Title

Toward A Multilevel Analysis of Human Attentional Networks

Permalink

https://escholarship.org/uc/item/9bn3p7kb

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 26(26)

ISSN

1069-7977

Authors

Wang, Hongbin
Fan, Jin
Yang, Yingrui

Publication Date

2004

Peer reviewed

Toward A Multilevel Analysis of Human Attentional Networks

Hongbin Wang (Hongbin.Wang@uth.tmc.edu)*
Jin Fan (Jin.Fan@mssm.edu)**
Yingrui Yang (yangyri@rpi.edu)***
*School of Health Information Sciences, University of Texas Health Science Center at Houston
**Department of Psychiatry, Mount Sinai School of Medicine
***Department of Cognitive Science, Rensselaer Polytechnic Institute

Abstract

Attention is a complex multilevel system subserved by at least three interacting attentional networks in the brain. This paper describes a multilevel computational model of attentional networks, developed in both the symbolic architecture of ACT-R and the connectionist framework of leabra. We evaluated the model using the Attentional Networks Test and the simulation results fitted the empirical data well. We argue that developing multilevel computational models helps to link findings at different levels.

Introduction

Suppose a student S was asked to solve the equation "2x + 3 = 9" (Figure 1A), and he used 2 seconds to produce the answer "x = 3". Both cognitive scientists X and Y were interested in understanding how S did it. Scientist X recorded S's detailed verbal protocol (Figure 1B), based on which, and other relevant behavioral measures, X hypothesized the possible knowledge structures underlying S's problem solving and developed a symbolic computational model that simulated the process (Figure 1C). On the other hand, scientist Y adopted sophisticated brain imaging techniques such as electroencephalograph (EEG) and functional Magnetic Resonance Imaging (fMRI) and acquired a high-resolution recording of S's brain dynamics during problem solving (Figure 1D). Based on some well-established neural computing principles, Y then developed a biologically realistic connectionist model to simulate the brain activities underlying S's performance (Figure 1E). Though both models fitted the data well, the two models are clearly different. While the symbolic model offers a description of the process with psychological plausibility and high behavioral relevance, the connectionist model emphasizes the process' biological realism and brain foundations. One question is, do we, cognitive scientists who endeavor to discover unified theories of cognition, have justifiable reasons to prefer one to another?

This question and similar others have led to a long debate in the rather brief history of cognitive science (e.g., Churchland & Sejnowski, 1992; Newell, 1990; Rumelhart & McClelland, 1986). Recently a BBS (Behavioral and Brain Sciences) target article was dedicated to this issue (Anderson & Lebiere, 2003). The authors adopted a set of 12 criteria, which they called "The Newell Test", to systematically compared and contrasted ACT-R, a rule-based cognitive architecture (Anderson & Lebiere, 1998),

and the connectionist modeling framework. Their conclusion was that both frameworks had great strengths as well as serious limitations as candidates of the unified theory of cognition.
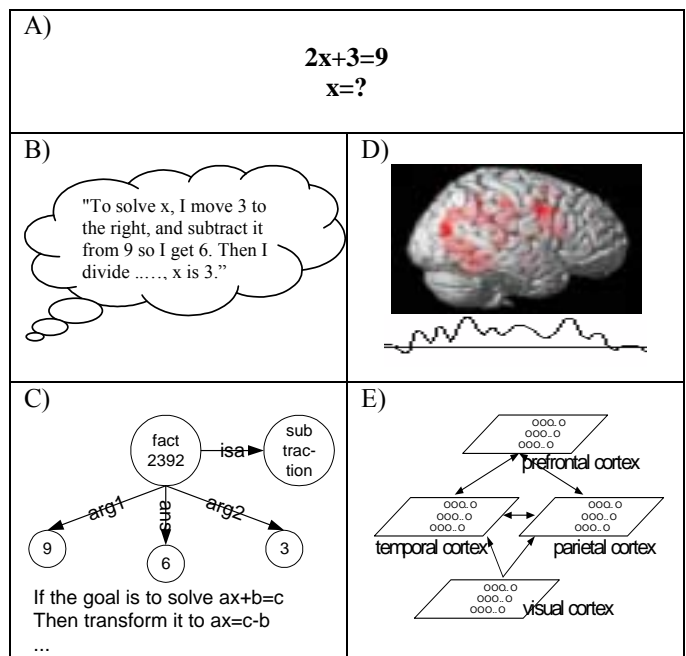
Figure 1. A hypothetical equation-solving problem is presented in A. Verbal protocol and brain imaging data are presented in B and D. Sketches of a symbolic model and a connectionist model of task are presented in C and E.

This is hardly surprising given the inherent complexity of the human mind itself. It has long been recognized that the mind is a multilevel construct and can be analyzed at different levels. Marr, for example, distinguished and separated among computational theory, representation and algorithm, and hardware implementation (Marr, 1982). Similar distinctions were made by Newell among different bands of cognitive functions (Newell, 1990). Newell argued that different bands utilize different basic operators, which have different time scales. More importantly, different bands form a hierarchy. Multiple lower lever basic operators can be combined to form higher level basic operators. In other words, lower level operators can be summarized up at higher level though this summarization may not be linear.

Single level analyses have been the dominant methodology in cognitive science. Experimental psychology and symbolic modeling, for example, largely depend on controlled experiments and behavioral observation. Recent advances in cognitive neuroscience allow us to directly observe, with high temporal-spatial resolutions, how an active brain functions during cognitive performance (Posner & Raichle, 1994). As a result, biologically realistic neural networks modeling has flourished (O'Reilly & Munakata, 2000). Efforts have also been made to probe the function of mind at lower molecular levels (e.g., Bellugi & George, 2001; Squire & Kandel, 2000). While all these levels of analyses tell us important aspects of the mind, neither of them alone is adequate to describe the whole picture. The human mind is a complex entity and may leave shadows at different levels when it works (Penrose, 1996). However, in order to achieve a unified theory all of the pieces have to be somehow linked together.

One approach would be to develop so called "hybrid systems", which typically combine symbolic and subsymbolic components together (e.g., Sun & Alexandre, 1997). We, for example, have developed a hybrid model of human abductive reasoning by combining a Soar component (a symbolic architecture) for hypothesis generation and a connectionist component for hypothesis evaluation (Johnson, Zhang, & Wang, 1997). Although hybrid systems take advantage of both types of components and can become quite powerful, they often bear little true psychological and neurophysiological significance due to the fact they are artificially assembled systems. While it is well agreed that human cognition involves mechanisms and operations at, among others, both psychological and neuronal networks levels, simply piecing them together is *ad hoc* and trivializes the problem (see also Wang, Johnson, & Zhang, 2003)

In this paper we argue that we need a multilevel modeling approach. That is, we need to develop well-fitted computational models at multiple levels for any given cognitive phenomenon. Because the mind manifests itself at multiple levels, each level is real and tells a unique story of the mind on its own. When we develop models for a specific phenomenon at multiple levels, we would be able to compare them, contrast them, and more importantly, mutually justify them. By doing so, we expect that a more complete picture of the mind might emerge.

This paper is organized as follows. We first briefly review findings on human attentional networks and introduce the Attentional Network Test (ANT) (Fan, MaCandliss, Sommer, Raz, & Posner, 2002). We then demonstrate the multilevel modeling approach by reporting and cross-validating two computational models for the same ANT task, one developed in ACT-R, and the other in leabra, a biologically realistic connectionist modeling framework (O'Reilly & Munakata, 2000). While both models fitted data well they emphasized different levels of explanations. Finally the implications of this practice are discussed.

Human Attentional Networks

Although "everyone knows what attention is" (James, 1890), how attention works remains one of the most challenging questions in science (Parasuraman, 2000; Pashler, 1998). Recent advances in cognitive psychology and cognitive neuroscience have suggested that there exist multiple attentional networks in the brain, each of which subserves different types of attention (Fan et al., 2002; Posner & Dehaene, 2000; Posner & Petersen, 1990). At least three attentional networks, for alerting, orienting, and executive control, have been distinguished at both cognitive and neuroanatomical levels (see Figure 2A). Specifically, alerting involves a change in the internal state to become ready for any incoming task-related events. Neuroimaging evidence has revealed that the alerting network consists of some frontal and parietal areas particularly of the right hemisphere. Orienting, closely related to the conventional selective visuo-spatial attention, involves selectively focusing on one or a few items out of many candidate inputs. Evidence has shown that the orienting network includes parts of the superior and inferior parietal lobe, frontal eye fields and such subcortical areas as the superior colliculus of the midbrain and the pulvinar and reticular nucleus of the thalamus. Finally, executive control of attention is related to monitoring and resolving conflicts. Executive control is often needed in higher level mental operations including planning, decision making, error detection, novel or not well-learned responses, and overcoming habitual actions. Converging evidence from neuroimaging and neuropathology studies has suggested that the executive control network consists of the midline frontal areas (anterior cingulate cortex), lateral prefrontal cortex, and the basal ganglia.

The ANT paradigm was recently developed to simultaneously measure the performance of the three attentional networks and evaluate their interrelationships (Fan et al., 2002). It is essentially a combination of a spatial cueing task (Posner, 1980) and a flanker task (Eriksen & Eriksen, 1974), as illustrated in Figure 2B. The stimulus consists of a row of 5 horizontal arrows and the participants' task is to report the pointing direction (left or right) of the center arrow (the target) by pressing a key. The four arrows surrounding the target, with two on each side, are called the flankers. These flanker arrows point either in the same direction as that of the target (the congruent condition), or in the opposite direction (the incongruent condition). An additional condition (the neutral condition) is also included in which the flankers are four straight lines with no arrowheads. To introduce an orienting component, the row can be presented at two locations, either above a fixation point (top) or below it (bottom). To introduce an alerting component, the row may be preceded by a cue (the cued condition) or may not (the no-cue condition). In addition, when there is a cue, this cue may be presented at the center fixation location (the center-cue condition), at the top or bottom location where the stimulus row is to appear (the

A)

B)
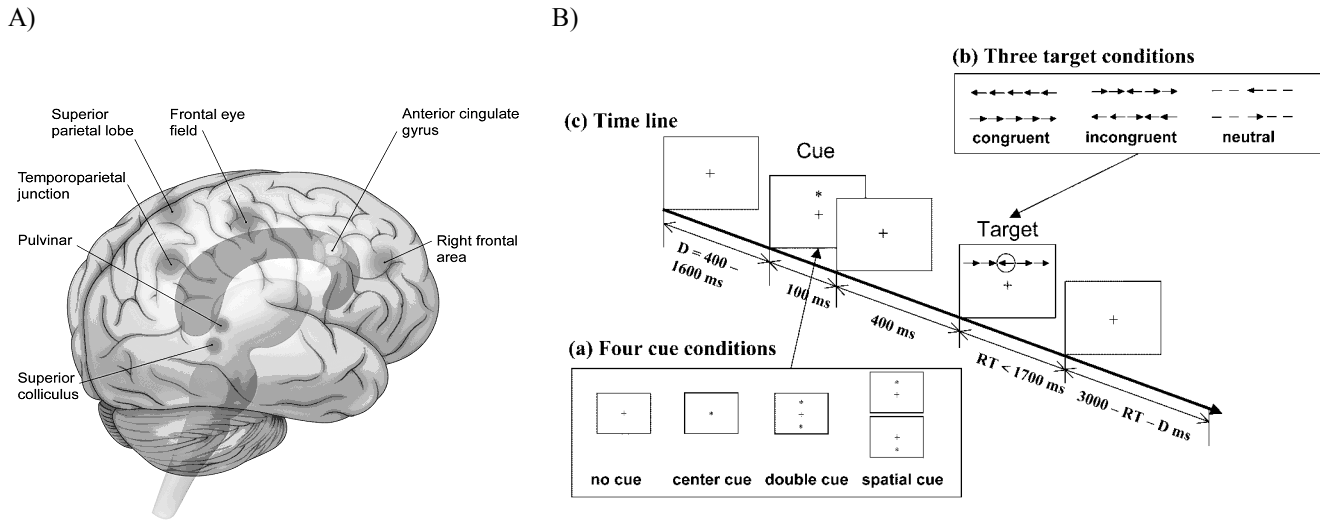
Figure 2. Human attentional networks (A) and the ANT task (B)

spatial-cue condition), or at both top and bottom locations (the double-cue condition). Note that while a spatial-cue precisely predicts where the stimulus is to appear, in both the center-cue condition and the double-cue condition the participant cannot infer that information from the cue.

Fan et al. (2002) tested 40 normal adult participants using the ANT paradigm. Their reaction time (RT) results are shown in Figure 3A. They then proposed the following formula as a measure of the efficiency of each of the three attentional networks:

· Alerting efficiency = RT(no-cue) – RT(double-cue),
· Orienting efficiency = RT(center-cue) – RT(spatial-cue),
· Conflict efficiency = RT(incongruent) – RT(congruent),

which resulted in the efficiency measures of 47 ± 18 ms, 51 ± 21 ms, 84 ± 25 ms, for alerting, orienting, and executive control, respectively.

Fan et al. (2001) also reported an fMRI study using the ANT paradigm. Their results were consistent with the general findings shown in Figure 2A.

Multilevel Computational Modeling of Human Attentional Networks

While both the behavioral and neuroimaging studies using the ANT paradigm revealed important psychological and neurophysiological characteristics of human attentional networks, there exists a gap between these two levels of analyses. In particular, how do these different attentional neural networks work together to generate psychologically meaningful behavior? It has been well agreed that the link between neural activities and psychological performance is nontrivial and must be taken into account seriously to avoid "neo-phrenology". Developing well-principled and constrained computational models help in the regard (Cohen & Tong, 2001).

Traditional computational modeling approaches to human attention have typically adopted various

connectionist modeling techniques (e.g., Cohen, Dunbar, & McClelland, 1990). While it has been fruitful, this practice fails to account for the manifestations of attention at symbolic/cognitive levels. As we illustrated earlier, attention, as an essential aspect of human cognition, is a complex multilevel construct. In order to understand the computational mechanisms of attention at different levels and the links among them, we need multilevel models.

We have developed a multilevel model for the ANT task. One sub-model was developed in the symbolic modeling framework of ACT-R and focused on the psychological aspects of the task. The other was developed in the connectionist modeling framework of leabra and emphasized the neurophysiological aspects of the task. A preliminary cross-validation of two models is discussed.

ANT on ACT-R

ACT-R is a production rule based cognitive modeling architecture developed by John Anderson and colleagues over a period of nearly two decades (see Anderson & Lebiere, 1998). In essence, ACT-R explains human cognition by proposing a model of the knowledge structures and knowledge deployment that underlie cognition. Although ACT-R consists of a nontrivial subsymbolic component for computations involving activation and association, it is fundamentally a symbolic modeling framework in that it relies extensively on various symbolic structures for knowledge representation. For example, ACT-R makes a fundamental distinction between declarative and procedural knowledge. Declarative knowledge corresponds to things people are aware of and can usually describe to others and is represented in ACT-R by chunks. Procedural knowledge is knowledge that people display in behavior but are not conscious of and is represented by production rules (condition-actions pairs). Both chunks and production rules are fundamental symbolic structures in ACT-R and are regarded as the atomic components of thought in the sense
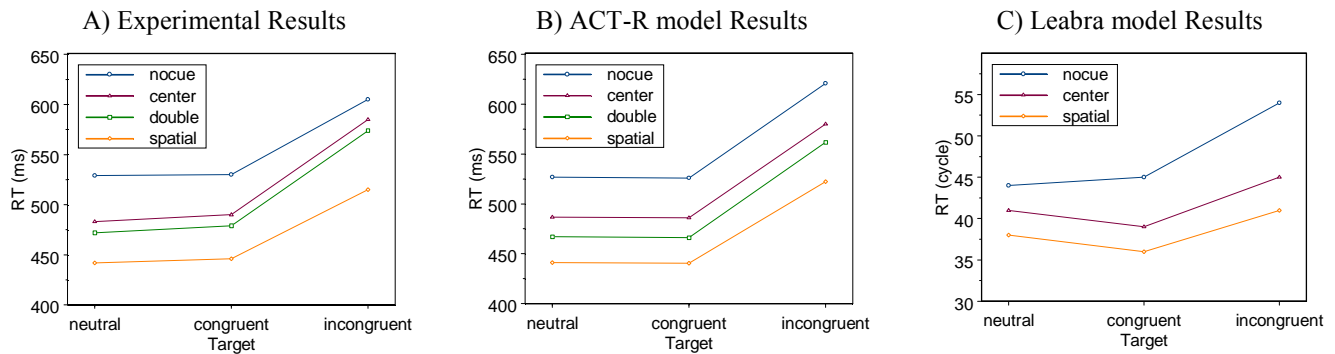
Figure 3. Experimental (A, based on Fan et al. (2002)) and modeling results (B and C).

that they are as far down as one can go in the symbolic decomposition of cognition. In ACT-R, on average every fifty (50) milliseconds, one production rule is chosen to fire, a few declarative chunks are processed, and cognition advances one step. Therefore, it is claimed that ACT-R captures the symbolic grain size of cognition.

We developed a computational model for the ANT task in the framework of ACT-R (Wang, Fan, & Johnson, 2004). Our purpose is two-fold. First, we want to explore how different types of attention work together in a single framework to produce the cognitive performance. Second, such a model offers a solid testbed for us to cross-validate those models based on various connectionist modeling results and neuroimaging data.

We started by analyzing the major functional components in the ANT task. We distinguished six major stages in a typical ANT trial: fixation and cue expectation; cue or stimulus judgment; cue processing; stimulus expectation; stimulus processing; and response. We then mapped these functional components onto 36 ACT-R production rules. With these rules our model could perform the ANT task and interact with the same experimental environment that human participants interact.

We evaluated the performance of the model by using the model as a "simulated subject" to perform the ANT experiment. The RT results of 100 "simulated subjects" are presented in Figure 3B. A correlation analysis shows very high correlations (0.99 for RTs and 0.97 for error rates) between the simulation and experimental results. We then followed the same procedure discussed early to estimate the effects of the three attentional networks based on the simulated RT data, resulting in the efficiency measures of 55 ± 7.4 ms, 45 ± 7.0 ms, 86 ± 7.4 ms, for alerting, orienting, and executive control, respectively. A close match between the two sets of data is apparent, with a notable exception that the simulated standard deviations are consistently smaller than the empirical ones. The reason is that we did not add any between-subject variance in our model. As a result, these simulated variances actually reflected those within-subject variations in performing the ANT task. Overall these results suggest that the model captured well the various attentional effects that the ANT task was designed to measure.

The concept of production rule is fundamental to our model of attention. One of the key features of the model is that it mapped the effects of attentional networks to production rules. Rules fire in sequence and operate at a rate of about 40-50 ms per production rule. As argued by ACT-R, production rules define the atomic components of thought at the symbolic level. When we examined the efficiency measures of attentional networks reported in Fan et al (2002) it seemed that they (51 ms, 47 ms, and 84 ms, for alerting, orienting, and executive control, respectively) fell well into the range of a few rule firings time period. Perhaps all we need is about one (for alerting and orienting) or two (for executive control) additional production rules to explain symbolically the work of attentional networks. This is indeed what our model demonstrated.

ANT on Leabra

Leabra (*l*ocal, *e*rror-driven and *a*ssociative, *b*iologically *r*ealistic *al*gorithm) is a connectionist modeling framework proposed recently by O'Reilly and Munakata (2000). There are at least three features that distinguish it from other connectionist modeling frameworks. First, it has sound neurological foundations. It is biologically realistic in multiple aspects. Its neurons compute based on membrane potentials and ion channels. Its neuronal connections are often bi-directional and cannot change signs (i.e., changing from an excitatory link to an inhibitory link, and vice versa). It uses biologically inspired learning rules such as Hebbian learning for unsupervised learning and the generalized recirculation algorithm (but not the biologically unrealistic backpropagation) for error-driven learning. Second, leabra is a coherently integrated framework. Many distinctions in traditional neural network modeling, including supervised vs unsupervised learning, feedforward vs recurrent networks, and pattern recognition vs self-organization maps, are all unified in a single coherent framework, based on well-supported biological principles. Third, partly due to its biological realism, it is now possible, for example, to designate a specific neural network to simulate a specific area of brain, and flexibly connect the multiple such networks, each of which can have its own properties such as the average activation level and the connection density, to simulate various brain pathways. As a result, it offers great

flexibility to build a hierarchy of neural networks and link network activities to higher-level symbols.

A connection model of the ANT task was developed in the framework of leabra. The structure of the model is shown in Figure 4. This model contains modules for all the three attentional networks. In addition, it contains modules for perception (visual input and primary visual cortex), object recognition (object pathway), and response (output). The networks are connected in such a way that they conform to the known functional an anatomical constraints as much as possible (Farah, 2000; O'Reilly & Munakata, 2000).

The model works as follows. When a cue comes on, the primary visual cortex module is activated, which in turn triggers the alerting network. This cue-induced alerting affects later stimulus processing because the alerting network will remain excited for a while which will activate the orienting network in general causing it to become ready for the incoming stimulus. In addition, when the cue is a spatial one (i.e., a cue that indicates where the target stimulus is to appear), it will further make the corresponding sub-region of the orienting network even more excited. This occurs because the orienting network adopts a retinotopy-based spatial representation of the environment. This extra excitation in the sub-region of the orienting network will facilitate the corresponding stimulus processing in the object pathway network, due to the connections between them. This accounts for the orienting effect. Finally, note that it is the object pathway network that is responsible for the arrow direction detection. When the incongruent stimulus (e.g., a left arrow flanked by four right arrows) is presented, the object pathway network may propose different responses, which compete for the final expression in the output network. The executive control network then activates making the center arrow defeating the flankers. This is where the executive control attention plays a role.
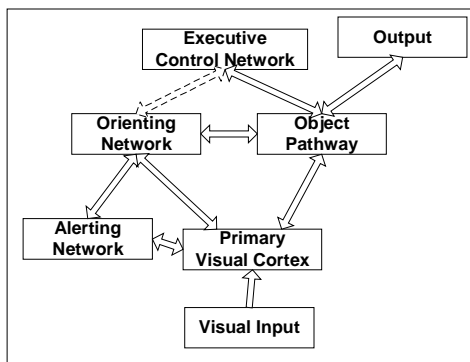
Figure 4. A leabra model of ANT.

The performance of the model was evaluated by using it to perform the ANT task. Stimuli are presented to the model in a similar way as to a human. Depending on the conditions, a cue, which can be either a center cue or a spatial cue, may be presented for a fixed time period before the stimulus presentation (note that the double cue condition was not simulated here since the current version of model

were not equipped with enough neurons). The number of cycles the output module takes to produce a stable response after the stimulus presentation serves as a measure of the reaction time. The simulation results are shown in Figure 3C. A regression analysis showed that

$$RT(ms) = 12 * RT(cycle)$$

with a R-square of 0.99. It is clear that the model fits the behavioral data reasonably well.

Discussion

Human attention is a multi-component multilevel construct. Both behavioral and neuroimging studies using the ANT paradigm revealed important aspects of the function of human attentional networks. Multilevel computational modeling helps to probe how these multiple components work together and manifest themselves at multiple levels.

The multilevel model we reported in this paper consisted of a sub-model developed in the framework of ACT-R and the other in the framework of leabra. While the former sub-model focused on the symbolic knowledge structure of cognitive performance and psychological plausibility, the latter focused on the subsymbolic neural information processing and biological realism. However, since both models simulated the same ANT task and fitted the empirical data well, the combined multilevel model offered a real possibility to cross-validate the models and probe the computational link among different levels.

First of all, the model illustrated interesting relationships between production rules and underlying neural computation. As demonstrated in the ACT-R model, rules are fundamental units of psychological reality and typically proceed serially. However, the underlying neural networks process information in parallel. The parallelism of neural computation and the serial nature of rule firing can be mapped against each other along the time line. Since both types of models decompose the cognitive performance into sub-units that occur at tens of millisecond scales, the mapping may be able to tell how rules are implemented in neural level computation. Based on the models, for example, we can map one ACT-R rule (40 ms in the current model) to about three leabra cycles (about 12 ms per cycle). Though such a simple and linear mapping should not be taken literally, it does provide a vivid footnote about how parallel neural computing is summarized psychologically by serial rule firings. It illustrates that we may not be able to find a "rule center" in the brain. Instead, rules can be implemented anywhere in the brain – they are simply pattern matching. For example, there is a symbolic rule that summarizes the conflict monitoring and detection operation typically subserved by the anterior cingulate cortex. The general neural priming underlying alerting in the alerting networks is summarized by another task switch rule.

Our model also demonstrates how functionally identical operations can be implemented by different mechanisms at different levels. One interesting finding from Fan et al. (2002) is the small but reliable difference in RT (about 11 ms) between the center-cue and the double-cue conditions.

A convenient explanation is that in the double-cue condition due to *diffused attention* both stimulus locations had been primed a little, which saved a little time when the stimulus appeared later. While it is easy to model priming and diffused attention in a connectionist model (e.g., our leabra model), how it is implemented at a symbolic rule level raises a challenge. Our ACT-R model adopted a mechanism in which several symbolic and psychologically meaningful move-attention operations were carried out sequentially. The simulated RT difference was 19 ± 8 ms.

The multilevel model for human attentional networks we reported in this paper has allowed us to compare/contrast the computational mechanisms at different levels and to probe the important computational links between psychologically meaningful mental operations and neural activities. It also enjoys potentially significant prediction power in that the model at one level can lead to nontrivial predictions about the operations at another level. However, we recognize that for this approach to work models at each level have to be independently and/or mutually validated. Further analyses and more detailed alignments of our current model remain to be done.

Acknowledgments

References

Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought*. Hillsdale, NJ: Lawrence Erlbaum Press.

Anderson, J. R., & Lebiere, C. (2003). The Newell Test for a theory of Mind. *Behavioral and Brain Science, 26*, 587-639.

Bellugi, U., & George, M. S. (Eds.). (2001). *Journey from cognition to brain to gene: Perspectives from Williams Syndrome*. Cambridge, MA: MIT Press.

Churchland, P. S., & Sejnowski, T. J. (1992). *The computational brain*. Cambridge, MA: MIT Press.

Cohen, J. D., Dunbar, K., & McClelland, J. L. (1990). On the control of automatic processes: A parallel distributed processing account of the stroop effect. *Psychological Review, 97*(3), 332-361.

Cohen, J. D., & Tong, F. (2001). The face of controversy. *Science, 293*, 2405-2407.

Fan, J., MaCandliss, B. D., Sommer, T., Raz, A., & Posner, M. I. (2002). Testing the efficiency and independence of attentional networks. *Journal of Cognitive Neuroscience, 14*(3), 340-347.

Fan, J., McCandliss, B. D., Flombaum, J. I., & Posner, M. I. (2001). *Imaging attentional networks*. Paper presented at the Society for Neuroscience 2001 Annual Meeting, San Diego, CA.

Farah, M. J. (2000). *The cognitive neuroscience of vision*. Malden, MA: Blackwell Publishers.

James, W. (1890). *Principles of psychology*. New York: Holt.

Johnson, T. R., Zhang, J., & Wang, H. (1997). A hybrid learning model of abductive reasoning. In R. Sun & F. Alexandre (Eds.), *Connectionist Symbolic Integration*. Hillsdale, NJ: Lawrence Erlbaum.

Marr, D. (1982). *Vision*. San Francisco, CA: Freeman.

Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.

O'Reilly, R. C., & Munakata, Y. (2000). *Computational explorations in cognitive neuroscience*. Cambridge, MA: MIT Press.

Parasuraman, R. (Ed.). (2000). *The attentive brain*. Cambridge, MA: MIT Press.

Pashler, H. E. (1998). *The psychology of attention*. Cambridge, MA: MIT Press.

Penrose, R. (1996). *Shadows of the mind: A search for the missing science of consciousness*. New York: Oxford University Press.

Posner, M. I., & Dehaene, S. (2000). Attentional networks. In M. S. Gazzaniga (Ed.), *Cognitive neuroscience: A reader*. Malden, MA: Blackwell Publishers.

Posner, M. I., & Petersen, S. E. (1990). The attention systems of the human brain. *Annual Review of Neuroscience, 13*, 25-42.

Posner, M. I., & Raichle, M. E. (1994). *Images of mind*. New York: Scientific American Library.

Rumelhart, D. E., & McClelland, J. L. (1986). *Parallel distributed processing: Explorations in the microstructure of cognition. Vol 1: Foundations*. Cambridge, MA: MIT Press.

Squire, L. R., & Kandel, E. R. (2000). *Memory: From mind to molecules*. New York: Scientific American Library.

Sun, R., & Alexandre, F. (Eds.). (1997). *Connectionist-symbolic interaction: From unified to hybrid approaches*. Mahwah, NJ: Lawrence Erlbaum Associates.

Wang, H., Fan, J., & Johnson, T. R. (2004). A Symbolic Model of Human Attentional Networks. *Cognitive Systems Research, 5,* 119-134.

Wang, H., Johnson, T. R., & Zhang, J. (2003). [Commentary] A multilevel approach to modeling human cognition. *Behaviroal and Brain Science, 26,* 626-627.