Introduction to Computational Cognitive Modeling

Ron Sun

-

Instead going straight into dealing with specific approaches, issues, and domains of computational cognitive modeling, it would be more appropriate to first take some time to explore a few general questions that lie at the very core of cognitive science and computational cognitive modeling.

What is computational cognitive modeling? What exactly can it contribute to cognitive science? What has it contributed thus far? Where is it going? Answering such questions may sound overly defensive to the insiders of computational cognitive modeling, and may even seem so to some other cognitive scientists, but they are very much needed in a volume like this—because they lie at the very foundation of this field. Many insiders and outsiders alike would like to take a balanced and rational look at these questions, without indulging in excessive cheer-leading, which, as one would expect, happens sometimes amongst computational modeling enthusiasts.

However, given the large number of issues involved and the complexity of these issues, only a cursory discussion is possible in this introductory chapter. One may thus view this chapter as a set of pointers to the existing literature, rather than a full-scale discussion.

1 What is Computational Cognitive Modeling?

Research in computational cognitive modeling, or simply computational psychology, explores the essence of cognition (broadly defined, including motivation, emotion, perception, and so on) and various cognitive functionalities through developing detailed, process-based understanding by specifying corresponding computational models (in a broad sense) of representations, mechanisms, and processes. It embodies descriptions of cognition in computer algorithms and programs, based on computer science (Turing 1950). That is, it imputes computational processes (in a broad sense) onto cognitive functions, and thereby it produces runnable computational models. Detailed simulations are then conducted based on the computational models (see, e.g., Newell 1990, Rumelhart et al 1986, Sun 2002). Right from the beginning of the formal establishment of cognitive science around late 1970's, computational modeling has been a mainstay of cognitive science. [1]

In general, models in cognitive science may be roughly categorized into computational, mathematical, or verbal-conceptual models (see, e.g., Bechtel and Graham 1998). Computational models (broadly defined) present process details using algorithmic descriptions. Mathematical models presents relationships between variables using mathematical equations. Verbal-conceptual models describe entities, relations, and processes in rather informal natural languages. Each model, regardless of its genre, might as well be viewed as a *theory* of whatever phenomena it purports to capture (as argued extensively before by, for example, Newell 1990, Sun 2005).

[1] The roots of cognitive science can, of course, be traced back to much earlier times. For example, Newell and Simon's early work in the 60's and 70's has been seminal (see, e.g., Newell and Simon 1976). The work of Miller, Galanter, and Pribram (1960) has also been highly influential. See the chapter by Boden in this volume for a more complete historical perspective (see also Boden 2006).

Although each of these types of models has its role to play, in this volume, we will be mainly concerned with computational modeling (in a broad sense), including those based on computational cognitive architectures. The reason for this emphasis is that, at least at present, computational modeling (in a broad sense) appears to be the most promising approach in many respects, and it offers the flexibility and the expressive power that no other approach can match, as it provides a variety of modeling techniques and methodologies and supports practical applications of cognitive theories (Pew and Mavor 1998). In this regard, note that mathematical models may be viewed as a subset of computational models, as normally they can readily lead to computational implementations (although some of them may appear sketchy and lack process details).

Computational models are mostly process based theories. That is, they are mostly directed at answering the question of how human performance comes about, by what psychological mechanisms, processes, and knowledge structures and in what ways exactly. In this regard, note that it is also possible to formulate theories of the same phenomena through so called "product theories", which provide an accurate functional account of the phenomena but do not commit to a particular psychological mechanism or process (Vicente and Wang 1998). We may also term product theories blackbox theories or input-output theories. Product theories do not make predictions about processes (even though they may constrain processes). Thus, product theories can be evaluated mainly by product measures. Process theories, in contrast, can be evaluated by using process measures when they are available and relevant (which are, relatively speaking, rare), such as eye movement and duration of pause in serial recall; or by using product measures, such as recall accuracy, recall speed, and so on. Evaluation of process theories using the latter type of measures can only be indirect, because process theories have to generate an output given an input based on the processes postulated by the theories (Vicente and Wang 1998).

Depending on the amount of process details specified, a computational model may lie somewhere along the continuum from pure product theories to pure process theories.

There can be several different senses of "modeling" in this regard, as discussed in Sun and Ling (1998). The match of a model with human cognition may be, for example, qualitative (i.e., nonnumerical and relative), or quantitative (i.e., numerical and exact). There may even be looser "matches" based on abstracting general ideas from observations of human behaviors and then developing them into computational models. Although different senses of modeling or matching human behaviors have been used, the overall goal remains the same, which is to understand cognition (human cognition in particular) in a detailed (process-oriented) way.

This approach of utilizing computational cognitive models for understanding human cognition is relatively new. Although earlier precursors might be identified, the major developments of computational cognitive modeling have occurred since the 1960's. It has since been nurtured by the Annual Conferences of the Cognitive Science Society (which began in the late 1970's), by the International Conferences on Cognitive Modeling (which began in the 1990's), as well as by the journals of Cognitive Science (which began in the late 1970's), Cognitive Systems Research (which began in the 1990's), and so on.

From Schank and Abelson (1977) to Minsky (1981), a variety of influential symbolic "cognitive" models were proposed in Artificial Intelligence. They were usually broad and capable of a significant amount of information processing. However, they were usually not rigorously matched against human data. Therefore, it was hard to establish cognitive validity of many of these models. Psychologists have also been proposing computational cognitive models, which are usually narrower and more specific. They were usually more rigorously evaluated in relation to human data. An early example is Anderson's HAM

(Anderson 1983). Many of such models were inspired by symbolic AI work at that time (Newell and Simon 1976).

The resurgence of neural network models in the 1980's brought another type of model into prominence in this field (see, e.g., Rumelhart et al 1986, Grossberg 1982). Instead of symbolic models that rely on a variety of complex data structures that store highly structured pieces of knowledge (such as Schank's scripts or Minsky's frames), simple, uniform, and often massively parallel numerical computation was used in these neural network models (Rumelhart et al 1986). Many of these models were meant to be rigorous models of human cognitive processes, and they were often evaluated in relation to human data in a quantitative way (but see Massaro 1988).

Hybrid models that combine the strengths of neural networks and symbolic models emerged in the early 1990's (see, e.g., Sun and Bookman 1994). Such models could be used to model a wider variety of cognitive phenomena due to their more diverse and thus more expressive representations (but see Regier 2003 regarding constraints on models). They have been used to tackle a broad range of cognitive data, often (though not always) in a rigorous and quantitative way (see, for example, Sun and Bookman 1994, Sun 1994, Anderson and Lebiere 1998, Sun 2002).

For overviews of some currently existing software, tools, models, and systems for computational cognitive modeling, the reader may refer to the following Websites (among others):

`http://www.cogsci.rpi.edu/~rsun/arch.html`

`http://books.nap.edu/openbook.php?isbn=0309060966`

`http://www.isle.org/symposia/cogarch/archabs.html`

as well as the following Websites for specific software, cognitive models, or cognitive architectures (e.g., Soar, ACT-R, and CLARION):

http://psych.colorado.edu/~oreilly/PDP++/PDP++.html

http://www.cogsci.rpi.edu/~rsun/clarion.html

http://act-r.psy.cmu.edu/

http://sitemaker.umich.edu/soar/home

http://www.eecs.umich.edu/~kieras/epic.html

2 What is Computational Cognitive Modeling Good for?

There are reasons to believe that the goal of understanding the human mind strictly from observations of human behavior is ultimately untenable, except for small and limited task domains. The rise and fall of behaviorism is a case in point. This point may also be argued on the basis of analogy with physical sciences (see Sun, Coward, and Zenzen 2005). The key point is that the processes and mechanisms of the mind cannot be understood purely on the basis of behavioral experiments, with tests that inevitably amount to probing only relatively superficial features of human behavior, which are further obscured by individual/group differences and contextual factors. It would be extremely hard to understand the human mind in this way, just like it would be extremely hard to understand a complex computer system purely on the basis of testing its behavior, if we do not have any a priori ideas about the nature, the inner working, and the theoretical underpinnings of that system (Sun 2005). For a simple example, in any experiment involving the human mind, there is a very large number of parameters that could influence the results, and these parameters are either measured or left to chance. Given the large number of parameters, many have to be left to chance. The selection of which parameters to control and which to leave to chance is a decision made by the experimenter. This decision is made on the basis of which parameters the experimenter thinks are important. Therefore, clearly, theoretical development need to go hand-in-hand with experimental tests of human behavior.

Given the complexity of the human mind, and its manifestation in behavioral flexibility, complex process-based theories, that is, computational models (in the broad sense of the term), are necessary to explicate the intricate details of the human mind. Without such complex process-based theories, experimentation may be blind—leading to the accumulation of a vast amount of data without any apparent purpose or any apparent hope of arriving at a succinct, precise, and meaningful understanding. It is true that even pure experimentalists may often be guided by their intuitive theories in designing experiments and in generating their hypotheses. So, it is reasonable to say that they are in practice not completely blind. However, without detailed theories, most of the details of an intuitive (or verbal-conceptual) theory are left out of consideration, and the intuitive theory may thus be somehow vacuous, or internally inconsistent, or otherwise invalid. These problems of an intuitive theory may not be discovered until a detailed model is developed (Sun, Coward, and Zenzen 2005, Sun 2005).

There are many reasons to believe that the key to understanding cognitive processes is often in fine details, which only computational modeling can bring out (Newell 1990, Sun 2005). Computational models provide algorithmic specificity: detailed, exactly specified, and carefully thought-out steps, arranged in precise and yet flexible sequences. Therefore, they provide both conceptual clarity and precision. As related by Hintzman (1990), "The common strategy of trying to reason backward from behavior to underlying processes (analysis) has drawbacks that become painfully apparent to those who work with simulation models (synthesis). To have one's hunches about how a simple combination of processes will behave repeatedly dashed by one's own computer program is a humbling experience that no experimental psychologist should miss" (p.111).

One viewpoint concerning the theoretical status of computational modeling and simulation is that they, including those based on cognitive architectures, should not be taken as theory. A simulation/model is a generator of phenomena and data. Thus it is a theory-building tool. Hintzman (1990) gave a positive assessment of the role of simulation/model in theory building: "a simple working system that displays some properties of human memory may suggest other properties that no one ever thought of testing for, may offer novel explanations for known phenomena, and may provide insight into which modifications that next generation of models should include" (p.111). That is, computational models are useful media for thought experiments and hypothesis generation. In particular, one may use simulations for exploring various possibilities regarding details of a cognitive process. Thus, a simulation/model may serve as a theory-building tool for developing future theories. A related view is that computational modeling and simulation are suitable for facilitating the precise instantiation of a pre-existing verbal-conceptual theory (e.g., through exploring various possible details in instantiating the theory) and consequently the careful evaluation of the theory against data. A radically different position (e.g., Newell 1990, Sun 2005) is that every simulation/model provides a theory. It is not the case that a simulation/model is limited to being built on top of an existing theory, being applied for the sake of generating data, being applied for the sake of validating an existing theory, or being applied for the sake of building a future theory. To the contrary, according to this view, a simulation/model *is* a theory by itself. In philosophy of science, constructive empiricism (van Fraasen 1980) may make a sensible philosophical foundation for computational cognitive modeling, consistent with the view of models as theories (Sun 2005).

Computational models may be necessary for understanding a system as complex and as diverse as the human mind. Pure mathematics, developed to describe the physical universe, may not be sufficient for understanding a system as different and as complex as the human mind (cf. Luce 1995, Coombs et al 1970). Compared with scientific theories developed in other disciplines (e.g., in physics), computational cognitive modeling may be mathematically less elegant—but the point is that the human mind itself is likely to be less mathematically elegant compared with the physical universe (see, e.g., Minsky 1985) and therefore an alternative form of theorizing is called for, a form that is more complex, more diverse, and more algorithmic in nature. Computational cognitive models provide a viable way of specifying complex and detailed theories of cognition. Consequently, they may provide detailed interpretations and insights that no other experimental or theoretical approach can provide.

In particular, a cognitive architecture denotes a comprehensive, domain-generic computational cognitive model, capturing the essential structures, mechanisms, and processes of cognition. It is used for a broad, multiple-level, multiple-domain analysis of cognition and behavior (Sun 2004, Sun, Coward, and Zenzen 2005, Sun 2005). It deals with componential processes of cognition in a structurally and mechanistically well defined way (Sun 2004). Its function is to provide an essential framework to facilitate more detailed modeling and understanding of various components and processes of the mind. A cognitive architecture is useful and important because it provides a comprehensive initial framework for further exploration of many different domains and many different cognitive functionalities. The initial assumptions may be based on either available scientific data (e.g., psychological or biological data), philosophical thoughts and arguments, or ad hoc working hypotheses (including computationally inspired such hypotheses). A cognitive architecture helps to narrow down possibilities, provides scaffolding structures, and embodies fundamental theoretical postulates. Note that the value of cognitive architectures has been argued many times before; see, for example, Newell (1990), Anderson and Lebiere (1998), Sun (2002), Anderson and Lebiere (2003), Sun (2004), Sun, Coward, and Zenzen (2005), Sun (2005), and so on. [2]

As we all know, science in general often progresses from understanding to prediction and then to prescription (or control). Computational cognitive modeling potentially may contribute to all of these three phases of science. For instance, through process-based simulation, computational modeling may reveal dynamic aspects of cognition, which may not be revealed otherwise, and allows a detailed look at constituting elements and their interactions on the fly during performance. In turn, such understanding may lead to hypotheses concerning hitherto undiscovered or unknown aspects of cognition and may lead to predictions regarding cognition. The ability to make reasonably accurate predictions about cognition can further allow prescriptions or control, for example, by choosing appropriate environmental conditions for certain tasks, or by choosing appropriate mental types for certain tasks and/or environmental conditions.

In sum, the utility and the value of computational cognitive modeling (including cognitive architectures) can be argued in many different ways (see Newell 1990, Sun 2002, Anderson and Lebiere 2003, and so on). These models in their totality are clearly more than just simulation tools or programming languages of some sorts. They are theoretically pertinent, because they represent theories in a unique and, I believe, indispensable way. Cognitive architectures, for example, are broad theories of cognition in fact.

[2] For information about different existing cognitive architectures, see, for example, http://www.cogsci.rpi.edu/∼rsun/arch.html. See also Sun (2006) for information on three major cognitive architectures.

3 Multiple Levels of Computational Cognitive Modeling

A strategic decision that one has to make with respect to cognitive science is the level(s) of analysis (i.e., level(s) of abstraction) at which one models cognitive agents. Computational cognitive modeling can vary in terms of level of process details and granularity of input and output, and thus may be carried out at multiple levels. Let us look into this issue of multiple levels of computational cognitive modeling, drawing upon the work of Sun, Coward, and Zenzen (2005).

We note that traditional theories of multi-level analysis holds that there are various levels each of which involves a different amount of computational details (e.g., Marr 1982). In Marr's theory, first, there is the *computational theory* level, in which one is supposed to determine proper computation to be performed, its goals, and the logic of the strategies by which the computation is to be carried out. Second, there is the *representation and algorithm* level, in which one is supposed to be concerned with carrying out the computational theory determined at the first level and, in particular, the representation for the input and the output and the algorithm for the transformation from the input to the output. The third level is the *hardware implementation* level, in which one is supposed to physically realize the representation and algorithms determined at the second level. According to Marr, these three levels are only loosely coupled; that is, they are relatively independent. Thus there are usually a wide array of choices at each level, independent of the other two. Some phenomena may be explained at only one or two levels. Marr (1982) emphasized the "critical" importance of formulation at the level of computational theory, that is, the level at which the goals and purposes of a cognitive process are specified and internal and external constraints that make the process possible are worked out and related to each other and to the goals of computation. His reason was that the nature of compu-

level	object of analysis
1	computation
2	algorithms
3	implementations

Figure 1: A traditional hierarchy of levels (Marr 1982).

level	object of analysis	type of analysis	computational model
1	inter-agent processes	social/cultural	collections of agents
2	agents	psychological	individual agents
3	intra-agent processes	componential	modular construction of agents
4	substrates	physiological	biological realization of modules

Figure 2: Another hierarchy of four levels (Sun, Coward, and Zenzen 2005).

tation depended more on the computational problems to be solved than on the way the solutions were to be implemented. In his own words, "an algorithm is likely to be understood more readily by understanding the nature of the problem being solved than by examining the mechanism (and the hardware) in which it is embodied." Thus, he preferred a top-down approach—from a more abstract level to a more detailed level. See Figure 1 for the three levels. It often appears that Marr's theory centered too much on the relatively minor differences in computational abstractions (e.g., algorithms, programs, and implementations; see Sun, Coward, and Zenzen 2005, Dayan 2003, Dawson 2002). It also appears that his theory represented an over-simplification of biological reality (for example, ignoring the species-specific or motivation-relevant representations of the environment and the close relationship between low-level implementations and high-level computation), and as a result represented an over-rationalization of cognition.

Another variant is Newell and Simon's three-level theory. Newell and Simon

(1976) proposed the following three levels: (1) The knowledge level, in which why cognitive agents do certain things is explained by appealing to their goals and their knowledge, and by showing rational connections between them. (2) The symbol level, in which the knowledge and goals are encoded by symbolic structures, and the manipulation of these structures implements their connections. (3) The physical level, in which the symbol structures and their manipulations are realized in some physical form. Sometimes this three-level organization was referred to as “the classical cognitive architecture” (Newell 1990). The point being emphasized here was very close to Marr’s view: What is important is the analysis at the knowledge level and then at the symbol level, that is, identifying the task and designing symbol structures and symbol manipulation procedures suitable for it. Once this analysis (at these two levels) is worked out, the analysis can be implemented in any available physical means.

In contrast, according to Sun, Coward, and Zenzen (2005), the differences (borrowed from computer programming) amongst “computation”, algorithms, programs, and hardware realizations, and their variations, as have been the focus in Marr’s (1982) and Newell and Simon’s (1976) level theories, are relatively insignificant. This is because, first of all, the differences among them are usually small, fuzzy, and subtle, compared with the differences among the processes to be modeled (that is, the differences among the sociological vs. the psychological vs. the intra-agent, etc.). Second, these different computational constructs are in reality closely tangled (especially in the biological world): One cannot specify algorithms without at least some considerations of possible implementations, and what is to be considered “computation” (i.e., what can be computed) relies on algorithms, especially the notion of algorithmic complexity, and so on. Therefore, one often has to consider computation, algorithms, and implementation together somehow (especially in relation to cognition). Third, according to Sun, Coward, and Zenzen (2005), the separation of these computational details failed to produce any major useful insight in relation to cognition, but theoretical baggage. A re-orientation toward a systematic examination of *phenomena*, instead of *tools* one uses for modeling them, is thus a step in the right direction.

The viewpoint of Sun, Coward, and Zenzen (2005) focused attention on the very phenomena to be studied, on their scopes, scales, degrees of abstractness, and so on. Thus, the differences among levels of analysis can be roughly cast as the differences among disciplines, from the most macroscopic to the most microscopic. These levels of analysis include: the sociological level, the psychological level, the componential level, and the physiological level. See Figure 2 for these levels. Different levels of modeling may be established in exact correspondence with different levels of analysis.

First of all, there is the sociological level, which includes collective behavior of agents (Durkheim 1895), inter-agent processes (Vygotsky 1986), sociocultural processes, as well as interaction between agents and their (physical and sociocultural) environments. Only recently, the field of cognitive science has come to grip with the fact that cognition is, at least in part, a social/cultural process (Lave 1988, Vygotsky 1986, Sun 2006). To ignore the sociocultural process is to ignore a major underlying determinant of individual cognition. The lack of understanding of sociological processes may result in the lack of understanding of some major structures and constraints in cognition. Thus, any understanding of individual cognition can only be partial and incomplete when sociocultural processes are ignored or downplayed. [3]

The next level is the psychological level, which covers individual behaviors, beliefs, knowledge, concepts, and skills (as well as motivation, emotion, perception, and so on). In relation to the sociological level, one can investigate the

[3] See Sun (2001, 2006) for a more detailed argument of the relevance of sociocultural processes to cognition and vice versa.

relationship of individual beliefs, knowledge, concepts, and skills with those of the society and the culture, and the processes of change of these beliefs, knowledge, concepts, and skills, independent of or in relation to those of the society and the culture. At this level, one can examine human behavioral data, and compare them with models and with insights from the sociological level and further details from the lower levels.

The third level is the componential level. It is important to note that in computational cognitive modeling, the computational process of an agent is mostly specified in terms of *components* of the agent, i.e., in terms of intra-agent processes. Thus, at this level, one may specify a cognitive architecture and components therein. In the process of analysis, one specifies essential computational processes of each component as well as essential connections among various components. Thus, analysis of capacity (functional analysis) and analysis of components (structural analysis) become one and the same at this level. However, at this level, unlike at the psychological level, work is more along the line of structural analysis than functional analysis (while the psychological level is mostly concerned with functional analysis). At this level, one models cognitive agents in terms of components, with the theoretical language of a particular paradigm, for example, symbolic computation or connectionist networks, or their combinations (Sun and Bookman 1994). That is, one imputes a computational process onto a cognitive function. Ideas and data from the psychological level—the psychological constraints from above, which bear on the division of components and possible implementations of components, are among the most important considerations. This level may also incorporate biological/physiological observations regarding plausible divisions and their implementations; that is, it can incorporate ideas from the next level down—the physiological level, which offers the biological constraints. This level results in cognitive *mechanisms*, although they are usually computational and thus abstract, compared with physiological-level specifications of details.

Although this level is essentially in terms of intra-agent processes, computational models developed therein may also be used to model processes at higher levels, including the interaction at a sociological level where multiple individuals are involved. This can be accomplished, for example, by examining interactions of multiple copies of individual agents (Sun 2006).

The lowest level of analysis is the physiological level, that is, the biological substrate, or biological implementation, of computation (Dayan 2003). This level is the focus of a range of disciplines including physiology, biology, computational neuroscience, cognitive neuroscience, and so on. Although biological substrates are not among our major concerns here, they may nevertheless provide valuable input as to what kind of computation is likely employed and what a plausible architecture (at a higher level) should be like. The main utility of this level is to facilitate analysis at higher levels, that is, to use low-level information to narrow down, at higher levels, choices in selecting computational architectures and choices in implementing componential computation.

Although computational cognitive modeling is often limited to within a particular level at a time (inter-agent, agent, intra-agent, or substrate), this need not always be the case: Cross-level analysis and modeling could be intellectually highly enlightening, and might be essential to the progress of computational cognitive modeling in the future (Sun, Coward, and Zenzen 2005, Dayan 2003). These levels described above do interact with each other (e.g., constraining each other) and may not be easily isolated and tackled alone. Moreover, their respective territories are often intermingled, without clear-cut boundaries.

For instance, the cross-level link between the psychological and the neurophysiological level has been strongly emphasized in recent years (in the form of cognitive neuroscience; see, e.g., LeDoux 1992, Damasio 1994, Milner and Goodale 1995). For example, Wilson et al. (2000) presented a model of human subjects perceiving the orientation of the head of another person. They accounted for the empirical findings from psychological experiments with a model based on a population code of neurons in the visual cortex, and thus the underlying neural structures were used to explain a psychological phenomenon at a higher level. For another instance of cross-level research, the psychological and the social level may also be crossed in many ways, in order to generate new insights into social phenomena on the basis of cognitive processes (e.g., Boyer and Ramble 2001, Sun 2006) and, conversely, to generate insights into cognitive phenomena on the basis of sociocultural processes (e.g., Hutchins 1995, Nisbett et al 2001). In all of these cases, the ability to shift appropriately between levels when needed is a critical part of the work.

Beyond cross-level analysis, there may be "mixed-level" analysis (Sun, Coward, and Zenzen 2005). The idea of mixed-level analysis may be illustrated by the research at the boundaries of quantum mechanics. In deriving theories, physicists often start working in a purely classical language that ignores quantum probabilities, wave functions, and so forth, and subsequently overlay quantum concepts upon a classical framework (Greene 1999, Coward and Sun 2004). The very same idea applies to mixing cognitive modeling and social simulation as well. One may start with purely social descriptions but then substitute cognitive principles and cognitive process details for simpler descriptions of agents (e.g., Sun and Naveh 2004). Relatedly, there has also been strong interplay between psychological models and neurophysiological models—for example, going from psychological descriptions to neurobiological details.

Note that Rasmussen (1986) proposed something similar to the view described above on levels. His hierarchy was a more general framework but had a number of constraining properties (see also Vicente and Wang 1998): (1) All levels deal with the same system, with each level providing a different description of the system; (2) each level has its own terms, concepts, and principles; (3)

the selection of levels may be dependent on the observer's purpose, knowledge, and interest; (4) the description at any level may serve as constraints on the operation of lower levels, whereas changes at a higher level may be specified by the effects of the lower levels; (5) by moving up the hierarchy, one understands more the significance of some process details with regard to the purpose of the system; by moving down the hierarchy, one understands more how the system functions in terms of the process details; (6) there is also a means-ends relationship between levels in a hierarchy.

Note also Ohlsson and Jewett's (1997) and Langley's (1999) idea of abstract cognitive model, which is relevant here as well. To guard against over-interpretation of empirical evidence and to avoid the usually large gaps between evidence and full-blown computational models, Ohlsson and Jewett (1997) proposed "abstract computational models", which were relatively abstract models that were designed to test a particular (high level) hypothesis without taking a stand on all the (lower level) details of a cognitive architecture. Similar ideas were also expressed by Langley (1999), who argued that the source of explanatory power of a model often lay at a higher level of abstraction.

In sum, there have been various proposals regarding multiple levels of computational cognitive modeling. Although details vary, the very notion of multiple levels of cognitive modeling appears to be useful. It can be expected to be of importance for the further development of this field.

4 Success Stories of the Past

There have been quite a few success stories of computational cognitive modeling, in a practical or a theoretical sense. They include, among many others:

- the various models of developmental psychology, including the connectionist models of verb past-tense learning and the controversies stemming from such models,

- the tutoring systems based on the ACT-R cognitive architecture,
- the model of implicit and explicit learning based on the CLARION cognitive architecture.

For instance, computational models of child development have been successful in accounting for, and in explaining, fine-grained developmental processes. In terms of widespread impact and associated theoretical interests and controversies, computational models of verb past-tense learning may be ranked as being at the top of all computational cognitive models (see, e.g., Rumelhart et al 1986).

Theoretically, successful development models have clarified a number of major issues. In developmental psychology, there is the dichotomy contrasting knowledge that the child acquires through interacting with the environment (nurture) with knowledge of phylogenic origin (nature). It was argued that mechanisms of gene expression and brain development did not allow for the detailed specification of neural networks in the brain as required by the nativist position. It has been argued that a more plausible role for innate knowledge is at the level of architectures and timing of development (see the chapter by Shultz and Sirois in this volume). In this regard, neural network models have provided new ways of thinking about innateness. That is, instead of asking whether or not something is innate, one should ask how evolution constrains the emergence of a brain function during individual development. This kind of theorizing has benefited from the use of neural networks (as detailed in the chapter by Shultz and Sirois).

Developmental psychologists have also been debating the distinction between learning and development. A static neural network can only learn what is within its representational power. Thus, when static neural networks are used, it is assumed that the ultimate brain network topology has already been developed (even if initial weights are random). However, this assumption implies representational innateness, which has been argued to be implausible. An alternative is to use neural network models that form their network topology as a result of their experience. Using constructive learning models also resolves the "paradox of development": It was argued that if learning was done by proposing and testing hypotheses, it was not possible to learn anything that could not already be represented. This argument becomes irrelevant in light of constructive learning models where learning mechanisms that construct representations are separate from the representation of domain-specific knowledge. A constructive model builds representational power that it did not previously possess. Thus, computational modeling suggests that development is functionally distinct from learning (as argued in the chapter by Shultz and Sirois).

Similarly, as another example, an interpretation of a broad range of skill learning data (including those from the implicit learning literature) was proposed based on the CLARION cognitive architecture (see Sun, Slusarz, and Terry 2005 and Sun 2002; see also the chapter by Taatgen and Anderson in this volume concerning cognitive architectures). At a theoretical level, this work explicates the interaction between implicit and explicit cognitive processes in skill learning, in contrast to the tendency of studying each type in isolation. It highlights the interaction between the two types of processes and its various effects on learning (including the so called synergy effects; see Sun 2002). At an empirical level, a model centered on such an interaction constructed based on CLARION was used to account for data in a variety of task domains: process control tasks, artificial grammar learning tasks, serial reaction time tasks, as well as some much more complex task domains (such as Tower of Hanoi and Minefield Navigation). The model was able to explain data in these task domains, shedding light on some apparently contradictory findings (including some findings once considered as casting doubt on the theoretical status of implicit learning). Based on the data and the match between the CLARION architecture and the data, this work argues for an integrated theory/model of skill learning that takes into account both implicit and explicit processes, as the data match pointed to the usefulness of incorporating both explicit and implicit processes in theorizing about cognition (Sun, Slusarz, and Terry 2005). Moreover, it argues for a bottom-up approach (first learning implicit knowledge and then explicit knowledge on its basis) in an integrated theory/model of skill learning, which was radically different from the then existing models (see Sun 2002; see also the chapter on skill learning by Ohlsson in this volume). So, in this case, the application of the computational cognitive architecture CLARION to the skill learning data helped to achieve a level of theoretical integration and explanation beyond the previous theorizing (Sun, Slusarz, and Terry 2005; Sun 2002). For yet another example of using cognitive architectures to provide theoretical interpretation and integration, see Meyer and Kieras (1997).

As a final example, a number of interesting tutoring systems have been constructed on the basis of the ACT-R cognitive architecture (Koedinger et al 1997; see also the chapter by Taatgen and Anderson in this volume). These tutoring systems were based on the analysis of the task units that were necessary to achieve competence in a number of domains of mathematics and computer programming. These units were represented as production rules. A typical course involves on the order of 500 production rules. On the assumption that learning in these domains involves the acquisition of such production rules, it is possible to diagnose whether students have acquired such production rules and provide instruction to remedy any difficulties they might have with specific rules. This led to the design of tutoring systems that ran production rule models in parallel with a student and attempted to interpret the student behavior in terms of these rules. Such systems tried to find some sequence of production rules that produced the behavior exhibited by a student. The model-tracing process allowed the interpretation of student behavior, and in turn the interpretation controlled the tutorial interactions. Thus, such tutoring systems are predicated on the validity of the cognitive model and the validity of the attributions that the model-tracing process makes about student learning. There have been a few assessments that established to some extent the effectiveness of these systems. The tutoring systems have been used to deliver instruction to more than 100,000 students thus far. They demonstrated the practical usefulness of computational cognitive modeling. Other examples of practical applications of computational cognitive modeling may be found in Pew and Mavor (1998), and many in the area of human-computer interaction.

5 Directions for the Future

Many accounts of the history and the current state of the art of computational cognitive modeling in different areas will be provided by the subsequent chapters in this volume. At this point, however, it may be worthwhile to speculate a little about future developments of computational cognitive modeling.

First of all, some have claimed that grand scientific theorizing has become a thing of the past. What remains to be done is filling in details and refining some minor points. Fortunately, many cognitive scientists believe otherwise. Indeed, many of them are pursuing integrative principles that attempt to explain data in multiple domains and in multiple functionalities (e.g., Anderson and Lebiere 1998, Sun 2002). In cognitive science, as in many other scientific fields, significant advances may be made through discovering (hypothesizing and confirming) deep-level principles that unify superficial explanations across multiple domains, in a way somewhat analogous to Einstein's theory that unified electromagnetic and gravitational forces, or String Theory that aims to provide even further unifications (see Green 1999). Such theories are what cognitive science needs, currently and in the foreseeable future.

Integrative computational cognitive modeling may serve in the future as an antidote to the increasing specialization of scientific research. In particular, cognitive architectures are clearly going against the trend of increasing specialization, and thus constitute an especially effective tool in this regard. Cognitive scientists are currently actively pursuing such approaches and, hopefully, will be increasingly doing so in the future. In many ways, the trend of over-specialization is harmful, and thus the reversal of this trend by the means of computational cognitive modeling is a logical (and necessary) next step toward advancing cognitive science (Sun et al 1999).

Second, related to the point above, while the importance of being able to reproduce the nuances of empirical data from specific psychological experiments is evident, broad functionality is also important (Newell 1990). The human mind needs to deal with the full cycle that includes all of the followings: transducing signals, processing them, storing them, representing them, manipulating them, and generating motor actions based on them. In computational cognitive modeling, there is clearly a need to develop generic models of cognition that are capable of a wide range of cognitive functionalities, to avoid the myopia often resulting from narrowly-scoped research (e.g., in psychology). In particular, cognitive architectures may incorporate all of the following cognitive functionalities: perception, categorization and concepts, memory, decision making, reasoning, planning, problem solving, motor control, learning, metacognition, motivation, emotion, language and communication, among others. In the past, this issue often did not get the attention it deserved in cognitive science (Newell 1990), and it remains a major challenge for cognitive science.

However, it should be clearly recognized that over-generality, beyond what is minimally necessary, is always a danger in computational cognitive modeling, and in developing cognitive architectures (Sun 2007). It is highly desirable to come up with a well constrained cognitive model with as few parameters as possible while accounting for as large a variety of empirical observations and phenomena as possible (Regier 2003). This may be attempted by adopting a broad perspective — philosophical, psychological, biological, as well as computational, and by adopting a multi-level framework going from sociological, to psychological, to componential, and to physiological levels, as discussed before (and as argued in more detail in Sun, Coward, and Zenzen 2005). Although some techniques have been developed to accomplish this, more work is needed (see, e.g., Sun and Ling 1998, Regier 2003, Sun 2007).

Third, in integrative computational cognitive modeling, especially in developing cognitive architectures with a broad range of functionalities, it is important to keep in mind a broad set of desiderata. For example, in Anderson and Lebiere (2003), a set of desiderata proposed by Newell (1990) was used to evaluate a cognitive architecture versus conventional connectionist models. These desiderata include flexible behavior, real-time performance, adaptive behavior, vast knowledge base, dynamic behavior, knowledge integration, natural language, learning, development, evolution, and brain realization (see Newell 1990 for detailed explanations). In Sun (2004), another, broader set of desiderata was proposed and used to evaluate a larger set of cognitive architectures. These desiderata include ecological realism, bio-evolutionary realism, cognitive realism, and many others (see Sun 2004 for details). The advantages of coming up with and applying these sets of desiderata in computational cognitive modeling include (1) avoiding overly narrow models and (2) avoiding missing important functionalities. We can reasonably expect that this issue will provide impetus for further research in the field of computational cognitive modeling in the future.

Fourth, the validation of process details of computational cognitive models has been a difficult, but extremely important, issue (Pew and Mavor 1998). This is especially true for cognitive architectures, which often involve a great deal of intricate details that are almost impossible to disentangle. This issue needs to be better addressed in the future. There have been too many instances in the past that research communities rushed into some particular model or some particular approach toward modeling cognition and human intelligence, without knowing exactly how much of the approach or the model was veridical or even useful. Theoretical (including mathematical) analysis often lagged behind. Thus, often without sufficient effort at validation and theoretical analysis, claims were boldly made about the promise of a certain model or a certain approach. Unfortunately, we have seen quite a few setbacks in the history of cognitive science as a result of this cavalier attitude toward the science of cognition. As in any other scientific field, painstakingly detailed work needs to be carried out in cognitive science, before sweeping claims can be made. Not only is empirical validation necessary, theoretical analysis, including detailed mathematical and computational analysis, is also necessary in order to better understand models and modeling approaches, before committing a large amount of resource (cf. Roberts and Pashler 2000). In particular, sources of explanatory power need to be identified and analyzed (as called for in Sun and Ling 1998). The issue of validation should be an important factor in directing future research in the field of computational cognitive modeling.

Related to that, the "design" space of computational cognitive models needs to be more fully explored (as pointed out in Sun and Ling 1998 and Sloman and Chrisley 2005). While we explore the behavioral space, in the sense of identifying the range and variations of human behavior, we also need to explore the design space (that is, all the possibilities for constructing computational models) that maps onto the behavioral space, so that we may gain a better understanding of the possibilities and the limitations of modeling methodologies, and thereby open up new avenues for better capturing cognitive processes. This is especially important for cognitive architectures, which are complex and in which many design decisions need to be made, often without the benefit of a clear understanding of their full implications in computational or behavioral terms. More systematic exploration of the design space of cognitive models is thus necessary. Future research in this field should increasingly address this issue (Sloman and Chrisley 2005).

Computational cognitive models may find both finer and broader applications, that is, both at lower levels and at higher levels, in the future. For example, some cognitive models found applications in large-scale simulations at a social and organizational level. For another example, some other cognitive models found applications in interpreting not only psychological data but also neuroimaging data (at a biological/physiological level). A review commissioned by the National Research Council found that computational cognitive modeling had progressed to a degree that had made them useful in a number of application domains (Pew and Mavor 1998). Another review (Ritter, Shadbolt, Elliman, Young, Gobet, and Baxter 2003) pointed to similar conclusions. Both reviews provided interesting examples of applications of computational cognitive modeling. Inevitably, this issue will provide impetus for future research, not only in applied areas of computational cognitive modeling, but also in theoretical areas of computational cognitive modeling.

In particular, cognitive modeling may be profitably applied to social simulation. An important recent development in the social sciences has been agent-based social simulation. [4] So far, however, the two fields of social simulation and cognitive modeling have been developed largely separately from each other (with some exceptions). Most of the work in social simulation assumed rudimen-

[4] This approach consists of instantiating a population of agents, allowing the agents to run, and observing the interactions among them.

tary cognition on the part of the agents. As has been argued before (e.g., Sun and Naveh 2004; Sun 2001, 2006; Zerubavel 1997), social processes ultimately rest on the decisions of individuals, and thus understanding the mechanisms of individual cognition can lead to better theories of social processes. At the same time, by integrating social simulation and cognitive modeling, we may arrive at a better understanding of individual cognition. By modeling cognitive agents in a social context (as in cognitive social simulation), we may learn more about how sociocultural processes influence individual cognition. (See the later chapter by Ron Sun in this volume regarding cognitive social simulation.)

Cross-level and mixed-level work integrating the psychological and the neurophysiological level, as discussed before, will certainly be an important direction for future research. Increasingly, researchers are exploring constraints from both psychological and neurobiological data. In so doing, the hope is that more realistic and better constrained computational cognitive models may be developed. (See, for example, the chapter by Norman et al in this volume for some such models.)

Finally, will this field eventually become a full fledged discipline—computational psychology, just like computational neuroscience or computational physics? This is an interesting but difficult issue. There are a number of open questions in this regard. For example, how independent can this field be from closely allied fields such as experimental psychology (and cognitive psychology in particular)? What will the relationship be between data generation and modeling? How useful or illuminating can this field be in shedding new light on cognition per se (as opposed to leading up to building intelligent systems)? And so on and so forth. These are the questions that will determine the future status of this field. So far, the answers to these questions are by no means clear-cut. They will have to be worked out in the future through the collective effort of the researchers in this field.

6 About This Book

The present volume, the *Cambridge Handbook of Computational Cognitive Modeling*, is part of the *Cambridge Handbook in Psychology* series. This volume is aimed to be a definitive reference source for the growing field of computational cognitive modeling. Written by the leading experts in various areas of this field, it is meant to combine breadth of coverage with depth of critical details.

This volume aims to appeal to researchers and advanced students in the computational cognitive modeling community, as well as to researchers and advanced students in cognitive science (in general), philosophy, experimental psychology, linguistics, cognitive anthropology, neuroscience, artificial intelligence, and so on. For example, it could serve well as a textbook for courses in social, cognitive, and behavioral sciences programs. In addition, this volume might also be useful to social sciences researchers, education researchers, intelligent system engineers, psychology and education software developers, and so on.

Although this field draws on many humanity and social sciences disciplines and on computer science, the core of the approach is based on psychology, and this is a constant focus in this volume. At the same time, this volume is also distinguished by its incorporation of one contemporary theme in scientific research: how technology (namely computing technology) affects our understanding of the subject matter—cognition and its associated issues.

This volume contains 26 chapters, organized into 4 parts. The first part (containing the present chapter) provides a general introduction to the field of computational cognitive modeling. The second part, *Cognitive Modeling Paradigms*, introduces the reader to broadly influential approaches in cognitive modeling. These chapters have been written by some of those influential scholars who helped to define the field. The third part, *Computational Modeling of Various Cognitive Functionalities and Domains*, describes a range of computational modeling efforts that researchers in this field have undertaken regarding major cognitive functionalities and domains. The interdisciplinary combination of cognitive modeling, experimental psychology, linguistics, artificial intelligence, and software engineering in this field has required researchers to develop a novel set of research methodologies. This part surveys and explains computational modeling research, in terms of detailed computational mechanisms and processes, on memory, concepts, learning, reasoning, decision making, skills, vision, motor control, language, development, scientific explanation, social interaction, and so on. It contains case studies of projects, as well as details of significant models in the computational cognitive modeling field. These chapters have been written by some of the best experts in these areas. The final part, *Concluding Remarks*, explores a range of issues associated with computational cognitive modeling and cognitive architectures, and provides some perspectives, evaluations, and assessments.

Although our goal has been to be as comprehensive as possible, the coverage of this volume is, by necessity, selective. The selectivity is made necessary by the length limitation, as well as by the amount of activities in various topic areas — we need to cover areas with large amounts of scholarly activities, inevitably at the cost of less active areas. Given the wide-ranging and often fast-paced research activities in computational cognitive modeling, I never had any trouble in finding interesting topics to include, but I often found myself in a position whereby I had to sacrifice some less active topics.

As research in this field has developed at an exciting pace in recent years, the field is ready for an up-to-date reference to the best and latest work. For this field, what has been missing is a true handbook. Such a handbook should bring together top researchers to work on chapters each of which summarizes and explains the basic concepts, techniques, and findings for a major topic area, sketching its history, assessing its successes and failures, and outlining the directions in which it is going. A handbook should also provide quick overviews for experts as well as provide an entry point into the field for the next generation of researchers. The present volume has indeed been conceived with these broad and ambitious goals in mind.

7 Conclusions

It is clear that highly significant progress has been made in recent decades in advancing research on computational cognitive modeling (i.e., computational psychology). However, it appears that there is still a very long way to go before we fully understand the computational processes of the human mind.

Many examples of computational cognitive modeling are presented in this volume. However, it is necessary to explore and study more fully various possibilities in computational cognitive modeling in order to further advance the state of the art in understanding the human mind through computational means. In particular, it would be necessary to build integrative cognitive models with a wide variety of functionalities, that is, to build cognitive architectures, so that they can exhibit and explain the full range of human behaviors (as discussed before). Many challenges and issues need to be addressed, including those stemming from designing cognitive architectures, from validation of cognitive models, and from the applications of cognitive models to various domains.

It should be reasonable to expect that the field of computational cognitive modeling will have profound impact on cognitive science, as well as on other related fields such as linguistics, philosophy, experimental psychology, and artificial intelligence, both in terms of better understanding cognition and in terms of developing better (more intelligent) computational systems. As such, it should be considered a crucial field of scientific research, lying at the intersection of a number of other important fields. Through the collective effort of this research community, significant advances can be achieved, especially in better understanding the human mind.

Acknowledgments

This work was carried out while the author was supported in part by ARI grants DASW01-00-K-0012 and W74V8H-04-K-0002 (to Ron Sun and Bob Mathews). Thanks are due to Aaron Sloman and Frank Ritter for their comments on the draft.

References

J. R. Anderson, (1983). *The Architecture of Cognition.* Harvard University Press, Cambridge, MA

J. R. Anderson and C. Lebiere, (1998). *The Atomic Components of Thought.* Lawrence Erlbaum Associates, Mahwah, NJ.

J. R. Anderson and C. Lebiere, (2003). The Newell Test for a theory of cognition. *Behavioral and Brain Sciences.* 26, 587-640.

W. Bechtel and G. Graham (eds.), (1998). *A Companion to Cognitive Science.* Blackwell Publishers, Cambridge, UK.

M. Boden, (2006). *Mind as Machine: A History of Cognitive Science.* Oxford University Press, Oxford, UK.

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

L. Coward and R. Sun, (2004). Criteria for an effective theory of consciousness and some preliminary attempts. *Consciousness and Cognition*, Vol.13, pp.268-301.

A. Damasio, (1994). *Descartes' Error: Emotion, Reason and the Human Brain.* Grosset/Putnam, New York.

W. Durkheim, (1895/1962). *The Rules of the Sociological Method.* The Free Press, Glencoe, IL.

M. Dawson, (2002). Computer modeling of cognition: Levels of analysis. In: Nadel, L. (ed.), *Encyclopedia of Cognitive Science.* pp. 635-638. Macmillan, London, UK.

P. Dayan, (2003). Levels of analysis in neural modeling. In: L. Nadel (ed.), *Encyclopedia of Cognitive Science.* Macmillan, London.

G. Greene, (1999). *The Elegant Universe.* Norton, New York.

S. Grossberg, (1982). *Studies of Mind and Brain: Neural Principles of Learning, Perception, Development, Cognition, and Motor Control.* Norwell, MA: Kluwer Academic Publishers.

D. Hintzman, (1990). Human learning and memory: Connections and dissociations. In: *Annual Review of Psychology*, pp.109-139. Annual Reviews Inc, Palo Alto, CA.

E. Hutchins, (1995). How a cockpit remembers its speeds. *Cognitive Science*, 19, 265-288.

K. Koedinger, Anderson, J. R., Hadley, W. H., and Mark, M. (1997). Intelligent tutoring goes to school in the big city. *International Journal of Artificial Intelligence in Education*, 8, 30-43.

P. Langley, (1999). Concrete and abstract models of category learning. In: *Proceedings of the 21st Annual Conference of the Cognitive Science Society.* Erlbaum, Mahwah, NJ.

J. Lave, (1988). *Cognition in Practice.* Cambridge University Press, Cambridge, England.

J. LeDoux, (1992). Brain mechanisms of emotion and emotional learning. In: *Current Opinion in Neurobiology.* Vol.2, No.2, 191-197.

R. D. Luce, (1995). Four tensions concerning mathematical modeling in psychology. *Annual Review of Psychology*, 46, 1-26. Annual Reviews Inc, Palo Alto, CA.

D. Marr, (1980). *Vision.* MIT Press, Cambridge, MA.

D. Massaro, (1988). Some criticisms of connectionist models of human performance. *Journal of Memory and Language*, 27, 213-234.

D. Meyer and D. Kieras, (1997). A computational theory of executive cognitive processes and human multiple-task performance: Part 1, basic mechanisms. *Psychological Review.* 104 (1), 3-65.

G. Miller, E. Galanter, and K. Pribram, (1960). *Plans and the Structure of Behavior.* Holt, Rinehart, and Winston, New York.

D. Milner and N. Goodale, (1995). *The Visual Brain in Action.* Oxford University Press, New York.

M. Minsky, (1981). A framework for representing knowledge. In: J. Haugeland (ed.), *Mind Design*, 95-128. MIT Press, Cambridge, MA.

M. Minsky, (1985). *The Society of Mind.* Simon and Schuster, New York.

A. Newell, (1990). *Unified Theories of Cognition.* Harvard University Press, Cambridge, MA.

A. Newell and H. Simon, (1976). Computer science as empirical inquiry: Symbols and search. *Communication of ACM*, 19, 113-126.

R. Nisbett, K. Peng, I. Choi, A. Norenzayan, (2001). Culture and systems of thought: holistic versus analytic cognition. *Psychological Review.* 108 (2), 291-310.

S. Ohlsson and J. Jewett, (1997). Simulation models and the power law of learning. In: *Proceedings of the 19th Annual Conference of the Cognitive Science Society*. Erlbaum, Mahwah, NJ.

R. W. Pew and A. S. Mavor (eds), (1998). *Modeling Human and Organizational Behavior: Application to Military Simulations*. National Academy Press, Washington, DC.

J. Rasmussen, (1986). *Information Processing and Human-Machine Interaction: An Approach to Cognitive Engineering*. North-Holland, Amsterdam, Netherlands.

T. Regier, (2003). Constraining computational models of cognition. In: L. Nadel (ed.), *Encyclopedia of Cognitive Science*. Macmillan, London. pp.611-615.

F. E. Ritter, Shadbolt, N., Elliman, D., Young, R., Gobet, F., and Baxter, G., (2003). *Techniques for Modeling Human Performance in Synthetic Environments: A Supplementary Review*. Human Systems Information Analysis Center, Wright-Patterson Air Force Base, Dayton, OH.

S. Roberts and H. Pashler, (2000). How persuasive is a good fit? A comment on theory testing. *Psychological Review*, 107 (2), 358-367.

D. Rumelhart, J. McClelland and the PDP Research Group, (1986). *Parallel Distributed Processing: Explorations in the Microstructures of Cognition*. MIT Press, Cambridge, MA.

R. Schank and R. Abelson, (1977). *Scripts, Plans, Goals, and Understanding: An Inquiry into Human Knowledge Structures*. Lawrence Erlbaum Associates, Hillsdale, NJ.

A. Sloman and R. Chrisley, (2005). More things than are dreamt of in your biology: Information processing in biologically-inspired robots. *Cognitive Systems Research*, 6 (2), 145-174.

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

R. Sun, (2001). Cognitive science meets multi-agent systems: A prolegomenon. *Philosophical Psychology*, 14 (1), 5-28.

R. Sun, (2002). *Duality of the Mind*. Lawrence Erlbaum Associates, Mahwah, NJ.

R. Sun, (2004). Desiderata for cognitive architectures. *Philosophical Psychology*, 17 (3), 341-373.

R. Sun, (2005). Theoretical status of computational cognitive modeling. Technical report, Cognitive Science Department, Rensselaer Polytechnic Institute, Troy, New York.

R. Sun (ed.), (2006). *Cognition and Multi-Agent Interaction: From Cognitive Modeling to Social Simulation*. Cambridge University Press, New York.

R. Sun, (2007). The importance of cognitive architectures: An analysis based on CLARION. *Journal of Experimental and Theoretical Artificial Intelligence*, in press.

R. Sun and L. Bookman (eds.), (1994). *Computational Architectures Integrating Neural and Symbolic Processes*. Kluwer Academic Publishers, Boston, MA.

R. Sun, A. Coward, and M. Zenzen, (2005). On levels of cognitive modeling. *Philosophical Psychology*, 18 (5), pp.613-637.

R. Sun, V. Honavar, and G. Oden, (1999). Integration of cognitive systems across disciplinary boundaries. *Cognitive Systems Research*, Vol.1, No.1, pp.1-3.

R. Sun and C. Ling, (1998). Computational cognitive modeling, the source of power and other related issues. *AI Magazine*. Vol.19, No.2, pp.113-120.

R. Sun and I. Naveh, (2004). Simulating organizational decision-making using a cognitively realistic agent model. *Journal of Artificial Societies and Social Simulation*, Vol.7, No.3, June, 2004. http://jasss.soc.surrey.ac.uk/7/3/5.html

R. Sun, P. Slusarz, and C. Terry, (2005). The interaction of the explicit and the implicit in skill learning: A dual-process approach. *Psychological Review*, 112 (1), 159-192.

A.M. Turing, (1950). Computing machinery and intelligence. *Mind*, Vol.LIX, No.236.

B. van Fraasen, (1980). *The Scientific Image.* Oxford University Press, Oxford, UK.

K. Vicente and J. Wang, (1998). An ecological theory of expertise effects in memory recall. *Psychological Review*, 105 (1), 33-57.

L. Vygotsky, (1986). *Mind in Society.* Lawrence Erlbaum Associates, Hillsdale, NJ.

H. Wilson, F. Wilkinson, L. Lin, and M. Castilo, (2000). Perception of head orientation. *Vision Research*, 10, 459-472.

E. Zerubavel, (1997). *Social Mindscape: An Invitation to Cognitive Sociology.* Harvard University Press, Cambridge, MA.