model can include the (learning) algorithm and its biases; the parameters of the algorithm; task representation (including the scope, the covered aspects, and the way the human data modeled are collected and interpreted); data representation (for example, attributes used); and learning and training regimes, especially their underlying assumptions and biases (for example, Jackson, Constandse, and Cottrell [1996]), including data sampling and selection and data presentation (the order and frequency of presentation).

Obviously, some of these aspects might not or should not be legitimate sources of power. Until now, it has been a matter of case-by-case considerations in the context of specific models and tasks, without any generic

criteria that can be applied universally. Understanding the source of power is useful in zeroing in on central aspects of a successful model, identifying commonalities of models tackling the same problems, and comparing different models. We should also try (at least) to verify and validate important aspects that serve as sources of power in a computational model on the basis of behavioral data and cognitive phenomena.

## Identifying the Source of Power

Assume that a model consists of  $n$  components, denoted as  $C_1, C_2, \dots, C_n$ . For example,  $C_i$  can be learning algorithms, parameters of the learning algorithm, task-representation formats, and training regimes. Each  $C_i$  will take a particular value from a range of possible values. For example, if  $C_1$  represents the learning algorithm, then it can be the decision tree-learning algorithm, feed-forward connectionist model, or nearest-neighbor learning algorithm. To identify the source of power in the model, we need to find out the  $C_i$ 's that are crucial to the success of the modeling. (See the discussion of criteria of successful modeling later in this section.)

If all  $C_i$ 's are independent, identifying the source of power in the model might not be too difficult. Basically, the task becomes much the same as the sensitivity analysis: To see if  $C_k$  is crucial in modeling, we hold the values of all other  $C_i$ 's constant and vary the value of  $C_k$  to see if the model is still as successful as before (again, see the discussion on the criteria of successful modeling). If other values of  $C_k$  do not change much of the model's behavior, then  $C_k$  (in terms of those values tested) is deemed not important to the model's success and, thus, cannot be the source of power in the model. However, if other values of  $C_k$  worsen the model's behavior, the current value of  $C_k$  might be crucial to the model's success and can be a source of power in the modeling. This process should be applied repeatedly to every  $C_i$  in the model.

At the workshop, Shultz, Buckingham, and Oshima-Takane (1996a) pre-

sented generative cascade-correlation connectionist models for two cognitive development tasks: (1) the acquisition of the relation between velocity, time, and distance and (2) the acquisition of personal pronouns. They systematically explored a variety of network topologies and variations in critical learning parameters and found that it is the growing aspect (inclusion of new hidden units) of cascade-correlation networks that produces the desired behaviors. They concluded that the generative feature of cascade-correlation networks is the source of power in modeling child development.

In another workshop paper, however, Miller (1996) pointed out, contradicting common belief, that a certain commonly accepted aspect of models might not be the source of power. It is widely accepted that graded performance of a model is the result of graded representations in the model, such as numeric weight representations in connectionist models. Miller (1996) demonstrated that a symbolic rule-based model symbolic-concept acquisition can produce the appropriate graded performance both in terms of accuracy and response time. The system does not rely on graded representations but, rather, on the process that acquires and accesses symbolic rules. Miller concluded that the source of graded performance might have little to do with explicit graded representations in the models.

As another example, Schmidt and Ling (1996) constructed a symbolic development model for the balance-scale task. Their model, like previous models of the balance-scale task, requires assumptions on representation, learning environment, and learning algorithm. To verify that their choice is crucial in leading toward a successful model, they systematically varied each component (while other components were held constant) to see if the resulting model produces implausible behavior. The components that were shown to play a key role in modeling shed light on the source of power in the model and provide meaningful predictions about the model in terms of representation, learning environment, and learning

algorithm. For example, several redundant attributes (such as an attribute that is true if the scale has equal weight and distance on the both sides) are introduced in representation. Without them, the model did not produce the right behavior (for example, the first two stages are missed). Thus, a meaningful prediction is made that the simple balance problem is particularly salient for the purpose of children's learning, which seems to be supported by the work of Johnson (1987). In terms of learning algorithms, connectionist networks were tested on the same problem, but the model failed to show the first two stages, indicating that the decision tree-learning algorithm is crucial to the model's success.

However, in most computational models,  $C_i$ 's are not independent. That is, if the value of  $C_k$  is changed, the values of other  $C_i$ 's might have to be changed correspondingly, or other components that did not exist before might be introduced. For example, if the current model uses the feed-forward connectionist learning algorithm ( $C_1$ ), it would normally require a distributed representation ( $C_2$ ). If  $C_1$  is changed to a decision tree-learning algorithm, and  $C_2$  is kept constant, it might not be the best choice for the decision tree-learning algorithm because it can apply to the symbolic discrete representation directly without introducing the distributed one. In this case,  $C_2$  might need to be changed (to symbolic representation) accordingly. The same effect occurs when a symbolic learning algorithm (which uses symbolic representation directly) is substituted with a connectionist learning algorithm. In this case,  $C_2$  has to be changed to the distributed representation. In addition, the connectionist learning algorithm has its own parameters (such as network architecture [Wiles and Elman 1996], number of hidden layers, number of nodes in each hidden layer, learning rate, range of initial random weights), which did not exist in the previous model. These parameters are now introduced as a by-product of changing the learning algorithm ( $C_1$ ) to a connectionist learning algorithm. Now the question is, if one cannot "keep everything else

the same," how can one determine if  $C_k$  is a source of power or not?

In principle, if the change of a certain component (such as  $C_1$ ) causes necessary changes in other components (such as  $C_2$ ) or introduces some new components to the model (such as  $C_2'$ ), one should choose new values of such components that optimize the performance of the model. As an example, if a symbolic decision tree-learning algorithm is replaced by a connectionist learning algorithm (a new value of  $C_1$ ), one has to make the necessary change in representation ( $C_2$ ). Some distributed representation, which optimizes the model's performance, should be chosen. Other components, such as sampling method and training-testing sets, should be kept the same. Additional components ( $C_2'$ ) introduced, such as the learning rate and network structure, should be assigned values that optimize the model performance in terms of the criteria of successful modeling (see discussion later).

If, with the change of  $C_1$  and the corresponding changes, the overall model's behavior remains the same, then we deem that  $C_1$  does not play a critical role in modeling and is not a source of power. If, however, such a change results in a deterioration of the model, then the value of  $C_1$  is critical in the successful modeling.

This procedure certainly oversimplifies an important and complex issue in the computational modeling—it only outlines a mechanical principle of identifying the source of power in computational models. For real problems, it might be hard to identify components in a model, determine if components are mutually dependent or not, vary all possible values of some components, assign values of dependent components that optimize the model's outcome, and evaluate if the model is more successful with the new components. In certain simple cases, one can use this procedure to identify if certain components of a model are or are not the source of power.

As an example, Ling (1994) implemented a series of head-to-head comparisons between his symbolic pattern associator (SPA) and feed-forward connectionist models in terms of general-

ization abilities on learning the past tense of English verbs. When using exactly the same training and testing example sets, as well as the same templated binary distributed representation (designed for connectionist models), the testing accuracies of SPA and the connectionist model are close (56.6 percent versus 54.4 percent). These figures indicate that the difference in learning algorithms is not significant using the distributed representation. However, symbolic learning algorithms such as SPA can take symbolic attributes directly, instead of using the distributed representation. Thus, Ling (1994) applied SPA on symbolic representation directly (but it kept other components unchanged). Surprisingly, the testing accuracy is much improved (from 54.4 percent to 76.3 percent). This improvement suggests that the symbolic algorithm (with symbolic representation) is a source of power in a more successful model in terms of the generalization ability, although connectionist models provided the initial insight into the modeling of this task.

This process is effectively an exploration of design space for cognitive models, as advocated by Sloman (1996) in his talk at the workshop. Although we are exploring the behavioral space in the sense of identifying the range and variations in human behavior, we also need to explore the design space (that is, the possibilities for constructing models) that maps onto the behavioral space, so that we can gain a better understanding of the possibilities and limitations of our modeling methodologies and open up new avenues for better capturing cognitive processes in computational terms. This is especially important for developing reasonable, realistic, and convincing cognitive architectures that are highly complex and in which many design decisions have to be made without the benefit of a clear understanding of their full implications in computational or behavioral terms (Sun 1994). Systematic, or at least methodic, exploration of design space is necessary in identifying the source of power in cognitive models.

Finally, abstract models are another approach to facilitate the identifica-

tion of the source of power in modeling. As we discussed earlier, abstract models omit details that are not essential to the phenomena, but computer simulation is still performed for modeling behaviors. This approach would reduce the number of components in computational models, and therefore, it becomes easier to analyze and recognize key components in the models. At the workshop, Langley (1996) presented an abstract model for the sensory-rich flight-control domain for which the traditional computational model has been found too complex to model the available data. In the abstract model, few parameters are incorporated, and the model's central assumptions are analyzed and tested in detail (Langley 1996).

## Criteria of Successful Modeling

Identifying the source of power in models relies on the criterion used for measuring successful modeling, which itself is an important issue in cognitive modeling. It is evident by now that any sufficiently powerful computational (AI) model can be made to capture data in any narrow domain (especially those sanitized, isolated data with a fixed representation). For example, in his workshop presentation, Mareschal (1996) reviewed three different models of object permanence. All three models capture the data sufficiently well. However, five aspects of the information processing in the models distinguish good development models from others: (1) whether the model has a transition mechanism, (2) whether there is gradual knowledge transition, (3) whether the model is directly coupled to the input, (4) whether the model extracts information from the entire scene, and (5) whether the model reflects individual differences. In another workshop paper, Thagard (1996) presented a review of four competing models for analogical reasoning. He outlined seven criteria—(1) genuineness, (2) breadth of application, (3) scaling, (4) qualitative fit, (5) quantitative fit, (6) comparison, and (7) compatibility—for evaluating computational models of cognition in general. The four competing models for analog-

ical reasoning were analyzed and compared using these criteria (also see Veale et al. [1996]).

Therefore, we need to look into deeper issues beyond simple goodness of fit. In general, such issues can include (1) explanatory power, degrees to which data are accounted for by a model. It is especially important to use real-world situations, not just sanitized, isolated data, because real-world situations can be vastly different from laboratory situations; (2) generality, to be discussed in more detail later; (3) economy, the succinctness with which a model explains behaviors; (4) consistency with other knowledge, including compatibility of the transient process of models with human learning and development data, compatibility with models (or principles) in other related domains, and compatibility with evidence from other disciplines (especially psychological and neurobiological evidence); and (5) computational power and complexity and the correspondences of the model with the human data in this regard. There is always a many-to-many mapping between computational models and the cognitive phenomena to be modeled, so the development and application of fundamental, abstract generalized criteria in analyzing and comparing models are essential.

Among these factors, generality is an especially important consideration. To measure generality, we propose to look at an abstract measure:  $r = \text{scope} / \text{degree of freedom}$ . Generally speaking, the greater the measure is, the better; that is, we don't want a model with too narrow a scope but with too many parameters. As an example, Schmidt and Ling's (1996) symbolic model of the balance-scale task can explain the task with any number of weights, but the connectionist models work only for the five-weight version. At the same time, the symbolic model has only one free parameter, but the connectionist models have more. Therefore, the symbolic model has a broader generalization ability than connectionist models on the same task.

However, the application of this abstract measure of generality in other more complex domains might not be easy, especially when comparing

generic models such as cognitive architectures (Sun 1994). First, how do we define scope? For one thing, the scope of an architecture widens gradually as more and more applications are developed. It is not static. Second, how do we define degree of freedom? Can it be the number of parameters? In production systems, is it each production (or each numeric parameter therein)? In neural networks, is it each weight? If we take a close look at the two types of system—symbolic and connectionist—we see that the two are actually not comparable; in general, it is not just the number of parameters that matters but also the roles that they play in the working of a model in both learning and performance. For example, weights are mostly learned, but productions are mostly hand coded (at least to begin with), thus entailing different amounts of predetermination. The earlier suggestion, then, is too crude a measure to capture the difference between these two types of system and to compare their degrees of freedom. We suggest that we can look into parameters of a model along the following dimensions: (1) information content of a parameter; (2) ease in obtaining (estimating) a parameter; (3) preset versus learned; (4) amount of hand coding in either case (for example, much more hand coding in a production [such as that in SOAR] [Rosenbloom, Laird, and Newell 1993]) than in weights of neural networks); (5) emergent computational process versus that directly captured, step by step, by parameters (for example, by productions in a production system); and (6) the contribution of each parameter (for example, each weight by itself in a backpropagation network might reveal little about the function of the network.)

For example, in a backpropagation neural network, when different random initial weights and different training regimes (data presentations) result in similar model behaviors, free parameters are few: the number of hidden units (assuming that one layer of hidden units is used) and one or two parameters for learning (such as the learning rate and the momentum parameter). However, in a production system (even one with the learning

capability), many more parameters are used. Such parameters include each initial production (which requires extensive hand coding and contains a great deal of information) and parameters in whatever learning algorithms are used (for example, chunking or proceduralization in Anderson [1993] and Rosenbloom, Laird, and Newell [1993], which are the most commonly used forms of learning in symbolic cognitive modeling), which might require a priori estimation. Most likely, such a system contains a large number of initial productions, on the order of hundreds, which means at least hundreds of free parameters that can be tuned. In addition, the computational process is strictly directed by the chaining of individual productions, and the steps of the process can be isolated and identified with individual productions. Thus, it is easier, in the sense of being less constrained, to construct a cognitive model that captures a certain amount of data using production systems than using a connectionist model. Some other types of symbolic model, such as decision trees, share with connectionist models this property of being highly constrained. Such highly constrained models, with fewer parameters, can naturally lead to more succinct explanations of cognitive phenomena and, therefore, better insights into the fundamentals of cognitive processes.

Another aspect to consider is how to incorporate all relevant cognitive dimensions. To construct reasonable computational models of cognition, we would like to capture some of the known cognitive distinctions: learning versus performance, implicit versus explicit learning, implicit versus explicit performance (for example, memory), controlled versus automatized performance, and short-term versus long-term memory. We should investigate how each dimension can be (potentially) captured or ignored in each individual computational model. These dimensions also provide useful clues and constraints in constructing computational models of cognition, especially those comprehensive models known as cognitive architectures.

Finally, intertheoretical reduction needs to be considered. At the top of

the cognitive modeling hierarchy, we have abstract models of behavior such as the kind of behavioral mathematical model that we discussed earlier, and beneath them in this hierarchy, more detailed computational models can provide process details. In fact, a hierarchy of more and more detailed computational models can coexist that can progressively move toward capturing more and more microscopic levels of cognitive behavior. This is much like intertheoretical reduction in physical science: Thermodynamics that describes mass properties at a macroscopic level can be reduced to Newtonian physics at the molecular level, which, in turn, can be reduced to atomic and subatomic physics, which can then be related to even lower levels describable by quantum physics, and so on. This point was emphasized in an excellent treatment of mathematical psychology by Luce (1995).

## Concluding Remarks

To recapitulate some main points, there is almost always a many-to-many mapping between models and cognitive processes to be modeled. Thus, we need generalized criteria for evaluating and comparing models, especially those of different paradigms (e.g., connectionist versus symbolic) and of different levels of abstraction (for example, mathematical models versus computational models). We need to analyze systematically a successful model to identify its source of power and, thus, to make meaningful predictions from the model. It is also important to pay serious attention to the need of capturing real-world situations and performance, not just isolated sanitized data. This is because real-world data can be vastly different from laboratory situations, as has been demonstrated in many different domains (Sun 1994), and thus, can broaden the scope and enhance the ecological realism of cognitive modeling.

## Acknowledgments

Ron Sun's work is supported in part by Office of Naval Research grant N00014-95-1-0440. Charles Ling's work is supported in part by Natural Sciences and Engineering Research

Council of Canada research grant OGP0046392. We thank members of the program committee (Pat Langley, Mike Pazzani, Paul Thagard, Kurt VanLehn, and Tom Shultz) for their roles in organizing the workshop. We also thank all the paper presenters and workshop participants, especially the invited speakers: Garrison W. Cottrell, Denis Mareschal, Aaron Sloman, Tom Shultz, and Paul Thagard.

## References

- Anderson, J. 1993. *Rules of the Mind*. Hillsdale, N.J.: Lawrence Erlbaum.
- Anderson, J. 1982. Acquisition of Cognitive Skill. *Psychological Review* 89:369-406.
- Chi, M.; Bassok, M.; Lewis, M.; Reimann, P.; and Glaser, P. 1989. Self-Explanation: How Students Study and Use Examples in Learning to Solve Problems. *Cognitive Science* 13:145-182.
- Collins, A., and Michalski, R. 1989. The Logic of Plausible Reasoning: A Core Theory. *Cognitive Science* 13(1): 1-49.
- Coombs, C.; Dawes, R.; and Tversky, A. 1970. *Mathematical Psychology*. Englewood Cliffs, N.J.: Prentice Hall.
- Epstein, S. L., and Gelfand, J. 1996. The Creation of New Problem-Solving Agents from Experience with Visual Features. In Papers from AAAI-96 Workshop on Computational Cognitive Modeling: Source of the Power, eds. C. X. Ling and R. Sun, Technical Report, TR-CS-96-0023, University of Alabama.
- Ericsson, K. A., and Simon, H. 1993. *Protocol Analysis*. Cambridge, Mass.: MIT Press.
- Holyoak, K., and Thagard, P. 1989. A Computational Model of Analogical Problem Solving. In *Similarity and Analogical Reasoning*, eds. S. Vosniadou and A. Ortony, 242-266. New York: Cambridge University Press.
- Jackson, D.; Constandse, R. M.; and Cottrell, G. W. 1996. Selective Attention in the Acquisition of the Past Tense. In Papers from AAAI-96 Workshop on Computational Cognitive Modeling: Source of the Power, eds. C. X. Ling and R. Sun, Technical Report, TR-CS-96-0023, University of Alabama.
- Johnson, M. 1987. *The Body in the Mind: The Bodily Basis of Meaning, Imagination, and Reason*. Chicago: University of Chicago Press.
- Langley, P. 1996. An Abstract Computational Model of Learning Selective Sensing Skills. In Papers from AAAI-96 Workshop on Computational Cognitive Modeling: Source of the Power, eds. C. X. Ling and R. Sun, Technical Report, TR-CS-96-0023, University of Alabama.
- Ling, C. X. 1994. Learning the Past Tense of English Verbs: The Symbolic Pattern Associator versus Connectionist Models. *Journal of Artificial Intelligence Research* 1:209-229.
- Ling, C. X., and Marinov, M. 1993. Answering the Connectionist Challenge: A Symbolic Model of Learning the Past Tense of English Verbs. *Cognition* 49(3): 235-290.
- Luce, L. D. 1995. Four Tensions Concerning Mathematical Modeling in Psychology. *Annual Review of Psychology* 46:1-26.
- Mareschal, D. 1996. Models of Object Permanence: How and Why They Work. In Papers from AAAI-96 Workshop on Computational Cognitive Modeling: Source of the Power, eds. C. X. Ling and R. Sun, Technical Report, TR-CS-96-0023, University of Alabama.
- Medin, D.; Wattenmaker, W.; and Michalski, R. 1987. Constraints and Preferences in Inductive Learning: An Experimental Study. *Cognitive Science* 11:299-339.
- Miller, C. S. 1996. The Source of Graded Performance in a Symbolic Rule-Based Model. In Papers from AAAI-96 Workshop on Computational Cognitive Modeling: Source of the Power, eds. C. X. Ling and R. Sun, Technical Report, TR-CS-96-0023, University of Alabama.
- Riesbeck, C., and Schank, R. 1989. *Inside Case-Based Reasoning*. Hillsdale, N.J.: Lawrence Erlbaum.
- Rosenbloom, P.; Laird, J.; and Newell, A. 1993. *The SOAR Papers: Research on Integrated Intelligence*. Cambridge, Mass.: MIT Press.
- Rumelhart, D., and McClelland, J. 1986. On Learning the Past Tenses of English Verbs. In *Parallel Distributed Processing, Volume 2*, eds. D. Rumelhart, J. McClelland, and the PDP Research Group, 216-271. Cambridge, Mass.: MIT Press.
- Schmidt, W. C., and Ling, C. X. 1996. A Decision Tree Model of Balance-Scale Development. *Machine Learning* 24: 203-230.
- Schunn, C. D., and Reder, L. M. 1996. Modeling Changes in Strategy Selections over Time. In Papers from AAAI-96 Workshop on Computational Cognitive Modeling: Source of the Power, eds. C. X. Ling and R. Sun, Technical Report, TR-CS-96-0023, University of Alabama.
- Shultz, T.; Buckingham, D.; and Oshima-Takane, Y. 1996a. Generative Connectionist Models of Cognitive Development: Why They Work. In Papers from AAAI-96 Workshop on Computational Cognitive Modeling: Source of the Power, eds. C. X. Ling and R. Sun, Technical Report, TR-CS-96-0023, University of Alabama.
- Shultz, T. R.; Mareschal, D.; and Schmidt, W. C. 1996b. Modeling Cognitive Develop-



# Simulating Organizations

## Computational Models of Institutions and Groups

*Edited by Michael J. Prietula, Kathleen M. Carley, and Les Gasser*

**T**he globalization of the economy, increasing number of transnational organizations, and rapid changes in robotics, information, and telecommunication technologies are just a few of the factors significantly altering organizational time scales, forms, complexity, and environments. Time scales have shrunk, new organizational forms are emerging, and organizational environments are expanding and mutating at unprecedented rates. Computational modeling affords opportunities to both understand and respond to these complex changes.

Paralleling developments in the physical sciences, computational modeling is emerging in the social and organizational sciences. Organizational researchers are using computational models to gain insights into organizational phenomena and to explore dynamic processes and configurations that are difficult or impossible to investigate with other methods. Many interesting insights have already resulted from this research, such as how group cooperation arises or dissipates in social dilemma settings, and how honesty and benevolence affect behavior in a group task. On the practical side, computational modeling is increasingly effective for organizational design, analysis, and reengineering.

Although much work remains to be done, eventually both theorists and practitioners will routinely state theories, design organizations, and derive their implications using widely shared computational tools. This volume brings together work from many of the leading researchers in the field.

6 x 9, 350 pp., \$45.00, ISBN 0-262-66108-X

*Prices higher outside the U.S. and subject to change without notice.*

**To order, call 800-356-0343 (US and Canada) or (617) 625-8569.  
Distributed by The MIT Press, 55 Hayward, Cambridge, MA 02142**

## NEW TECHNICAL REPORTS FROM AAAI PRESS

### Qualitative Reasoning:

#### The Twelfth International Workshop

*Edited by Feng Zhao and Kenneth Yip*

178 pp., \$25.00

Technical Report WS-98-01

ISBN 0-1-57735-054-5

#### Satisficing Models: Papers from the 1998 AAAI Spring Symposium

*Sandip Sen, Chair*

Technical Report SS-98-05

94 pp., \$25.00

ISBN 1-57735-050-2

#### Multimodal Reasoning: Papers from the 1998 AAAI Spring Symposium

*Eugene Freuder, Chair*

Technical Report SS-98-04

174 pp., \$25.00

ISBN 1-57735-049-9

#### Interactive and Mixed-Initiative Deci- sion-Theoretic Systems: Papers from the 1998 AAAI Spring Symposium

*Peter Haddawy and*

*Steve Hanks, Cochairs*

Technical Report SS-98-03

151 pp., \$25.00

ISBN 1-57735-048-0

#### Intelligent Environments: Papers from the 1998 AAAI Spring Symposium

*Michael Coen, Chair*

Technical Report SS-98-02

181 pp., \$25.00

ISBN 1-57735-047-2

#### Applying Machine Learning to Discourse Processing: Papers from the 1998 AAAI Spring Symposium

*Jennifer Chu-Carroll and*

*Nancy Green, Cochairs*

Technical Report SS-98-01

128 pp., \$25.00

ISBN 1-57735-046-4

To order, call 650-328-3123 or  
send e-mail to [orders@aaai.org](mailto:orders@aaai.org).

Further information is available at  
[www.aaai.org/TechReports/](http://www.aaai.org/TechReports/)  
AAAI Press • 445 Burgess Drive  
• Menlo Park, CA 94025

ment on Balance-Scale Phenomena.  
*Machine Learning* 16:57-86.

Slovan, A. 1996. What Sort of Architecture Is Required for a Humanlike Agent? In Papers from AAAI-96 Workshop on Computational Cognitive Modeling: Source of the Power, eds. C. X. Ling and R. Sun, Technical Report, TR-CS-96-0023, University of Alabama.

Sun, R. 1996. Commonsense Reasoning with Rules, Cases, and Connectionist Models: A Paradigmatic Comparison. *Fuzzy Sets and Systems* 82:187-200.

Sun, R. 1995. Robust Reasoning: Integrating Rule-Based and Similarity-Based Reasoning. *Artificial Intelligence* 75:241-296.

Sun, R. 1994. *Integrating Rules and Connectionism for Robust Commonsense*. New York: Wiley.

Thagard, P. 1996. Evaluating Computational Models of Cognition: Notes from the Analogy Wars. In Papers from AAAI-96 Workshop on Computational Cognitive Modeling: Source of the Power, eds. C. X. Ling and R. Sun, Technical Report, TR-CS-96-0023, University of Alabama.

Thagard, P., and Holyoak, K. 1990. Analog Retrieval by Constraint Satisfaction. *Artificial Intelligence* 46:259-310.

VanLehn, K. 1995. Cognitive Skill Acquisition. In *Annual Review of Psychology, Volume 47*, eds. J. Spence, J. Darly, and D. Foss. Palo Alto, Calif.: Annual Reviews.

Veale, T.; Smyth, B.; O'Donoghue, D.; and Keane, M. 1996. Representational Myopia in Cognitive Mapping. In Papers from AAAI-96 Workshop on Computational Cognitive Modeling: Source of the Power, eds. C. X. Ling and R. Sun, Technical Report, TR-CS-96-0023, University of Alabama.

Walker, N., and Catrambone, R. 1993. Aggregation Bias and the Use of Regression in Evaluating Models of Human Performance. *Human Factors* 35:397-411.

Wiles, J., and Elman, J. 1996. States and Stacks: Doing Computation with a Recurrent Neural Network. In Papers from AAAI-96 Workshop on Computational Cognitive Modeling: Source of the Power, eds. C. X. Ling and R. Sun, Technical Report, TR-CS-96-0023, University of Alabama.



**Ron Sun** is currently an associate professor of computer science at the University of Alabama. He received his Ph.D. in 1991 from Brandeis University. Sun's research interest centers on the studies of intelligence and cognition, especially in the areas of

reasoning, learning, and connectionist models. He is the author of 50+ papers and has written, edited, or contributed to 8 books. For his paper on integrating rule-based reasoning and connectionist models, he received the 1991 David Marr Award from the Cognitive Science Society. He (co)chaired several workshops on hybrid models and cognitive modeling. He was the guest editor of the special issue of *Connection Science* focusing on architectures for integrating neural and symbolic processes and the special issue of *IEEE Transactions on Neural Networks* focusing on hybrid intelligent models. His home page is [cs.us.edu/~rsun](http://cs.us.edu/~rsun).



**Charles X. Ling** obtained his B.Sc. in computer science at Shanghai JiaoTong University, China, in 1985, and a Ph.D. from the Department of Computer and Information Science, University of Pennsylvania,

in 1989. Since then, he has been a faculty member in the Department of Computer Science, The University of Western Ontario, where he is currently an associate professor. He has done extensive research in computational modeling of benchmark cognitive learning tasks. He has also worked in other areas of machine learning, including inductive logic programming, learning to control dynamic systems, knowledge base refinement, nearest-neighbor algorithms, and artificial neural networks. His home page is [www.csd.uwo.ca/faculty/ling](http://www.csd.uwo.ca/faculty/ling).