

A symbolic model of human attentional networks

Action editor: Ron Sun

Hongbin Wang^{a,*}, Jin Fan^b, Todd R. Johnson^a

^a *School of Health Information Sciences, University of Texas Health Science Center at Houston,
7000 Fannin, Suite 600 Houston, TX 77030, USA*

^b *Sackler Institute for Developmental Psychobiology, Weill Medical College of Cornell University,
1300 York Ave., Box 140, New York, NY 10021, USA*

Received 15 August 2003; accepted 8 December 2003

Abstract

An increasing body of evidence has shown that attention is a multi-type and multilevel cognitive faculty. The dominant computational modeling approaches to attention have often focused on one specific type of attention at one specific level. In particular, various connectionist modeling techniques at the subsymbolic level have been widely adopted. In this paper, we report a symbolic computational model of the Attentional Network Test, which simultaneously involves different types of attention (alerting, orienting, and executive control), each subserved by distinctive attentional networks in the brain. The model was developed in ACT-R, a rule-based cognitive architecture. The results show that the model, by sequentially firing rules at a rate of about one every 40 ms, was able to capture the effect of each attentional network. The model implies that while the attentional networks can be distinguished at both neuroanatomical and behavioral levels, different attentional networks may adopt similar computational operations at least at a symbolic rule level.

© 2004 Elsevier B.V. All rights reserved.

1. Introduction

Although “everybody knows what attention is” (James, 1890), the nature of attention remains elusive after more than one hundred years of active research and hundreds of publications in this subject. This is hardly surprising given the com-

plexity of the phenomenon itself. A large body of evidence, in the broad areas of cognitive psychology and cognitive neuroscience, has demonstrated that attention is not a unitary but a diversified faculty of the human cognitive system (Parasuraman & Davies, 1984; Posner & Raichle, 1994) – it is subserved by multiple interrelated attentional networks in the brain and manifests itself in a variety of types and at different levels in almost every aspect of human behavior, from perception and motor control to working memory, skill acquisition, response selection, and consciousness

* Corresponding author. Tel.: +1-713-500-3911; fax: +1-713-500-3929.

E-mail address: Hongbin.Wang@uth.tmc.edu (H. Wang).

(e.g., Luck, Woodman, & Vogel, 2000; Pashler, 1998; Posner, DiGirolamo, & Fernandez-Duque, 1997).

Partly because of this complexity, traditional computational modeling approaches to human attention have typically focused on one specific type of attention at one specific level. In particular, various connectionist modeling techniques at the subsymbolic level have been widely adopted (Churchland & Sejnowski, 1992; Rumelhart & McClelland, 1986). Examples include Cohen and colleagues' well-known connectionist model of executive control of attention for the Stroop task (Cohen, Dunbar, & McClelland, 1990) and Mozer's connectionist modeling work of selective spatial attention in object recognition (Mozer, 1991). For more recent development and reviews in these lines of research, see (Botvinick, Braver, Barch, Carter, & Cohen, 2001; Mozer & Sitton, 1998; O'Reilly & Munakata, 2000).

While it has been very fruitful, this practice is problematic mainly due to two reasons. First, by focusing on only a few specific types of attention, it fails to adequately appreciate the possible interactions among different varieties of attention and different attentional networks. Second by emphasizing only the subsymbolic level of analysis, it encounters difficulties in understanding the link between neural information processing and the operation of attention at the symbolic/cognitive level, which is psychologically as real as the underlying neurophysiological mechanisms of attention (e.g., Newell, 1990).

In this article, we attempt to address these issues by reporting a symbolic computational model of attention that simultaneously simulated the essential operations of multiple attentional networks. The model was developed in ACT-R, a production rule based hybrid cognitive architecture (Anderson, 1993; Anderson & Lebiere, 1998). We believe that such a model offers an opportunity for us to explore the psychologically plausible symbolic foundations of attention and, more importantly, to compare/contrast it with neurally-based connectionist models of attention.

It should be noted that developing symbolic models of attention is not new. Byrne and Anderson (1998) augmented ACT-R with a set of

perceptual-motor modules, which implemented some rudimentary visuospatial attention functions. Altmann and Davidson (2001) and Lovett (2002) have developed ACT-R models of the Stroop task, in which executive control of attention plays a major role. Our model differs from these models in two aspects. First, our model simultaneously simulates multiple attentional networks, thus capturing the operational features of different types of attention, including alerting, orienting, and executive control. Second, both Altmann and Davidson's and Lovett's models relied extensively on the subsymbolic mechanisms of ACT-R to implement parallel processing and executive control. In our model, symbolic operations at the production rule level are emphasized.

This article is structured as follows. We first briefly introduce the attentional networks account of human attention and an experimental paradigm that was designed to simultaneously measure the performance of multiple attentional networks. We then briefly introduce ACT-R. Finally, an ACT-R model of human attentional networks is reported and discussed.

2. Human attentional networks

Recent advances in cognitive psychology and cognitive neuroscience have suggested that there exist multiple attentional networks in the brain, each of which subserves a different type of attention (Fan, Raz, & Posner, 2003c; Posner & Dehaene, 2000; Posner & Petersen, 1990; Posner & Raichle, 1994). At least three attentional networks, for *alerting*, *orienting*, and *executive control*, have been distinguished at both cognitive and neuro-anatomical levels. Specifically, *alerting* involves a change in the internal state to become ready for any incoming task-related events. Alerting is an important source of attention in the sense that maintaining an adequate level of alertness is critical for optimal performance. Neuroimaging evidence has revealed that the alerting network consists of some frontal and parietal areas particularly of the right hemisphere. Lesions of these areas reduce alertness. *Orienting*, closely related to the conventional selective visuospatial attention,

involves selectively focusing on one or a few items out of many candidate inputs. Apparently, orienting can be voluntary (top-down and controlled) or involuntary (bottom-up and automatic), overt (with head/eye movement) or covert (without head/eye movement), location-based (orienting to spatial locations) or object-based (orienting to objects). 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 in the presence of competing information. 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 neuropathological studies has suggested that the executive control network consists of the midline frontal areas (anterior cingulate cortex) and lateral prefrontal cortex. See Posner and Dehaene (2000) for a more detailed description of these attentional networks.

An experimental paradigm called the Attentional Network Test (ANT) was recently devel-

oped to simultaneously measure the performance of the three attentional networks and evaluate their interrelationships (Fan, MaCandliss, Sommer, Raz, & Posner, 2002). It is essentially a combination of a spatial cueing task (Posner, 1980) and a flanker task (Eriksen & Eriksen, 1974), as illustrated in Fig. 1. In each ANT trial, the participants look at a fixation point in the center of a computer screen and wait for the stimulus to appear. The *stimulus* consists of a row of five horizontal arrows and the participants' task is to report the direction (left or right) that the center arrow (the *target*) points to by pressing a corresponding key as quickly as possible. The reaction time (RT) is recorded. 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. One manipulation is that the stimulus row is not presented at the fixation location. Instead, it can be presented at two locations, either above a fixation point (top) or below it (bottom). Therefore, to identify the direction of the target,




Fig. 1. A sketch of the design of ANT.

presumably the participants have to shift their attention either upward or downward to the stimulus row. Yet another manipulation is the cue condition. That is, the stimulus row may be preceded by a cue (the cued condition) or may not (the no-cue condition). In particular, 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 spatial-cue condition), or at both top and bottom locations (the double-cue condition). Therefore, 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.

The ANT paradigm adopts the following formula to measure the efficiency of each of the three attentional networks:

Alerting efficiency

$$= \text{RT}(\text{no-cue}) - \text{RT}(\text{double-cue}),$$

Orienting efficiency

$$= \text{RT}(\text{center-cue}) - \text{RT}(\text{spatial-cue}),$$

Executive control efficiency

$$= \text{RT}(\text{incongruent}) - \text{RT}(\text{congruent}).$$

The underlying rationale is as follows. First, since the appearance of a cue in the double-cue condition precisely alerts the participants that the stimulus row is going to occur after 500 ms and there is no such alerting in the no-cue condition, the RT difference between the no-cue condition and the double-cue condition can be used as a measure of the alerting efficiency.¹ Second, in both the center-cue condition and the spatial-cue condition the participants are alerted when the stimulus row is going to occur. The difference is where it is going to occur: while the spatial-cue precisely tells the participants where the stimulus row is going to appear,

the center-cue does not. Presumably, knowing the stimulus location before the stimulus appears allows the participants to be better prepared for response in that they can shift attention in advance to the desired location. As a result, the RT difference between the center-cue condition and spatial-cue condition serves a measure of the orienting efficiency. Finally, compared to the congruent condition, the presence of incongruent flankers in the incongruent condition results in interference in response. Evidence has shown that this type of interference shares the largely identical brain foundations with that in the Stroop effect and requires similar executive control of attention to resolve (see Fan, Flombaum, McCandliss, Thomas, & Posner, 2003a). Therefore, the executive control cost is measured as the RT difference between the incongruent and congruent conditions.

Fan et al. (2002) tested 40 normal adult participants using the ANT paradigm. They adopted a balanced within-subject factorial design and presented 3 blocks of trials to each participant with 96 trials in each block (3 flanker conditions $\times$ 4 cue conditions $\times$ 2 possible stimulus locations $\times$ 2 possible target arrow directions $\times$ 2 repetitions). Their reaction time results led to the efficiency measures of 47 ± 18 ms, 51 ± 21 ms, and 84 ± 25 ms, for alerting, orienting, and executive control, respectively. A further correlation test showed that the effects of the three attentional networks were statistically independent.

ANT as a general and robust framework for measuring attentional efficiency has been evaluated in a variety of studies. For example, Rueda et al. (in press) adopted the ANT paradigm to study human attentional development. They measured the attentional efficiency of children in different age groups (from 6 to 10 years old) and compared the results with adults. They found that each network showed a different developmental course. ANT has also been used in various patient populations, including attention deficit hyperactivity disorder (ADHD, Booth, Carlson, & Tucker, 2001), and borderline personality disorders (Posner et al., 2002). More recently, the attentional efficiency measures from ANT have been linked to neuroimaging results and genetic variations (Fan, Fossella, Sommer, Wu, & Posner, 2003b).

¹ Note that both the double-cue and center-cue conditions alert the participants that the stimulus row is going to occur but do not tell where. The reason that the double-cue condition instead of the center-cue condition is used to define the alerting efficiency is that the double-cue condition involves diffused attention, which is more similar to the situation in the no-cue condition. Similar reasoning underlies the definition of the orienting effect: both the center-cue and spatial-cue conditions involve non-diffused cued attention.

3. 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 (Anderson, 1983, 1990, 1993; Anderson & Bower, 1973; 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 relies extensively on various symbolic structures for knowledge representation. For example, ACT-R makes a fundamental distinction between declarative and procedural knowledge (see also Schacter & Tulving, 1994). Declarative knowledge corresponds to things people are aware of and can usually describe to others. Procedural knowledge is knowledge that people display in behavior but are not conscious of. Declarative knowledge in ACT-R is represented in terms of chunks. Procedural knowledge is represented in terms of production rules, which are condition-action 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 that they are as far down as one can go in the symbolic decomposition of cognition. In ACT-R, on average every 50 ms, 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.

Canonical ACT-R is mainly a theory of higher-level cognition, excelling in modeling phenomena in human memory, learning, and problem solving. One recent advance in ACT-R is the release of ACT-R 5.0, which intends to close the gap between ACT-R and the external environment by augmenting ACT-R with a set of perceptual-motor modules (Byrne & Anderson, 1998). To implement the visual system, ACT-R 5.0 implements some rudimentary attention functions. For example, there is a *move-attention* command that higher cognition can send to the visual system. This command allows the visual system to shift atten-

tion to a designated location in the environment and then start visual recognition. The time cost of this command is specified by a parameter, which is typically set to be 85 ms, a reasonable setting based on many previous studies (see Egeth & Yantis, 1997).

Although ACT-R 5.0 represents important progress in the symbolic modeling of attention, one limitation is that it only handles the spatial orienting type of attention. It is not clear how ACT-R 5.0 can be used to model other types of attention such as alerting and executive control at the symbolic rule level.

4. ANT on ACT-R

4.1. The model

We developed a computational model for the ANT task in the framework of ACT-R 5.0. Our purpose is two-fold. First, we want to explore how alerting and executive control of attention can be modeled symbolically in ACT-R 5.0 and how the three types of attention work together to produce the cognitive performance. In particular, we hope to show how the effects of different attentional networks can be explained by just firing different sets of production rules but without resorting to various subsymbolic mechanisms. Second, such a model offers a possibility for us to cross-validate those models based on various connectionist modeling results and neuroimaging data, and by doing so we hope that we can probe the possible connections about the function of attention among different levels of description.

We started by analyzing the major functional components in the ANT task. Based on the design presented in Fig. 1, we distinguished 6 stages in performing a generic ANT trial (see Fig. 2):

1. *Fixation and cue expectation*. In this stage, the participant fixates at the fixation point and prepares for something to occur. Note that at this time the expectation of the participant was uncertain – either a cue (in the cued conditions) or a stimulus row (in the no-cue condition) could appear. As a result, we hypothesize that the participant has to engage in an uncertain




Fig. 2. A functional decomposition of the ANT task. The number in parentheses in each stage represents the number of production rules that are used to implement these functions.

preparatory state that is appropriate for either possibility.

2. *Cue or stimulus*. This second stage starts when a visual object appears, which can be either a cue (in the cued conditions) or a stimulus row (in the no-cue condition). A decision has to be made based on what appears. If it is a cue, go to stage 3 for cue processing. If it is a stimulus row, go directly to stage 5 for stimulus processing.
3. *Cue processing*. In this stage the cue is processed and attention is shifted to the corre-

sponding cued location. More specifically, for the center-cue, attention remains at the center fixation location. For the spatial-cue, attention is moved to the cued location (top or bottom) where a stimulus row is to appear. For the double-cue condition, instead of assuming that attention is diffused to both the cued locations, we hypothesize that attention shifts randomly to one of the cued locations. We will discuss this issue in more detail in the model discussion section.

4. *Stimulus expectation.* In this stage the participant is expecting a stimulus to appear, with attention remaining at the cued location(s). Note that in this stage the participant has engaged in a preparatory state that is particularly suitable for stimulus processing.
5. *Stimulus processing.* In this stage the stimulus has to be processed and the direction of the target arrow identified. Specifically, in the spatial-cue condition, since attention has already been allocated to the location where the stimulus appears, presumably a more accurate and rapid stimulus processing can occur. In the center-cue condition, an additional attentional movement needs to be initiated to shift attention to the stimulus location. Depending upon the result from stage 3, similar operations need to be performed in the double-cue condition. Again, we leave the details of this condition to the model discussion section.
6. *Response.* In this final stage, the participant responds by initiating a key pressing process based on the result from the previous stage.

We implemented our model by mapping these functional components into ACT-R production rules. We used 36 rules to fully implement these functions and cover all the conditions. The distribution of the rules in each stage is also shown in Fig. 2. The name and a brief description of each rule are provided in Appendix A.

With these rules our model could perform the ANT task and interact with the same experimental environment that human participants use. In a typical simulation trial, visual objects, including the fixation, the cue(s), and the stimulus, are presented in a computer window in a timed sequence according to the design as shown in Fig. 1. The model can “see” what is currently in the window through its visual system. At any time point, based on the goal and what the model currently “sees”, one of the best matching rules is selected to fire, which may lead to a chain of more rule firings. In the current model each rule firing is set to cost 40 ms instead of the default 50 ms (see the model discussion section for more details). The trial ends when the model makes a response by “pressing” a response key. The time from the stimulus pre-

sensation to the key-press is then recorded as the reaction time.

As an example, a running trace generated when the model performs a center-cue incongruent ANT trial is listed in Appendix B. It shows that the trial starts at time 0 and the fixation lasts for 1255 ms, during which two production rules have fired. A center cue appears at time 1255 and disappears at time 1355 (for a duration of 100 ms). Two production rules have fired during this cue processing period. At time 1755 (after a 400 ms wait) the stimulus appears. It then takes eight production rules firing in sequence for the model to make a response and initiate a key-press at time 2165. The key-press is not finished until time 2375, resulting in a reaction time of 620 ms ($2375 - 1755$).²

The ACT-R code of the model can be downloaded online from <http://www.sahs.uth.tmc.edu/hwang/antmodeling.htm>.

4.2. Model evaluation

We evaluated the performance of the model by using the model as a simulated human participant to perform the ANT experiment. The experimental design was exactly same as that used by Fan et al. (2002), with 3 blocks of 96 randomly mixed trials. The RT results of testing 100 “simulated subjects” are summarized in Table 1 and Fig. 3, along with the experimental results from Fan et al. (2002). A correlation analysis shows very high correlations (0.99 for RTs and 0.97 for error rates) between the simulation and experimental results. From the results in Table 1, we estimated the effects of the three attentional networks, which are shown in Table 2, together with Fan et al. (2002)’s results. It shows a close match between the two sets of data, 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

² Note that the RT of 620 ms shown in this trace is different from the RT of the same condition in Table 1, which is 580 ms. The reason is that 580 ms is an average result of many trials. In some trials, the RT may be lower depending on the outcome of conflict resolution.

Table 1
Means RT and error rates under each condition

Congruency	Cue type			
	No-cue	Center-cue	Double-cue	Spatial-cue
(a) Mean RTs $\pm$ Standard deviations in ms from the experiment and (the model simulation)				
Neutral	529 $\pm$ 47 (527 $\pm$ 3)	483 $\pm$ 46 (487 $\pm$ 3)	472 $\pm$ 44 (467 $\pm$ 5)	442 $\pm$ 39 (441 $\pm$ 4)
Congruent	530 $\pm$ 49 (526 $\pm$ 4)	490 $\pm$ 48 (486 $\pm$ 3)	479 $\pm$ 45 (466 $\pm$ 6)	446 $\pm$ 41 (441 $\pm$ 4)
Incongruent	605 $\pm$ 59 (621 $\pm$ 14)	585 $\pm$ 57 (580 $\pm$ 14)	574 $\pm$ 57 (562 $\pm$ 15)	515 $\pm$ 58 (522 $\pm$ 16)
(b) Error rates in % from the experiment and (the model simulation)				
Neutral	1.17 (0.96)	0.93 (0.92)	1.56 (0.71)	0.78 (0.79)
Congruent	0.73 (0.75)	0.54 (1.00)	0.59 (0.79)	0.44 (0.83)
Incongruent	3.49 (3.25)	4.88 (3.79)	4.27 (3.50)	3.51 (2.67)




Fig. 3. Mean RTs in various conditions from (a) the experimental study and (b) our model simulation.

Table 2
Effects of attentional networks

Effects (ms) (mean $\pm$ SD)	Attentional networks		
	Alerting	Orienting	Executive control
Experiment	47 $\pm$ 18	51 $\pm$ 21	84 $\pm$ 25
Simulation	55 $\pm$ 7.4	45 $\pm$ 7.0	86 $\pm$ 7.4

result, these simulated variances actually reflected those within-subject variations in performing the ANT task. Although we can add some between-subject noise to get a better fit of the variances, it leads to a more complicated model.

One of the key results in Fan et al. (2002) concerns the interrelationship among the three attentional networks. They found that the efficiency measures of the three attentional networks were independent of each other based upon the statistics of correlations. Similar independences were found in the simulation results, which are shown in Table 3, along with empirical data from Fan et al. (2002). It is clear that the two sets of data reveal similar correlation structures among the measures. An examination of the model suggests that the independence resulted from the dif-

Table 3
Correlations between attentional networks^a

Results	Alerting	Orienting	Executive control
Orienting			
Experiment	0.08		
Simulation	0.04		
Executive control			
Experiment	0.05	-0.12	
Simulation	-0.07	-0.14	
Mean RT			
Experiment	0.09	0.29	0.44 ^b
Simulation	-0.14	-0.15	0.80 ^b

^a Correlations are calculated based on the relevant measures of all participants (or simulated participants).

^b Correlation is significant at the 0.01 level.

ferent sets of rules underlying different testing conditions. The significant positive correlation between the executive control effect and the overall mean RT, shown in both the experiment and the simulation, is interesting. It probably reflects the strong influence of the RTs in the incongruent condition on the overall mean RT. Overall these results suggest that the model captured well the various attentional effects that the ANT task was designed to measure.

4.3. Model discussion

What are the mechanisms underlying the model's attention-related performance? Here we briefly report several major features of the model that we think are critically responsible for the performance.

First, to perform any one ANT trial, in stage 5 (stimulus processing), one *key-press* process (for responding) and at least one *move-attention* process (for encoding the target) are necessary. Based on well-established evidence ACT-R presets the time cost of these processes to be 210 and 85 ms, respectively. We did not change these settings. Subtracting these costs from the RT measures gives us a rough range of time that the production rules have to explain. Note that although the model consists of 36 production rules, in any one trial not all of them will fire (or are eligible to fire) because they are programmed to cover all different conditions. When multiple rules are applicable in a situation, our model relies on ACT-R's built-in conflict resolution mechanisms to choose one to fire.

Taken together, these settings allow us to predict how many production rules are recruited in each trial. More specifically, the model uses, depending upon the conditions, about 4–8 production rules to perform a trial, for total costs of about 160–320 ms. The trace listed in Appendix B demonstrates this. It shows that after the stimulus appears, eight production rules fire in sequence (*notice-stimulus-with-centercue-and-shift*, *attend-to-at-large-target*, *attended-item-is-left-to-the-target*, *refocus-again-if-incongruent*, *harvest-target*, *goahead-responding-if-it-is-the-target*, *decide-left*, *respond*), resulting in two *move-attention* and one *key-press* commands. The final reaction time is 620 ms.

Second, although the experimental results in Fan et al. (2002) show a performance with very high accuracy, participants made errors, in all conditions. How could the model make errors? We speculated that error could come in two major sources: random noise in the response and the distraction caused by the incongruent flankers. While the first source might be responsible for the baseline (neutral condition) error rate, which is about 1.11%, the second source might account for the higher error rate in the incongruent conditions, which is about 4.04%. We modeled the first source of error through two production rules (*random-left* and *random-right*), which say to respond left (or right) regardless of the direction of the target arrow. They compete with the *decide-left* or *decide-right* rule to determine the response in each trial, with their chances of winning being very low (the odds are set to be 1:20). The second source is modeled through a production called *hurryup-responding-no-matter-whether-target-or-not*. This rule rushes to respond as soon as an arrow, which may not be the center target arrow, is encoded. This rule competes, again with relatively low winning probability, with more cautious rules that emphasize accuracy (e.g., checking and refocusing) and causes errors if it wins in the incongruent condition (the odds are 1:3). As shown in Table 1, the model error rates are about 0.85 and 3.30, for the baseline average and the incongruent condition respectively, indicating a good match to the experimental data.

Third, how has the alerting effect been explained in the model? The empirical measure of 47 ms is about the time cost of one ACT-R production rule. This is indeed one major mechanism underlying how our model generated the alerting effect. There is one critical production rule that can fire in the no-cue condition but not in any other cued condition. This rule is called *not-cue-so-switch-state-and-shift-attention* (rule 4 in Appendix A). As we mentioned earlier, since either a cue or a stimulus row could occur following stage 1, the participant had to engage in an uncertain preparatory state. This was different from the specific and certain preparatory state (for stimulus-processing) following stage 4 in the cued conditions. This rule summarizes the cost associated with the

state change (from expecting either a cue or a stimulus to specifically expecting a stimulus) and is responsible for a *major* part of the alerting effect.³ It is important to note that this alerting effect did not come from the time saving due to an anticipatory shift of attention but from an expectancy state switching. We believe that this account for the alerting effect captures the essential characteristics of alerting as a change of preparatory state (from a uncertain state to a better prepared state), which may often facilitate later processes.

Fourth, the orienting effect reflects the difference between the center-cue condition and the spatial-cue condition, with an empirical measure of about 51 ms. Presumably this is due to the time saved from knowing in advance where the stimulus is going to appear. In our model this difference is, again, modeled by a production rule named *notice-stimulus-with-centercue-and-shift* (rule 17 in Appendix A), which basically says that if a stimulus appears while attention is on the center location, try to shift attention to the stimulus. We assume that in the spatial-cue condition attention has already been allocated to the correct spatial location before the stimulus appears, whereas in the center-cue condition the firing of this additional production rule is necessary to bring the system to a comparable level of stimulus processing. This additional step costs 40 ms and is the *major* source of the orienting effect.

Fifth, the effect of executive control reflects the effect of the flankers on the task performance. The empirical measure is about 84 ms. In our model this effect is modeled using two mechanisms. First, the result of *move-attention* is not perfect. When attention is directed to one location, an object nearby may be selected, especially when the scene is crowded or the objects are similar (Pashler, 1998). This kind of imprecision is one fundamental reason

for the flanker effect. Second, there are three production rules that specifically distinguish between the incongruent and congruent condition in the cases when a distracter arrow, but not the center target, is focused on. When this situation arises, we hypothesize that instead of performing an automatic but expensive (costly) re-focusing, a simpler congruency detection is conducted. Specifically, if a congruent condition is detected, the system moves on to respond since a re-focusing is not necessary. This is implemented by the rule named *goahead-responding-if-congruent* (rule 28 in Appendix A). If an incongruent condition is detected, the system can either move on to do a direct re-encoding (*harvest-target-directly-if-incongruent*, rule 30 in Appendix A, equivalent to a simple direction reversal) or perform a re-focusing again (*refocus-again-if-incongruent*, rule 29 in Appendix A). Compared to the congruent condition, the first strategy costs an additional 85 ms (a direct *move-attention*) and the second strategy costs an additional 125 ms (one additional rule firing for refocusing first and then a *move-attention*). With the second strategy a little more likely to be selected than the first one (the odds are set to be about 2:1), this leads to an additional cost of about 111 ms on average for the incongruent conditions. It is this cost, accrued quite often since there are four times more flankers than the target, that is *mainly* responsible for the executive control effect.

One interesting empirical result from Fan et al. (2002)'s study is the small but reliable difference (11 ms) between the center and double-cue conditions. This can be conveniently explained by a diffused attention assumption. More specifically, the center-cue induces the participant to focus attention on the fixation location while the double-cue makes the participant diffuse attention at both the top and bottom locations. As a result, compared to the center-cue condition where the stimulus location receives no priming, the stimulus location in the double-cue condition is primed a small amount, which may speed up later stimulus processing. This explanation, however, provides a challenge to a symbolic model of attention. How can attention be diffused symbolically when we only have in hand a *move-attention* command, which presumably shifts the *focus* of attention to a

³ Note that the cost of the production rule (40 ms) does not match the simulated alerting effect (55 ms, see Table 2). The reason is that the simulated effect is an average number that incorporates multiple sources although the cost of the production rules accounts for the major part of the effect. Similar mechanisms apply to the orienting and executive control effects discussed next.

pre-specified spatial location? While a parallel global spatial priming process is plausible from a neural computing perspective, we argue that it is equally plausible to speculate a series of symbolic *move-attention* operations occurring quickly and sequentially. This speculation also captures *psychologically* how we focus our spatial attention in our everyday life (*keep your attention here!*). Our model uses this mechanism to account for what occurs in the double-cue condition. Instead of using a diffused attention mechanism, we assume that attention is moved twice, each time to one of the two cues. The final location is where attention resides to wait for the stimulus to appear. This is like a betting strategy. If the betting is correct (i.e., the stimulus appears in the final waiting location), a situation similar to that in the spatial-cue condition arises. If the betting is wrong (i.e., the stimulus actually appears in the location opposite to where one is waiting), a situation similar to the center-cue arises – an additional attentional shifting (*notice-stimulus-with-doublecue-and-shift*) has to be performed when the stimulus appears. Since the betting is random, on average this gives rise to half of the cost of a production rule firing, which is about 20 ms. The simulation result is 19 ± 8 ms.

Finally, we would like to point out that our model manipulates five free parameters to obtain the goodness of fit. The first four fine-tune ACT-R's standard conflict resolution mechanisms and are mainly responsible for fitting the error rates. Specifically, two parameters – the odds of winning the conflict resolution for rule *random-left random-right* (1:20) and rule *hurryup-responding-no-matter-whether-target-or-not* (1:3) – are implemented in the model by setting their *p* values, which in ACT-R designate the probabilities that their firings will lead to eventual task success. The other two parameters, which are expected gain noise (*egs* = 3) and utility threshold (*ut* = -100), are set to add some random noise in ACT-R's conflict resolution so that the model can sometimes fire wrong rules and make mistakes in response. The fifth parameter, the firing time of each rule, is more critically responsible for fitting the RTs and attentional efficiencies. We have shortened the cost of each rule firing from ACT-R's default 50 to 40 ms. One

criticism of our model is that this parameter is widely accepted as one of a few of ACT-R's fundamental architectural primitives and changing it is indicative of a misuse of the architecture. Here, we offer two main justifications for this parameter change. First, the ANT task is extremely simple and has tight temporal constraints, which may speed up the rate of rule firing. Second and more importantly, this parameter should not be taken for granted. In particular, as long as the change is within a limited range and consistent across all the productions rules in a model, it may not be a misuse of the architecture. One underlying justification for the default setting, offered by EPIC, another rule-based cognitive architecture that excels in modeling human perceptual-motor performance (Meyer & Kieras, 1997a, 1997b), is that 50 ms conforms well with the mean period between zero crossings in the brain's 10 Hz α rhythm (e.g., Callaway & Yeager, 1960). Note, however, that the α rhythm is typically not fixed at 10 Hz but varies within a range (e.g., 7–13 Hz). From this perspective, shortening the rule firing cost to 40 ms conforms well with a slightly faster 12.5 Hz brain rhythm, still within the α wave's range but consistent with the tight temporal constraints of the task. Clearly, further investigations are necessary to evaluate our modification and clarify the implication of the parameter.

5. General discussion

Though attention is a complex multilevel cognitive function, the dominant approaches to computational modeling of attention have emphasized its close association with lower level processes such as perception and thus often adopted some connectionist and neural modeling techniques. In addition, typically only one type of attention is focused on. In this article, we reported a symbolic computational model of the ANT task, which incorporates the work of multiple attentional networks including alerting, orienting, and executive control. The model was developed in ACT-R. Using basic symbolic knowledge modeling structures such as rules and strategies, all of which often have clear and straightforward psychological

meaning, we were able to show that the operation of attentional networks could be understood at a symbolic level. In this section we would like to briefly discuss some implications of developing such a model.

First, our model demonstrates that as a multi-level construct attention can be modeled at both the subsymbolic and symbolic levels, and more importantly, that there might exist different mechanisms underlying the operation of attention at each level. For example, one distinct feature of our model is that it is a strict serial processing model: rules fire in sequence with one firing leading to another. This is different from various connectionist models, which extensively adopt parallel and distributed processing. While parallel distributed processing is certainly plausible at the neural and neural networks levels, serial processing is psychologically justifiable due to the limited capacity of human working memory (Baddeley, 1986; Baddeley & Logie, 1999). Being able to model the same behavioral regularities at both levels using different mechanisms offers another strong support for the multilevel nature of attention.

The concept of the production rule is fundamental to our model of attention. Rules fire in sequence and operate at a rate of about 50 ms (40 ms in the current model) per production rule. As argued by ACT-R, production rules define the atomic components of thought at the symbolic level. With this claim in mind, when we examine the efficiency measures of attentional networks reported in Fan et al. (2002) it seems that they make good sense at a rule level: these measures (51, 47, and 84 ms, for alerting, orienting, and executive control, respectively. See Table 2.) are just in 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 rule firings to explain symbolically the work of attentional networks. This is indeed what our model demonstrated.

While our model simulates the ANT results reasonably well, one question is whether there is any independent justification for the particular set of production rules used in our model. We address this question in two ways. First, as we reported

earlier, a functional analysis of the ANT task was performed before we developed the model. Such an analysis identified a list of psychologically meaningful functional components underlying the ANT task, which motivated and constrained the selection of production rules in the model. Second, our model can make nontrivial predictions and be applied to account for phenomena beyond the normal ANT task reported in this article. An example is attentional deficiency. The model predicts that if the rule *not-cue-so-switch-state-and-shift-attention* (rule 4 in Appendix A) is somehow nonfunctional (e.g., there is no such a rule or it takes a long time to fire), the alerting operation will be affected. A recent study (Booth et al., 2001) using ANT on ADHD children shows that these children revealed a damaged alerting function (i.e., a larger alerting effect), which may suggest that these children possessed a dysfunctional rule 4. By altering the operation of rule 4, we may quantitatively simulate these ADHD children's attentional performance. Another interesting prediction our model can make is about the effect of executive control. If the strength of the rule *hurryup-responding-no-matter-whether-target-or-not* (rule 25 in Appendix A) is increased due to certain psychological and/or pathological conditions, for example in schizophrenic patients, then the model would predict a deficit executive control due to a lower neural activities in the executive control attentional networks. It seems that these rules provide an interesting symbolic summary of the function of the underlying networks.

While a large body of evidence from neuroimaging and neuropathology has suggested that these distinct types of attention are subserved by independent attentional networks (see Fan et al., 2003c, for a review), our model provides some interesting implications about the operations of the attentional networks and the interrelationship among the different types of attention. The current model suggests that alerting can be explained by a rule that captures and performs a preparatory state change, orienting can be explained by a rule that moves attention before the appearance of the stimulus, and executive control can be explained by a set of rules that detect and resolve a conflict. Therefore, the neuroanatomical inde-

pendence of different attentional networks is simulated at a symbolic rule level by the interaction of different sets of rules fired in different task conditions. In this sense, it seems that at least at a symbolic rule level different attentional networks may adopt similar computational mechanisms.

Finally, modeling attention at a symbolic rule level through serial processing mechanisms raises an interesting question about the role of attention in cognitive architectures. As ACT-R moves towards a modular architecture (Anderson, Bothel, Byrne, & Lebiere, submitted), it seems that attention may be more properly modeled as an add-on module like other perceptual-motor modules (e.g., Byrne & Anderson, 1998). However, while a peripheral module may be appropriate for modeling the visuospatial type of attention (orienting), it appears unsuitable for alerting and executive control types of attention due to their intriguing relationship with central cognition. To us, production rules summarize nicely the conscious efforts and psychological reality associated with being alerted and resolving conflicts, as we demonstrated in our model. However, the need to reduce the cost of each rule firing from the ACT-R default of 50 ms to 40 ms in this model raises questions about the adequacy of adopting a purely serial symbolic approach to modeling attention. In general we envision that to model the full range of human attentional functions in ACT-R we may have to combine modular, subsymbolic, and rule-level approaches.

Acknowledgements

This research was supported in part by the Grant N00014-01-1-0074 from the Office of Naval Research Cognitive Science Program. We would like to thank Dr. Michael I. Posner for providing constructive comments that have improved the quality of the paper. Correspondence and request for reprints should be sent to Hongbin Wang, School of Health Information Sciences, University of Texas Health Science Center at Houston, 7000 Fannin Suite 600, Houston, TX 77030, USA. E-mail: Hongbin.Wang@uth.tmc.edu.

Appendix A. The production rules used in the model

Stage 1. Fixation and cue expectation

1. P notice-fixation
“An ANT trial starts when a fixation appears”
2. P encode-fixation-and-waiting
“Encode the fixation and start expecting for a cue”

Stage 2. Cue or stimulus?

3. P notice-something-appeared
“Some visual object appears on the screen”
4. P not-cue-so-switch-state-and-shift-attention
“Resist the surprise and shift attention to the stimulus”

Stage 3. Cue Processing

5. P notice-a-center-cue
“notice a center cue”
6. P notice-a-top-cue
“notice a top cue”
7. P notice-a-bottom-cue
“notice a bottom cue”
8. P given-a-top-cue-find-a-bottom-cue
“check if it is in fact a double cue”
9. P given-a-bottom-cue-find-a-top-cue
“check if it is in fact a double cue”
10. P find-no-more-cue-so-spatialcue
“No more cue so it is a spatial cue”
11. P find-more-cue-so-doublecue
“It is a double cue indeed”

Stage 4. Stimulus expectation

12. P anticipating-the-stimulus
“Wait and anticipate the stimulus”

Stage 5. Stimulus processing

13. P notice-stimulus-at-cued-top-location-and-attend
“Notice an arrow at the cued top location so encode it”
14. P notice-stimulus-at-cued-top-location-but-a-neutral-item-is-selected
“Notice a neutral item at the cued top location so refocus on the center item”
15. P notice-stimulus-at-cued-bottom-location-and-attend
“Notice an arrow at the cued top location so encode it”

16. P notice-stimulus-at-cued-bottom-location-but-a-neutral-item-is-selected
“Notice a neutral item at the cued top location so refocus on the center item”
 17. P notice-stimulus-with-centercue-and-shift
“The stimulus row appears while attention is on the center location so shift attention to it”
 18. P notice-stimulus-with-doublecue-and-shift
“The stimulus row appears in a location that is different from the current attended location in the doublecue condition so shift attention to it”
 19. P notice-stimulus-with-doublecue-and-an-arrow-is-focused-on-so-attend
“The stimulus row appears in a location that is same as the current attended location in the doublecue condition so encode the arrow”
 20. P notice-stimulus-with-doublecue-but-a-neutral-item-is-focused-on-so-shift
“The stimulus row appears in a location that is same as the current attended location in the doublecue condition but it seems a neutral distracter is selected so refocus”
 21. P attend-to-at-large-target
“Encode the arrow that is currently focused on”
 22. P shift-to-at-large-target-from-a-neutral-item
“If the currently focused on item is a neutral distract, refocus”
 23. P harvest-target
“Encode the target after refocusing”
 24. P goahead-responding-if-it-is-the-target
“If the encoded item is the center arrow, goahead to respond”
 25. P hurryup-responding-no-matter-whether-target-or-not
“As long as the encoded item is an arrow, hurry up to respond”
 26. P attended-item-is-right-to-the-target
“Hold on, the currently attended arrow is to the right of the target”
 27. P attended-item-is-left-to-the-target
“Hold on, the currently attended arrow is to the left of the target”
 28. P goahead-responding-if-congruent
“But move on to respond since they are congruent”
 29. P refocus-again-if-incongruent
“Refocus on the center location”
 30. P harvest-target-directly-if-incongruent
“Encode the center arrow directly since its location is available”
- Stage 6. Response*
31. P decide-left
“Decide to make a left response since the target is a left arrow”
 32. P decide-right
“Decide to make a right response since the target is a right arrow”
 33. P random-left
“Randomly decide to make a left response”
 34. P random-right
“Randomly decide to make a right response”
 35. P respond
“Respond by pressing the decided key”
 36. P refixating-and-wait-for-next-trial
“refocusing on the fixation and be ready for the next trial”

Appendix B. The trace in a center-cued incongruent trial

? (do-trial :cue centercue :flanker incongruent)

;; show fixation at 0 and for 1255 ms

Time 0.000: Notice-Fixation Selected

Time 0.040: Notice-Fixation Fired

Time 0.040: Module :vision running command
move-attention

Time 0.125: Module :vision running command
focus-on

Time 0.125: Encode-Fixation-And-Waiting Se-
lected

Time 0.165: Encode-Fixation-And-Waiting
Fired

Time 1.255: * Running stopped because time
limit reached.

;; show cue (or nocue) at 1255

Time 1.255: Notice-A-Center-Cue Selected

Time 1.295: Notice-A-Center-Cue Fired

Time 1.295: Module :vision running command
move-attention
Time 1.295: Anticipating-The-Stimulus Selected
Time 1.335: Anticipating-The-Stimulus Fired
Time 1.355: * Running stopped because time
limit reached.
;; *the cue disappears and continue fixating at 1355*
Time 1.380: Module :vision running command
focus-on
Time 1.755: * Running stopped because time
limit reached.
;; *the stimulus appears at 1755*
Time 1.755: Notice-Stimulus-With-Centercue-
And-Shift Selected
Time 1.795: Notice-Stimulus-With-Centercue-
And-Shift Fired
Time 1.795: Module :vision running command
find-location
Time 1.795: Attend-To-At-Large-Target Selected
Time 1.835: Attend-To-At-Large-Target Fired
Time 1.835: Module :vision running command
move-attention
Time 1.835: Attended-Item-Is-Left-To-The-Tar-
get Selected
Time 1.875: Attended-Item-Is-Left-To-The-Tar-
get Fired
Time 1.875: Module :vision running command
find-location
Time 1.920: Module :vision running command
focus-on
Time 1.920: Refocus-Again-If-Incongruent Se-
lected
Time 1.960: Refocus-Again-If-Incongruent Fired
Time 1.960: Module :vision running command
find-location
Time 1.960: Harvest-Target Selected
Time 2.000: Harvest-Target Fired
Time 2.000: Module :vision running command
move-attention
Time 2.000: Goahead-Responding-If-It-Is-The-
Target Selected
Time 2.040: Goahead-Responding-If-It-Is-The-
Target Fired
Time 2.085: Module :vision running command
focus-on
Time 2.085: Decide-Left Selected
Time 2.125: Decide-Left Fired
Time 2.125: Respond Selected

Time 2.165: Respond Fired
Time 2.165: Module :motor running command
press-key
Time 2.165: Module :vision running command
find-location
Time 2.165: Refixating-And-Wait-For-Next-
Trial Selected
Time 2.205: Refixating-And-Wait-For-Next-
Trial Fired
Time 2.205: Module :vision running command
move-attention
Time 2.290: Module :vision running command
focus-on
Time 2.315: Module :motor running command
preparation-complete
Time 2.365: Module :motor running command
initiation-complete
Time 2.375: Device running command output-
key
Time 2.465: Module :motor running command
finish-movement
Time 2.465: Checking for silent events.
Time 2.465: * Nothing to run: No productions, no
events.
;; *response at 2375, RT = 620 ms*
Time 4.090: * Running stopped because time limit
reached.
;; *trial ends at 4090*

References

- Altmann, E. M., & Davidson, D. J. (2001). An integrative approach to stroop: Combining a language model and a unified cognitive theory. In *Proceedings of the 23rd annual meeting of the Cognitive Science Society* (pp. 21–26). Hillsdale, NJ: Erlbaum.
- Anderson, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.
- Anderson, J. R. (1990). *The adaptive character of thought*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Anderson, J. R. (1993). *Rules of the mind*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Anderson, J. R., Bothel, D., Byrne, M., & Lebiere, C. (submitted). An integrated theory of the mind. *Psychological Review*.
- Anderson, J. R., & Bower, G. H. (1973). *Human associative memory*. Mahwah, NJ: Lawrence Erlbaum Associates.

- Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought*. Hillsdale, NJ: Lawrence Erlbaum Press.
- Baddeley, A. D. (1986). *Working memory*. Oxford: Oxford University Press.
- Baddeley, A. D., & Logie, R. H. (1999). Working memory: The multiple-component model. In A. Miyake & P. Shah (Eds.), *Models of working memory: Mechanisms of active maintenance and executive control*. New York: Cambridge University Press.
- Booth, J., Carlson, C. L., & Tucker, D. (2001). Cognitive inattention in the ADHD subtypes. In *Paper presented at the tenth scientific meeting of the International Society for Research in Child and Adolescent Psychopathology (ISRCAP)*, Vancouver.
- Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological Review*, 108, 624–652.
- Byrne, M. D., & Anderson, J. R. (1998). Perception and action. In *The atomic components of thought*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Callaway, E., & Yeager, C. L. (1960). Relationship between reaction time and electroencephalographic alpha base. *Science*, 132, 1765–1766.
- 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.
- Egeth, H. E., & Yantis, S. (1997). Visual attention: Control, representation, and time course. *Annual Review of Psychology*, 48, 269–297.
- Eriksen, B. A., & Eriksen, C. W. (1974). Effects of noise letters upon the identification of a target letter in a nonsearch task. *Perception and Psychophysics*, 16, 143–149.
- Fan, J., Flombaum, J. I., McCandliss, B. D., Thomas, K. M., & Posner, M. I. (2003a). Cognitive and brain consequences of conflict. *Neuroimaging*, 18(1), 42–57.
- Fan, J., Fossella, J. A., Sommer, T., Wu, Y., & Posner, M. I. (2003b). Mapping the genetic variation of executive attention onto brain activity. *Proceedings of the National Academy of Sciences USA*, 100(12), 7406–7411.
- Fan, J., McCandliss, 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., Raz, A., & Posner, M. I. (2003c). Attentional mechanisms. In *Encyclopedia of Neurological Sciences*. San Diego, CA: Academic Press.
- James, W. (1890). *Principles of psychology*. New York: Holt.
- Lovett, M. C. (2002). Modeling selective attention: Not just another model of stroop (NJAMOS). *Cognitive Systems Research*, 3(1), 67–76.
- Luck, S. J., Woodman, G. F., & Vogel, E. K. (2000). Event-related potential studies of attention. *Trends in Cognitive Sciences*, 4(11), 432–440.
- Meyer, D. E., & Kieras, D. E. (1997a). A computational theory of executive cognitive processes and multiple-task performance: Part 1. Basic mechanism. *Psychological Review*, 104(1), 3–65.
- Meyer, D. E., & Kieras, D. E. (1997b). A computational theory of executive cognitive processes and multiple-task performance: Part 2. Accounts of psychological refractory-period phenomena. *Psychological Review*, 104(4), 749–791.
- Mozer, M. C. (1991). *The perception of multiple objects: A connectionist approach*. Cambridge, MA: MIT Press.
- Mozer, M. C., & Sitton, M. (1998). Computational modeling of spatial attention. In H. E. Pashler (Ed.), *Attention*. East Sussex, UK: Psychology Press.
- 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., & Davies, D. R. (1984). *Varieties of attention*. San Diego, CA: Academic Press.
- Pashler, H. E. (1998). *The psychology of attention*. Cambridge, MA: MIT Press.
- Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology*, 32, 3–25.
- Posner, M. I., & Dehaene, S. (2000). Attentional networks. In M. S. Gazzaniga (Ed.), *Cognitive neuroscience: A reader*. Malden, MA: Blackwell Publishers.
- Posner, M. I., DiGirolamo, G. J., & Fernandez-Duque, D. (1997). Brain mechanisms of cognitive skills. *Consciousness and Cognition*, 6, 267–290.
- 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.
- Posner, M. I., Rothbart, M. K., Vizueta, N., Levy, K. N., Evans, D. E., Thomas, K. M., & Clarkin, J. F. (2002). Attentional mechanisms of borderline personality disorder. *Proceedings of the National Academy of Sciences USA*, 99(25), 16366–16370.
- Rueda, M. R., Fan, J., McCandliss, B. D., Halparin, J. D., Gruber, D. B., Lercari, L. P., & Posner, M. I. (in press). Development of attentional networks in childhood. *Neuropsychologia*.
- Rumelhart, D. E., & McClelland, J. L. (1986). Parallel distributed processing: Explorations in the microstructure of cognition. *Foundations* (Vol. 1). Cambridge, MA: MIT Press.
- Schacter, D. L., & Tulving, E. (1994). *Memory systems 1994*. Cambridge, MA: MIT Press.