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

Ron Sun
Dept. of Computer Science
University of Alabama
Tuscaloosa, AL 35487, USA
rsun@cs.ua.edu
Charles X. Ling
Department of Computer Science
The University of Western Ontario
London, Ontario N6A 5B7, Canada
ling@csd.uwo.ca

February 11, 1997

Abstract

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

1 Introduction

The workshop on "Computational Modeling: the Source of Power", which we co-chaired and organized, was held at AAAI'96 on August 5, 1996 at 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 source of power in various models are especially prominent issues in the discussions. To summarize and continue the theme of this workshop, we would like to offer the highlights of a few important issues of cognitive modeling, including why computational models work, measures of success, and systematic ways of identifying the source of power in the models.

Computational cognitive modeling is an important aspect in cognitive science since it plays a central role in computational understanding of the mind. A common methodology of cognitive science is to express the theory about human cognition in a computer program and then compare the program's behavior with human cognitive behavior. However, many fundamental questions remain unanswered: To what extent shall 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 this workshop and the present report will stimulate more discussions, and that a better understanding will be reached in the future.

Cognitive modeling has traditionally been underrepresented in AAAI conferences and journals. However, its importance should not be underestimated. It would provide key insights into models of intelligence, and thus provide great impetus for the further development of AI. In 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.

2 Types of Cognitive Models

Cognitive modeling can be of two major types. One is by (to some extent) detailed, computational models (or AI models), and the other is by (behavioral) mathematical models. The former seeks to capture internal computational (mechanistic) processes that generate overt cognitive behavior (e.g., Anderson, 1982, 1993; 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, & Tversky, 1970).

Mathematical models can reveal fundamental regularities and structures in cognitive behavior, in very 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 those 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 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 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 the apparent lack of regularities. The real-world cognitive behavior is extremely complex and varied, and it is thus difficult to find good behavioral measures that can provide deep, theoretically interesting insight.

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 in a highly hypothetical way — often numerous assumptions need to be made and parameters to be set. This is because in most cases we do not have sufficient cognitive data that can lead directly and *unambiguously* to a computational model. Thus models are often underconstrained from data, even with all the methodologies of protocol analysis developed (Ericsson & Simon, 1993) and other processes. 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 experiments. It is also difficult to identify the key aspects in a model (the source of power) that lead to capturing a cognitive phenomenon. Another shortcoming of computational models is that they often failed to account for individual difference and thus serve only as an "average subject", which is non-existent and ultimately meaningless. Or they may even average over some important variables and thus fail to account for important variance in human behavior (cf. Walker & Catrambone, 1993). In addition, an individual model usually cannot account for all the relevant data within its claimed scopes, if the scope is sufficiently broad; as a matter of fact, a model usually can only account for a small portion therein. So, given these shortcomings of each of the two approaches, it is clear that some combination of the two approaches may be useful or even necessary (Anderson, 1993).

Abstract models may serve as a compromise between mathematical and computational models, as proposed by Langley (1996). In abstract models, assumptions and details that are not essential to the phenomena are omitted, yet computer programs are still necessary to generate model behaviors. E.g., Langley (1996) presented in the workshop an abstract model for the sensory-rich flight-control domain. We will discuss further the role of abstract models in identifying the source of power in models later (also see Schunn & Reder, 1996).

3 Matching Human Cognitive Phenomena

Computational (AI) models, because of their indirect capturing of behavioral data, can be difficult to match against cognitive phenomena. 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:

- Behavioral outcome modeling: in this case, a computational model produces roughly the same types of behaviors 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 decisions as a human decision maker (as, e.g., in expert systems, or commonsense reasoning models; Collins & 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 extract (as, e.g., in AI natural language processing systems).
- 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/deteriorates when one or more control variables are changed; if a model shows the same changes, given the same manipulations, we may say that the model captures the data qualitatively (see, e.g., Medin, Wattenmaker, & Michalski, 1987; Sun, 1994).
- Quantitative modeling: a model produces exactly the same quantitative behaviors as 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, e.g., Anderson, 1982, 1993; Rosenbloom, Laird, & Newell, 1993).

Computational models so far have some successes in terms of modeling a wide variety of cognitive phenomena in one of the above three senses (especially in the first two senses). These phenomena include learning, such as learning concept descriptions (Medin et al., 1987), skill learning (Anderson, 1982; VanLehn, 1995), and decision-tree, or connectionist learning models of development (Shultz, Mareschal, & Schmidt, 1996a; Schmidt & Ling, 1996). They also include everyday commonsense reasoning, such as in logic-based models (Collins & Michalski, 1989), case-based models (Riesbeck &

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 (Holyoak & Thagard, 1989; Thagard & Holyoak, 1990), acquiring verb past-tense forms (Rumelhart & McClelland, 1986; Ling & Marinov, 1993), game playing (such as go or tic-tac-toe; Epstein & Gelfand, 1996), expertise (expert vs. non-expert performance; Chi, Bassok, Lewis, Reimann, & Glaser, 1989), and so on.

4 Source of Power in Computational Models

It is extremely difficult to verify empirically all the aspects of a computational model; this is particularly true for qualitative and quantitative models. Therefore, when a computational model is judged to be successful, without proper justification and detailed analysis, it may not be clear at all which aspects of the model are the main reasons for success. Many such claims have been made in the past, however, and are probably still being made at the present.

Therefore, it is essential to consider ways of identifying the source of power in a computational model, in order to account for an apparent success and to make meaningful predictions of the model. The source of power in a model can include

1. the (learning) algorithm and its bias,
2. the parameters of the algorithm,
3. task representation (including the scope, the covered aspects, and the way the human data modeled is collected and interpreted),
4. data representation (e.g., attributes used), and
5. learning/training regimes, especially their underlying assumptions and biases (e.g., Jackson, Constandse, & Cottrell, 1996), including data sampling and selection and data presentation (the order and frequency of presentation).

Obviously, some of these aspects may not or should not be legitimate sources of power. Up 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 universally applied. Nevertheless, the importance of understanding the source of power should not be underestimated. Understanding the source of power can be 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 important aspects that serve as sources of power in a computational model on the basis of behavioral data and cognitive phenomena.

5 Identifying the Source of Power

All cognitive models consist of several components, and each can play a specific role in modeling. Assume that a model consists of n components, denoted as $C_1, C_2, ..., C_n$. Specifically, C_i can be the learning algorithm, parameters of the learning algorithm, task representation format, training regimes, and the like. Each C_i will take on a particular "value" in a particular modeling task. For example, if C_1 represents the learning algorithm, then it can be the decision-tree learning algorithm, feedforward connectionist model, or nearest neighbouring algorithm in a particular computational model. To identify the source of power in the model, we need to find which C_i's are crucial to the success of the modeling (see the discussion of criteria of successful modeling next section).

If all C_i's are independent, identifying the source of power in the model may not be too difficult. Basically the task becomes very much the same as simple 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 next section on criteria of successful modeling). If other values of C_k do not change much of the model's behavior, then C_k is deemed not important in the model's success, and thus cannot be the source of power in the model. On the other hand, if other values of C_k worsen the model's behavior, the current value of C_k may 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 (1996b) presented generative cascade-correlation connectionist models for two cognitive development tasks: acquisition of the relation between velocity, time, and distance, and 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 may not be the source of power. It is widely accepted that graded performance of a model is due to graded representations in the model, such as numerical weight representations in connectionist models. Miller (1996) demonstrated that a symbolic rule-based model SCA (symbolic concept acquisition) can produce 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 may 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. The model, like previous models of the balance scale task, requires assumptions about representation, learning environment, and learning algorithm. To verify that their choice is crucial in leading towards a successful model, they systematically very each component (while other components are held constant) to see if the resulting model produces implausible behavior. The components that are 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 which is true if the balance scale has equal weight and distance on both sides) are introduced in representation. Without them, the model does not produce the right behavior (e.g., the first two stages are missed). This makes a meaningful prediction 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 the first two stages. This indicates that the decision tree learning algorithm is crucial in the model's success.

However, things are often not that simple. 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 may have to be changed correspondingly, or other components which did not exist before may be introduced. For example, if the current model uses the feedforward connectionist learning algorithm (C_1), it would normally require a distributed representation (C_2). If you change C_1 to be a decision-tree learning algorithm, and keep C_2 constant, it may not be the best choice for the decision-tree learning algorithm, since it can apply to the symbolic discrete representation directly without introducing the distributed one. In this case, C_2 may need to be changed (to symbolic representation) correspondingly. The same effect occurs when you substitute a symbolic learning algorithm (which uses symbolic representation directly) 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 (e.g., Wiles & Elman, 1996), number of hidden layers, number of nodes in each hidden layer, learning rate, range of initial random weights, etc.), which did not exist in the previous model. Those parameters are now introduced as a by-product of changing the learning algorithm (C_1) to be 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, we need to do the following: 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 the new values of such components to 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 criteria of successful modeling (see the next section).

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 causes deterioration of the model, then the value of C_1 is critical in the successful modeling.

As an example, Ling (1994) implemented a series of head-to-head comparisons between his Symbolic Pattern Associator (SPA) and feedforward connectionist models in terms of generalization 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 connectionist model are very close (56.6% vs 54.4%). This indicates 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 (while keeping other components unchanged). Surprisingly, the testing accuracy is much improved (from 54.5% to 76.3%). This 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 is effectively an exploration of design space for cognitive models, as advocated by Sloman (1996) in his talk at the workshop. While 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 may gain a better understanding of possibilities and limitations of our modeling methodologies and open up new avenues for better capturing of 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 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 identification 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).

6 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 within 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: whether or not the model has a transition mechanism, whether or not there is a gradual knowledge transition, whether or not the model is directly coupled to the inputs, whether or not the model extracts information from the entire scene, and whether or not the model reflects individual differences. In another workshop paper, Thagard (1996) presented an excellent review of four competing models for analogical reasoning. He outlined seven criteria (genuineness, breadth of application, scaling, qualitative fit, quantitative fit, comparison, and compatibility) for evaluating computational models of cognition in general. The four competing models for analogical reasoning were analyzed and compared using those criteria (also see Veale, Smyth, O'Donoghue, & Keane, 1996).

Therefore, we need to look into deeper issues beyond simple goodness of fit. In general, such issues may include

1. 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 could be vastly different from laboratory situations,

2. generality of the model,

3. compatibility of the transient process of models with human learning and development data,
4. compatibility with models (or principles) in other related domains,
5. compatibility with evidence from other disciplines (especially psychological and neurobiological evidence)
6. computational power, and the correspondence of the model with the human data in this regard, and
7. computational complexity and the correspondence 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 criteria in analyzing and comparing models are essential.

Among these factors, generality is an especially important consideration. In terms of measuring generality, we propose to look at an abstract measure:

$$r = \frac{Scope}{DegreeOfFreedom}$$

Generally speaking, the greater the measure, 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 (1996)'s symbolic model of the balance scale task can explain the task with any number of weights, while connectionist models work only for the 5-weight version. At the same time, the symbolic model has only one free parameter, while 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 may not be easy, especially when comparing generic models such as cognitive architectures (Sun, 1994). First of all, how do we define scope? For one thing, the scope of an algorithm widens gradually, as more and more applications are developed. It is not static. Second, how do we define a degree of freedom? Can it be the number of parameters? In production systems, is it each production (or each numerical parameter therein)? In neural networks, is it each weight? If we take a close look at the two types of systems, symbolic and connectionist models, 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 pre-determination. So the earlier suggestion is too crude a measure to capture the difference between these two types of systems and to compare their degrees of freedom. We suggest that we can look into parameters of a model along the following dimensions.

- information content of a parameter
- ease in obtaining (estimating) a parameter
- preset vs. learned
- amount of hand-coding in either case; e.g., much more hand-coding is involved in a production (such as that in SOAR; Rosenbloom et al., 1993) than in weights of neural networks
- the computational process is emergent vs. directly captured step-by-step by parameters (e.g., by productions in a production system)
- the contribution of each parameter (e.g. each weight by itself in a backpropagation network may reveal very 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 1 or 2 parameters for learning (such as the learning rate and the momentum parameter). On the other hand, 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 used (e.g., chunking or proceduralization in Anderson, 1993; Rosenbloom et al., 1993, which are the most commonly used forms of learning in symbolic cognitive modeling), which may require a priori estimation. Most likely, such a system contains a large number of initial productions, in the order of hundreds, which means 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 models, 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 vs. performance
- implicit vs. explicit learning
- implicit vs. explicit performance (e.g., memory)
- controlled vs. automatized performance
- short-term vs. 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 computational architectures.

Finally, inter-theoretical reduction needs to be considered. At the top of the cognitive modeling hierarchy, we shall 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 co-exist that can progressively move towards capturing more and more microscopic levels of cognitive behavior. This is very much like inter-theoretical 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).

7 Concluding Remarks

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

Acknowledgements

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 an NSERC (Natural Sciences and Engineering Research Council of Canada) Research Grant (Grant Number 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.

Reference

Anderson, J. (1982). Acquisition of cognitive skill. *Psychological Review*, *89*, 369 – 406.

Anderson, J. (1993). *Rules of the Mind.* Lawrence Erlbaum Associates. Hillsdale, NJ.

Chi, M., Bassok, M., Lewis, M., Reimann, P., & Glaser, P. (1989). Self-explanation: How students study and use examples in learning to solve problems. *Cognitive Science*, *13*, 145 – 182.

Collins, A., & Michalski, R. (1989). The logic of plausible reasoning: A core theory. *Cognitive Science*, *13*(1), 1 – 49.

Coombs, C., Dawes, R., & Tversky, A. (1970). *Mathematical Psychology.* Prentice Hall, Englewood-Cliffs, NJ.

Epstein, S. L., & Gelfand, J. (1996). The creation of new problem solving agents from experience with visual features. In Ling, C. X., & Sun, R. (Eds.), *Working Notes of AAAI-96 Workshop on Computational Cognitive Modeling: Source of the Power.*

Ericsson, K. A., & Simon, H. (1993). *Protocol Analysis.* MIT Press, Cambridge, MA.

Holyoak, K., & Thagard, P. (1989). A computational model of analogical problem solving. In Vosniadou, S., & Ortony, A. (Eds.), *Similarity and Analogical Reasoning.* Cambridge University Press, New York, NY.

Jackson, D., Constandse, R. M., & Cottrell, G. W. (1996). Selective attention in the acquisition of the past tense. In Ling, C. X., & Sun, R. (Eds.), *Working Notes of AAAI-96 Workshop on Computational Cognitive Modeling: Source of the Power.*

Johnson, M. (1987). *The Body in the Mind: The Bodily Basis of Meaning, Imagination, and Reason.* Chicago: The University of Chicago Press.

Langley, P. (1996). An abstract computational model of learning selective sensing skills. In Ling, C. X., & Sun, R. (Eds.), *Working Notes of AAAI-96 Workshop on Computational Cognitive Modeling: Source of the Power.*

Ling, C. X. (1994). Learning the past tense of English verbs: the Symbolic Pattern Associator vs. connectionist models. *Journal of Artificial Intelligence Research*, *1*, 209 – 229.

Ling, C. X., & 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 Ling, C. X., & Sun, R. (Eds.), *Working Notes of AAAI-96 Workshop on Computational Cognitive Modeling: Source of the Power.*

Medin, D., Wattenmaker, W., & 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 Ling, C. X., & Sun, R. (Eds.), *Working Notes of AAAI-96 Workshop on Computational Cognitive Modeling: Source of the Power.*

Riesbeck, C., & Schank, R. (1989). *Inside Case-based Reasoning.* Lawrence Erlbaum Associate, Hillsdale, NJ.

Rosenbloom, P., Laird, J., & Newell, A. (1993). *The SOAR Papers: Research on Integrated Intelligence.* MIT Press, Cambridge, MA.

Rumelhart, D., & McClelland, J. (1986). On learning the past tenses of English verbs. In Rumelhart, D., McClelland, J., & the PDP Research Group (Eds.), *Parallel Distributed Processing, Vol. 2*, pp. 216 – 271. Cambridge, MA: MIT Press.

Schmidt, W. C., & Ling, C. X. (1996). A decision-tree model of balance scale development. *Machine Learning*, *24*, 203 – 230.

Schunn, C. D., & Reder, L. M. (1996). Modeling changes in strategy selections over time. In Ling, C. X., & Sun, R. (Eds.), *Working Notes of AAAI-96 Workshop on Computational Cognitive Modeling: Source of the Power.*

Shultz, T. R., Mareschal, D., & Schmidt, W. C. (1996a). Modeling cognitive development on balance scale phenomena. *Machine Learning*, *16*, 57 – 86.

Shultz, T., Buckingham, D., & Oshima-Takane, Y. (1996b). Generative connectionist models of cognitive development: Why they work. In Ling, C. X., & Sun, R. (Eds.), *Working Notes of AAAI-96 Workshop on Computational Cognitive Modeling: Source of the Power.*

Sloman, A. (1996). What sort of architecture is required for a human-like agent?. In Ling, C. X., & Sun, R. (Eds.), *Working Notes of AAAI-96 Workshop on Computational Cognitive Modeling: Source of the Power.*

Sun, R. (1994). *Integrating Rules and Connectionism for Robust Commonsense.* John Wiley and Sons, New York, NY.

Sun, R. (1995). Robust reasoning: integrating rule-based and similarity-based reasoning. *Artificial Intelligence*, *75*, 241 – 296.

Sun, R. (1996). Commonsense reasoning with rules, cases, and connectionist models: a paradigmatic comparison. *Fuzzy Sets and Systems*, *82*, 187 – 200.

Thagard, P., & Holyoak, K. (1990). Ananlog retrieval by constrain satisfaction. *Artificial Intelligence*, *46*, 259 – 310.

Thagard, P. (1996). Evaluating computational models of cognition: Notes from the analogy wars. In Ling, C. X., & Sun, R. (Eds.), *Working Notes of AAAI-96 Workshop on Computational Cognitive Modeling: Source of the Power.*

VanLehn, K. (1995). Cognitive skill acquisition. In Spence, J., Darly, J., & Foss, D. (Eds.), *Annual Review of Psychology. Vol.47.* Annual Reviews Inc, Palo Alto, CA.

Veale, T., Smyth, B., O'Donoghue, D., & Keane, M. (1996). Representational myopia in cognitive mapping. In Ling, C. X., & Sun, R. (Eds.), *Working Notes of AAAI-96 Workshop on Computational Cognitive Modeling: Source of the Power.*

Walker, N., & Catrambone, R. (1993). Aggregation bias and the use of regression in evaluating models of human performance. *Human Factors*, *35*, 397 – 411.

Wiles, J., & Elman, J. (1996). States and stacks: Doing computation with a recurrent neural network. In Ling, C. X., & Sun, R. (Eds.), *Working Notes of AAAI-96 Workshop on Computational Cognitive Modeling: Source of the Power.*

Biography

Ron Sun is currently an assistant professor of computer science at the University of Alabama. He received his Ph.D in 1991 from Brandeis University. Dr. Ron Sun's research interest centers around 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 Cognitive Science Society. He (co)chaired several workshops on hybrid models and on cognitive modeling. He was the guest editor of the special issue of Connection Science focussing on architectures for integrating neural and symbolic processes and the special issue of IEEE Transaction on Neural Networks focussing on hybrid intelligent models. http://cs.ua.edu/faculty/sun/sun.html

Charles X. Ling obtained his B.Sc. in the Department of Computer Science, Shanghai Jiao Tong University, China, in 1985, and Ph.D. in 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. 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 (ILP), learning to control dynamic systems, knowledge-base refinement, nearest neighbor algorithms, and artificial neural networks. His WWW home page is http://www.csd.uwo.ca/faculty/ling from which most of his publications can be accessed.