Integrating Fast and Slow Cognitive Processes

William G. Kennedy (wkennedy@GMU.Edu)
George Mason University, 4400 University Drive
Fairfax, VA 22032 USA

Magdalena Bugajska (Magda.Bugajska@NRL.Navy.Mil)
Naval Research Laboratory, Code 5515, 4555 Overlook Ave. SW
Washington, DC 20585 USA

Abstract

Human reactions appear to be controlled by two separate types of mental processes: one fast, automatic, and unconscious and the other slow, deliberate, and conscious. With the attention in the literature focused on the taxonomy of the two processes, there is little discussion of how they interact. In this paper, we focus on modeling the slower process's ability to inhibit the fast process. We present computational cognitive models in which different strategies allow a human to consciously inhibit an undesirable fast response. These general strategies include (a) blocking sensory input, (b), blocking or interrupting the fast process's response, and (c) slowing down or delaying processing by introducing additional task. Furthermore, we discuss an approach to learning such strategies based on the inference of the causes and effects of the fast process.

Keywords: dual-processes, impulse control, inhibition, social behavior

Introduction

People appear to have two processes or systems controlling their actions: one fast, unconscious, or automatic and one slow, conscious, and deliberative (Kahneman 2003). Thus far the focus in the literature has been on discussing the differences in the processes in support of developing dual process theories of cognition (Evans 2008).

Evans (2008) provides an excellent review of the dual process theories of reasoning and decision-making. Although researchers use different terms for the two systems, almost all distinguish one system as "unconscious, rapid, automatic, and high capacity" while the other as "conscious, slow, and deliberative" (Evans, 2008). Researchers also differentiate between the systems saying the faster process is implicit and automatic and the slower is explicit and controlled. Many researchers also include the point that the faster process's control of behavior occurs without our being aware of the fact. The faster processing was described as "associative" and the slower process as "rule-based". Another theme reported was that the faster process was more concrete and situation specific and the slower, rational process more abstract and general. The key concept here is the characterization of the two systems by awareness and volition.

Our focus is on building a computational model of the interaction of these processes; specifically, we look at the ability of the slow, conscious process to inhibit the faster, automatic process. Blinking, for example, is one such fast, automatic action that with some effort can be inhibited. Under normal circumstance, blinking is an unconscious process occurring periodically whose rate is influenced by environmental conditions as well as internal, emotional state. But it is also well known that we can resist blinking. However, it is best described as "resisting" because it takes cognitive effort to not blink. The maintenance of our concentration is an example of the slow, cognitive process's inhibition on the blinking behavior. But when the concentration is broken, the fast, unconscious, and automatic process is back in control.

We propose that there are general strategies that humans use to inhibit the undesirable fast processes based on our ability to infer the causes and to detect the effects of those processes. We propose that a learned conscious process can effectively control the execution of the faster process through the control of the focus of attention and the deliberate common-resource management.

With this introduction, we will first discuss how the slow process can perceive the fast process and how the slow process can inhibit the fast process. We will then propose a general model integrating the fast and slow cognitive processes, present three instantiations of that general model, and discuss learning in these models before concluding.

Perception of a Fast Process

As Evans reported, many researchers noted that the faster process occurs without our awareness. Even though we may not be cognitively aware of the faster process while it is in progress, we can note its effect and infer its cause. When physical motion is involved, we have ability to attend to our own movement. In other words, we can sometimes sense the resulting action as soon as after it has been initiated, and definitely sense it after it has been completed. This is subject to the speed and the extent of the response as well as our focus of attention. Furthermore, Gladwell (2005) provided evidence that such fast, unexplainable processes can be the result of deep expertise we cannot easily articulate, but have ability to control including using them to our benefit as well as to inhibit them.

Humans are also capable of inferring a cause of a response. Whether it is attending to an environmental stimulus resulting in a movement, or an association between a memory and our emotional state resulting in an expression change, we can make the association.

For example, consider those nearly thoughtless responses to what we see, such as ducking a fast moving object, to what we hear, such as jumping at an unexpected sound, or even what we feel, such as uttering expletives or grimacing when we stub our toe, or smiling at a pleasant memory.

The ability to detect such effects and to infer the causes of the fast processes allows us to learn strategies to inhibit these fast processes. These general strategies for inhibiting them include (a) blocking the sensory input, (b) blocking (or interrupting) the response, and (c) running an additional process concurrently with the fast process. A general model of interaction of the two processes is shown in Figure 1. The undesirable fast process is represented as a direct Sense-Act thread while the desirable but slow process is shown below as a Sense-*Think*-Act thread. In the figure, the radar circle indicates the extent of changes to the focus of attention and the vertical lines are the boundaries between the cognitive model and the outside world. Attending to our own actions including vocalizations or facial expressions (indicated by the question mark icon in the figure), supports a deliberate choice or development of a control strategy.

Control of a Fast Process

To present how we envision a slow process can control a fast process, we begin by grounding both processes within a cognitive architecture. We will present three implementations of the general model as computational models within the ACT-R cognitive architecture (Anderson, 2007; Anderson et all, 2004). ACT-R is a symbolic and sub-symbolic, production-based cognitive architecture. The internal modules of ACT-R represent relatively specific cognitive functions (and regions of the brain) including declarative and procedural memory, auditory and visual perception, vocalization, and motor functions (based on the hand).

During each cycle, modules representing sensors fill buffers with representations of the environment. Like many production systems, ACT-R repeatedly matches production conditions with the contents of the buffers, but only selects a single production to fire, and then executes that production resulting in changes to internal buffers and module requests.

ACT-R, and more recently, jACT-R (Harrison & Trafton, 2010), have been embodied on a robotic platform which necessitated extension of motor functionality to control face muscles, head and limbs movements. For this project, we also added a rudimentary "emotional module" to allow us to keep track of the internal state of the robot. The emotions are based on appraisals according to the Appraisal Theory (Scherer, 2001; Marinier, et al, 2009), which are provided during the execution of the model. For example, unexpected stimulus is recorded automatically as it is being attended to, but the modeler could also issue an appraisal within a production to signify a successful completion of a goal or a failure. The intensity of the emotion is based on the number and recency of the appraisals along the dimensions indicative of the specific emotion. Unless the emotion is fueled after the initial event, it will decay over time; we

modeled the activation of the emotion on the base-level activation equation used in the recall of declarative memory (Anderson, 2004).
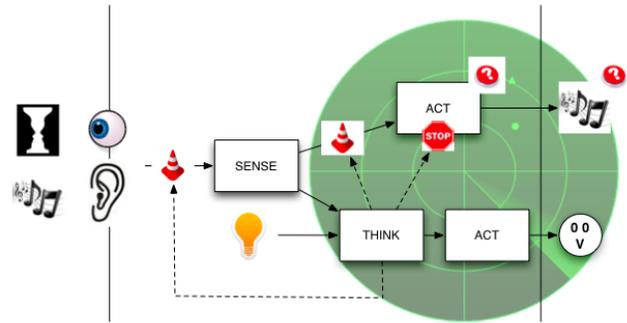

Figure 1. A General Model of Fast and Slow Process Integration.

Our theory of control of the fast process centers on the points at which its execution can be foiled. The alternative strategies leading to inhibition of the fast process are: (1) to block the perception of or attending to the relevant stimulus, and (2) to block the reaction to the stimulus, as indicated by the traffic cones graphic in Figure 1, and (3) running an additional process concurrently with the fast process, as indicated by the light bulb. It is also possible to interrupt or override, to certain degree, actions in progress, such as most large motions including face expressions.

Recall in the discussion of blinking, a slow, cognitive process could inhibit the fast, automatic blinking, but it took cognitive effort. We propose that, in general, it takes sustained cognitive effort to block fast responses. The blocking may not be completely effective in that there is evidence that like interrupting the non-blinking concentration, fleeting micro-expressions of emotion will still occur (Ekman & Friesen, 1969). An extreme example of blocking involves the patellar reflex test (the knee-jerk reaction). A patient can inhibit the normal knee jerk reaction but interrupting the patient's concentration allows the normal reaction to be observed. The common technique to break this concentration is the Jendrassik's Maneuver initially described in 1883 (Zehr & Stein 1999).

We propose that the slow process can both inhibit the faster process through the following alternative strategies:

(1) Intentionally blocking the stimulus by physically removing the stimulus, for example: by closing eyes or covering the ears, or by shifting the perceptual attention.

(2) Intentionally blocking the response by keeping the efferent processor busy, for example: performing another movement or subvocalizing to render the processor unavailable for other processes, or

(3) Intentionally performing another task at the same time. ACT-R supports this model of process interaction through:

(a) Allowing productions of various specificities.

(b) Buffer status queries including buffer contents and status at various phases of motor processing.

(c) Serialization of processing.

Below is a sample ACT-R production implementing a fast movement in response to an unexpected sound, which could be undesirable in context of many office tasks:

```
(p fast-response-to-sound ;production name
   =aural-location>       ;aural module detects
     isa    audio-event   ;  a sound
   ?aural-location>       ;the sound was
     buffer unrequested   ;  not expected
   ?manual>               ;the motor controller
     state  free          ;  is free, (not busy)
==>                       ;THEN
   +manual>               ;initiate a manual
     isa    press-key     ; action, press
     key    "return"      ; "return" key
)
```

For this production, the strategy to block the sensory input would be any action that would block the detection of an auditory event, such as covering one's ears with one's hands. To block the reaction part of this production, one needs to engage and keep the motor module unavailable because it is busy. Furthermore, due to ACT-R's adherence to serial processing, any other production whose utility is greater than this production would decrease the probability of the undesirable response.

Note that these strategies are temporary and require continuous attention, i.e., cognitive effort, to maintain the strategy. If the cognitive focus is interrupted and the sensory input is still present, the original fast response production will be able to fire.

Model Implementation

We will demonstrate the applicability of the general model by discussing its instantiation in three different models, specifically: (1) inhibiting the Stroop Effect through deliberate shift of visual attention, (2) inhibiting the startle reflex with respect to eye blinking, and (3) inhibiting socially unacceptable response in an emotional situation. Due to space constraints, we will present the model of only one of the alternate control strategies for each of these tasks, but other strategies are applicable as well.

Task: Inhibiting Stroop effect by blocking stimulus

Stroop (1935) identified a large increase in the time taken by participants to complete the color reading in the experiment that presented the participant with incongruent ink color and text, as compared to the naming of the colors of basic shapes. Original experiment has been extended and thoroughly studied over the years to determine in excess of 18 other effects (MacLeod, 1991). In this work we focus on the interpretation of the behavior within the dual processes presented earlier.

Our ACT-R model only captures relative speed difference between the color naming and word reading. Other researchers (Lovett 2002; van Maanen, van Rijn, & Porst, 2008) provide better models of an actual response times in the task, but ability to detect one's errors and to improve the performance at the cost of the response time is a focus of our model's implementation of the dual process theory.

When the fast word-reading process generates an incorrect response and it is detected due to a disparity between fast verbal response and the result of the intentional, but slower color naming process. As the response is being vocalized or as it was heard depending on the duration of the color vocalization process, an alternative strategy can be initiated. The easiest strategy simply calls for delaying, or in essence blocking the response, by pausing before giving the verbal response allowing time to reevaluate the color of the text.

As another strategy, Besner (2001) provides evidence that priming a location of a letter within the word eliminates the Stroop Effect. It stands to reason that a good, and in fact optimal, strategy would be for the participant to adjust visual attention accordingly hence blocking the word reading entirely. An easy way to achieve this is to upon or even prior to presentation of the stimulus, to shift attention to the right-most character of the text. With no competing response there is no need to confirm the answer and response can be given immediately.

To block the stimulus in our model, the automatic left-to-right visual search production competes with an intentional visual search production for the right-most symbol from the current location. As long as the expected location is attended to, the word reading (fast process) will not have a chance to happen resulting in a single and correct response.
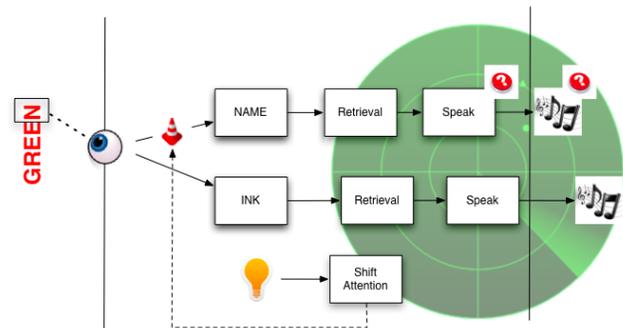

Figure 2. Inhibiting Stroop Effect by shifting gaze.

Task: Inhibiting startle reflex by blocking response

The startle reaction, also startle reflex, is the response to a sudden unexpected stimulus, such as a flash of light, a loud noise, or a quick movement near the face. These reactions include movement away from the stimulus, a contraction of arm and leg muscles, a verbal response, and often blinking. It also includes blood pressure, respiration, and breathing changes that are often described as being startled or scared.

In this section, we focus on the acoustic startle reflex, a response to an unexpected, loud, and near sound on the order of 40ms in duration. Specifically, we present an ACT-R model in which intentionally keeping eyes open inhibits blink-response to the acoustic event. Like other strategies described in this paper, muscle contraction is only a temporary strategy since it requires constant focus to maintain; any lapse in attention will result in muscle relaxing and ability for any process including the startle or routine physical maintenance reflex to control the muscle. Our ability to control blinking is often tested in a staring

contest. Due to the speed of the response, which on average, takes between 300 and 400 milliseconds to complete, this strategy works best when initiated before the stimulus is heard to act to prevent rather than override the reflex or fast response.

Our ACT-R system is capable of perceiving and attending to a sound. The general model strategy to engage the muscle in expectation of the stimulus translates in ACT-R to keeping motor module busy. Assuming the concentration can be maintained and the muscle stays engaged, the fast process's impulse to blink will be blocked. To capture the cognitive effort involved in this strategy, we allow the goal to be removed from focus of attention and the motion to be no longer than 350 ms. The model detects the unintentional motion, based on lack of the intention to move the muscle and presence of the motion.
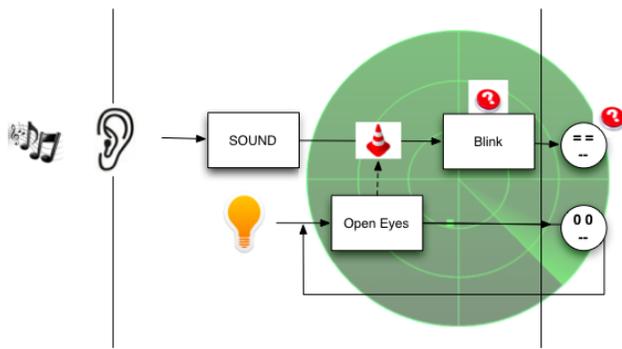

Figure 3. Preventing blinking.

This is definitely not the only strategy that can be used. Interestingly, Fillon, et al. (1993) presented an experiment which showed that an attended pre-pulse, a weaker pre-stimulus, produced greater blink inhibition at the 120 ms lead interval than an ignored pre-pulse. Obviously, covering your ears (or closing your eyes in the case of visual stimulus) is an effortless strategy and guarantees better performance, but is only feasible when task allows for it.

Both of these instantiations of the general model involve blocking the fast process. The next instantiation of the general model develops an acceptable alternative to an emotional response.

Task: Inhibiting emotional response by distraction

Thomas Jefferson is credited with having said "When angry, count to ten before you speak. If very angry, to a hundred," which even nowadays is considered a sound advice since time and distraction are key to anger management. An emotional response is a fast process behavior that rarely leads to positive result, especially in social interactions. However, given time to calm down, most people can get a handle on their initial impulses.

Evans reported that although some researchers ignore emotions in their discussions of the two systems, others place emotions within the faster process and some contemporary work includes an emotional influence in the slower, more deliberative process. Due to this lack of

consistency, Evans considered emotions outside the scope of his review of dual systems theory, but we will regard the basic, spontaneous emotional responses as the fast processes.

Ekman identified basic emotions including joy and anger, as being universally recognized from facial expressions (Ekman, 1992; Ekman, 1999). The automatic nature of his basic emotions included specification that the processing was very fast, between 150 and 250 ms. Another researcher, Griffiths (1997), suggested some emotions are higher-level introspective processes, i.e., belonging to the slower, more deliberative process. Others have suggested classifying emotions based on the part of the brain that is activated by the emotion, either the amygdala or prefrontal cortex (Evans, 2001; Frank, 2009). This later differentiation is useful here because although both classifications involve the brain in the response to emotions, the separation of the high-level cognitive function from the low-level processes based on the region of the brain involved, serves our purposes.

While an emotion can be treated as either a stimulus or as a response, for the sake of our argument, we will consider an emotion state as a perceivable stimulus. The emotional responses vary widely and include changes in vocalization characteristics and content, flailing arms or legs, and obviously as facial expressions. For ease of explanation, in the current instantiation of the model, we assume that emotions can be perceived as form of an internal state akin to perception of time (Taatgen, Van Rijn, & Anderson, 2007).

In this instantiation of the general model, we simulate the behavior of an individual that is impatiently waiting for a stimulus to appear (e.g. imagine waiting for a bus or a friend while time is wasting). Since we will be focusing on blocking the undesirable response, the actual stimulus that is cause of the anger is not relevant. Upon stimulus presentation, specifically, the bus or friend's arrival, the subject vocalizes the response based on the emotional state of the model. (See Figure 4.) The model monitors its emotional state as well as the response. A negative reward is associated with the undesirable response (or positive reward is associated with the socially acceptable response).
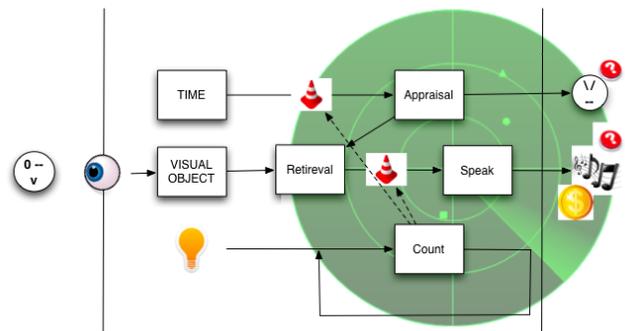

Figure 4. Preventing an emotional response.

As the passage of time is attended to, a negative appraisal is recorded and the model becomes angry. When the stimulus is detected, a fast response process is initiated. At first, the

process does not include the counting to ten and results in a negative, unacceptable response. The counting process triggered by intention to speak while angry, has the property of delaying the response to allow the emotion to decay, and it also distracts the perception of time process from "adding fuel to the fire."

A similar delay tactic can be employed during Stroop task to reinforce the color-naming process. Before giving the answer, the participant could confirm that the response is indicative of the task, which would force the color information of the stimulus to be processed independently. Detecting the conflict is resolved by the conduct (repeat) of a deliberate process to produce the correct answer. Our model of this strategy rewards the response from the deliberate process and may explain the observed brain activity associated with conflict detection and cognitive control (Egner & Hirsch, 2005).

Role of Learning

The feasibility of the strategies discussed in the previous section relies on two forms of learning. First, the alternative, slow process has to be crafted based on the input and output characteristics of the fast process. Second, the model has to learn that the alternative process is useful.

Our general model calls for learning of a control strategy upon detection of an unexpected and undesirable condition. The strategies presented in the task models were hand-crafted. We expect that a problem-solving process focused on addressing the causes of the undesired behavior can develop these strategies. Based on the realization that the causes involve both a stimulus and a response, we expect to be able to learn strategies that involve blocking both the stimulus as in the model of the first task and the response as in the model of the second task. Additionally, introducing a delay or distraction process can be learned if it can be inferred that the causes are time sensitive. This is, of course, subject for future research.

Once the control strategy, i.e., the slow, conscious process, has been crafted, it will eventually become procedurelized and compete with the fast, unconscious process productions. ACT-R utility learning provides the necessary mechanism. In accordance with the ACT-R theory, the utility of a production is determined based on its presence and position in the sequence leading to the reward; specifically, a negative reward issued upon detection of an unexpected and undesirable model behavior leads to relative increase of alternate processes. Since, in the tasks presented here, the fast process is the cause of the unexpected events, this reward mechanism results in the reinforcement of the slower processing path. For example, by punishing the sequence of productions leading up to undesirable response, we lower their utility allowing the counting process to have the higher utility and be included in execution on subsequent runs. Due to this approach, our task models make testable predictions that human error rates in experiments like the Stroop Effect should decrease over time and the response times should be representative of the shift between the two processes.

Essential to both forms of learning is detection of an incorrect or undesirable response. We define an error as an inconsistency between the fast and slow processes' responses indicating a need to decide which is the intended response. Within an ACT-R model, such inconsistencies are described by contents of the relevant buffers. For example, as we have described in the startle reflex task, the detection of a movement when none is expected indicated that a fast, unconscious process was being executed. It should be noted that attending to these cues requires additional processing and given the dynamics of the processing, such cues can be easily missed. Due to this approach, our task models make testable predictions that learning can be part of repeated tests of the Stroop Effect and that learning will not occur if the task dynamics preclude detection and adaptation.

Discussion

In the tasks modeled here, the fast process provided the wrong or undesirable response; this is not true in general. Humans have long depended on these impulses or reflexes to keep us safe as well as to provide the fast responses required in many tasks. Essentially, while slow, rational thinking has its role in our behavior, so does actually allowing the fast, irrational process guide us *in a controlled manner*. We have described how the slow process can control the fast process. However, this is only a beginning.

However, we have not yet presented evidence that our integration of the two processes matches experimental data. Several experiments are suggested by this work including re-visiting the Stroop Effect looking for learned strategies and performance over time.

Conclusions

We have shown that what has been widely discussed as a dual processes, one fast, automatic, and unconscious and the other slow, deliberate, and conscious, can be implemented within a single cognitive architecture and we provided a general model of their integration. We instantiated this general model using the ACT-R architecture and showed the slow process's control of the fast process in three different tasks. The general model's fast-process-control strategies we implemented and demonstrated included: (a) blocking the sensory input for the fast process, (b) blocking (or interrupting) the response from the fast process, and (c) substituting a slow process for the fast process. Finally, we discussed the architectural ability to reinforce the slow process's control of the fast process and an approach to learning the alternate processes.

Acknowledgments

References

Anderson, J. R. (2007). *How Can the Human Mind Occur in the Physical Universe?* Oxford University Press.

Anderson, J. R., Bothell, D., Byrne, M. D., Douglas, S., Lebiere, C., and Qin, Y. (2004). An integrated theory of mind. *Psychological Review* 111(4): 1036-1060.

Berger, A. J. (2008). What causes muscle atonia in REM? *Sleep, 31.* 1477-1478.

Besner, D. (2001). "The myth of ballistic processing: Evidence from Stroop's paradigm." *Psychonomic Bulletin & Review, 8(2)*, 324-330.

Chase M. H. & Morales, F. R. (2005). Control of motoneurons during sleep. In M. H. Kryger, T. Roth and W. C. Dement (Eds.) *Principles and practice of sleep medicine*. 4th ed. Philadelphia: WB Saunders. 154-68.

Egner, T. & Hirsch, J. (2005). The neural correlates and functional integration of cognitive control in the Stroop task. *NeuroImage, 24, 539-547.*

Evans, D. (2001). *Emotion: The science of sentiment*. Oxford, U.K.: Oxford University Press.

Evans, J. S. B. T. (2008). Dual-processing accounts of reasoning, judgment and social cognition. *Annual Review of Psychology*, *59*, 255-278.

Ekman, P. (1992). An argument for basic emotions. *Cognition and Emotion*, *6*, 169-200.

Ekman, P. (1999). Basic emotions. In T. Dalgleish and M. Power (Eds.) *Handbook of Cognition and Emotion*. Sussex, U.K.: John Wiley & Sons, Ltd.

Ekman, P. (2003). Darwin, Deception, and Facial Expression. *Annals N. Y. Academy of Science, 1000*, 205-221.

Ekman, P. & Friesen, W. V. (1969). Nonverbal leakage and clues to deception. *Psychiatry, 32,* 88-105.

Frank, M. J., Cohen, M. X., & Sanfey, A. G. (2009). Multiple systems in decision making: A neurocomputational perspective. *Current Directions in Psychological Science*, *18*, 73-77.

Gladwell, M. (2005). *Blink: The power of thinking without thinking.* New York, NY: Little, Brown and Company.

Griffiths, P. E. (1997). *What emotions really are: the* problem of psychological categories. Chicago: The University of Chicago Press.

Harrison, A.M., & Trafton, J.G. (2010). Cognition for action: an architectural account for "grounded interaction" *Proceedings of the Annual Conference on Cognitive Science*. Aug. 11-14, 2010.

Kahneman D. (2003). A perspective on judgment and choice. *American Psychologist, 58*, 697-720.

Lovett, M. C. (2002) Modeling selective attention: Not just another model of Stroop (NJAMOS). *Journal of Cognitive Systems Research, 3*, 67-76.

MacLeod, C. M. (1991). "Half a century of research on the Stroop effect: an integrative review". *Psychological Bulletin, 109 (2)*, 163–203.

Marinier, R., Laird, J. & Lewis, R. (2009). A Computational Unification of Cognitive Behavior and Emotion. *Journal of Cognitive Systems Research, 10, 48-68*.

Salo, R., Henik, A., & Robertson, L. C. (2001) Interpreting Stroop interference: An analysis of differences between task versions. *Neuropsychology, 15(4)*, 462-471.

Scherer, K. (2001). Appraisal considered as a process of multilevel sequential checking. In K. Scherer, A. Schorr, & T. Johnstone (Eds.), Appraisal processes in emotion: Theory, methods, research. New York: Oxford University Press.

Stroop, J. R. (1935). Studies of Interference in Serial Verbal Reactions. *Journal of Experimental Psychology, 18*, 643-662.

Taatgen, N., Van Rijn, H., Anderson, J. R. (2007). An integrated theory of prospective time interval estimation: The role of cognition, attention, and learning. *Psychological Review, 114 (3)*, 577-59

Van Maanen, L., Van Rijn, D. H., & Porst, J. P. (2008). Stroop and picture-word interference are two sides of the same coin. *Psychonomic Bulletin & Review, 18*, 987-999.

Zehr, E. P., & Stein, R. B., Interaction of the Jendrassik maneuver with segmental presnaptic inhibition. *Experimental Brain Research, 124 (4)*, 474-480.