TRANSIENT ATTENTIONAL ENHANCEMENT DURING THE ATTENTIONAL BLINK: EEG CORRELATES OF THE ST^2 MODEL

Srivas Chennu*, Patrick Craston, Brad Wyble and Howard Bowman

Centre for Cognitive Neuroscience and Cognitive Systems,
University of Kent, Canterbury CT2 7NF, United Kingdom
**Email: sc315@kent.ac.uk*
Web: www.cs.kent.ac.uk/people/rpg/sc315

There has been considerable recent interest in the identification of neural correlates of the Attentional Blink (AB), and the development of neurally explicit computational models. A prominent example is the Simultaneous Type Serial Token (ST^2) model, which suggests that when the visual system detects a task-relevant item, a spatially specific Transient Attentional Enhancement (TAE), called *the blaster*, is triggered. This paper reports on our investigations into EEG activity during the AB, and a hypothesized correlation between the blaster and the N2pc ERP component. Specifically, we demonstrate that the temporal firing pattern of the blaster in the model matches the N2pc component in human ERP recordings, for targets that are seen and missed inside and outside the attentional blink window. Such a correlation between a computational account of the AB and ERP data provides useful insights into the processes underlying selectivity in temporal attention.

Keywords: Attentional Blink; ST^2 Model; Transient Attentional Enhancement; ERP; N2pc.

1. Introduction

The Attentional Blink (AB)[1,2] is a well studied temporal attention phenomenon, particularly suitable for investigating the nature and limits of conscious perception. The AB employs Rapid Serial Visual Presentation (RSVP), in which a sequence of items is presented at the same spatial location at a rate of around 10 items per second, with each item rapidly replacing the previous one. At such speeds, the items presented, some of which are targets to be detected and others distractors, yield only fleeting mental representations. The AB describes the finding that the detection of a second target (T2) following a correctly identified first target (T1)

is significantly impaired if T2 follows T1 within 200–600 ms. There has been considerable recent interest in the identification of neural correlates of the AB, and the development of neurally explicit models, a prominent one being the Simultaneous Type Serial Token (ST^2) model.[3] In addition to incorporating a computationally explicit account of visual processing, attentional selection and working memory (WM) encoding, the model proposes the episodic distinctiveness hypothesis, i.e., that the AB reflects the visual system attempting to allocate unique episodic contexts to targets. Importantly, it suggests that when the visual system detects an item that may be task relevant, a spatially specific Transient Attentional Enhancement (TAE), called the *blaster* is triggered. For a fleeting stimulus, the contribution of this enhancement is critical in enabling it to be encoded into WM.

This paper reports on our investigations into EEG activity during the AB, and a hypothesized correlation between the blaster and the N2pc ERP component. The N2pc describes a negative deflection of the ERP at around 200–300 ms after the presentation of a laterally offset target, and is most strongly visible at parietal electrodes contralateral to the position of the target. Previous research has associated the N2pc with the selection of a target in the presence of competing distractors,[4] and as an indicator of the moment-to-moment deployment of attention.[5] In this paper, we describe an experiment employing a dual stream AB paradigm designed to record lateralized electrophysiological (EEG) activity during RSVP. We discuss experimental data that, when examined in light of predictions from the ST^2 model, suggests a correlation between the triggering of the blaster and the manifestation of the N2pc. Specifically, we demonstrate that the temporal firing pattern of the blaster in the ST^2 model matches the N2pc component in human ERP recordings, for targets that are seen and missed inside and outside the attentional blink window. This connection supports the hypothesis that the N2pc reflects the selective attentional enhancement that is embodied by the blaster.

This paper is organized as follows. Section 2 begins with a brief description of the ST^2 model, specifically focusing on the blaster and its influence on model dynamics. Section 3 shifts focus to a discussion of how blaster activation traces can be meaningfully compared to the N2pc. Section 4 describes the dual stream experiment designed to record the N2pc and elaborates on the correlational evidence gathered to connect it to the blaster. Section 5 concludes with a general discussion and proposes a potentially reciprocal relationship between cognitive modelling and ERPs.

2. The ST^2 Model

The Simultaneous Type Serial Token theory of temporal attention, developed and explained in-depth in Ref. 3, provides a specific account of temporal information processing in the visual system and the formation of WM representations. Its connectionist realization, the neural-ST^2 model as depicted in a simplified form in Figure 1, implements its principles in a fixed-weight neural network. The architecture of the model can be divided into two stages of processing, both of which interact with the *blaster*, which provides a brief but powerful attentional enhancement in response to the occurrence of targets. This two-stage design, inspired by Ref. 1, supports the hypothesis that the AB is characterized by a late-stage processing bottleneck as described below.

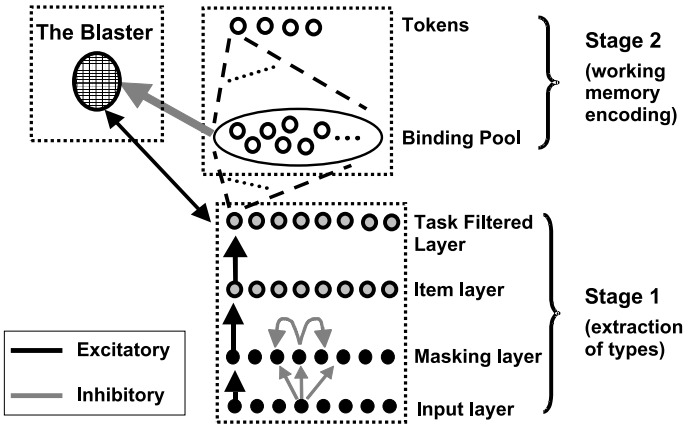

Fig. 1. The Neural-ST^2 Model

Stage 1 abstractly models early visual processing common to all stimuli, including visual masking and semantic categorization. Stimuli are processed in parallel and filtered based on task specific salience as they pass through Stage 1. Its pipelined design implies that the representation of a given item in Stage 1 is cascaded across multiple layers at any time.[6] In its final layer, Stage 1 generates rapidly decaying representations of featural characteristics of items, which in this framework are called *types*.[7] Only targets trigger the blaster, which provides additional excitation to support their successful encoding into WM.

Stage 2 creates and maintains durable representations of targets in WM. Targets at the end of Stage 1 pass into Stage 2 by binding their type to a *token*. Unlike in Stage 1, this *tokenization* process is strictly sequential, and attempts to associate distinct episodic contexts to targets. This serialization of target encoding is enforced by actively inhibiting the blaster and preventing it from refiring during ongoing tokenization.

The blaster is triggered by the detection of a target at the end of Stage 1, which in turn elevates activation levels across the final layers of Stage 1. This transient attentional enhancement provided by the blaster lasts for a brief (150ms) window of time following target detection. The availability of this attentional boost from the blaster is essential for most targets because of the fleeting representations generated in RSVP, which usually do not have sufficient bottom up strength to initiate WM encoding. A key feature of the blaster is that it fires only once per tokenization. Once a target (T1) triggers the blaster and the process of binding is initiated, it is held offline by inhibition from Stage 2 till the process is complete. This inhibition collapses the attractor previously set up, and attempts to associate distinct episodic contexts to targets. It does so by preventing a second target (T2) in close temporal proximity to T1 from interfering with its tokenization. T2 must "wait" till the T1 tokenization is complete for the blaster to fire again. This implies that only T2s with strong bottom up strength have enough activation to "outlive" T1 tokenization. It is this mechanism, embodying the episodic distinctiveness hypothesis, that enables the ST^2 model to simulate the attentional blink.

Model Performance Figure 2 compares the performance of the ST^2 model to human behavioural data relating to the AB.[1] It focus on performance in the basic AB scenario, when T1 is followed by a blank, and when T2 is at the end of the RSVP stream. In addition to these scenarios, the model reproduces a broad spectrum of AB data. See Ref. 3 for a detailed comparison.

3. Connecting Modelling and ERPs

Computational modelling in the context of the AB has tended to focus on replicating behavioural data. The ST^2 model itself was conceptualized and designed to simulate patterns of behavioural data collected in AB experiments. In this paper, we look at whether the connection between modelling and experimentation can be extended from the behavioural to the electro-

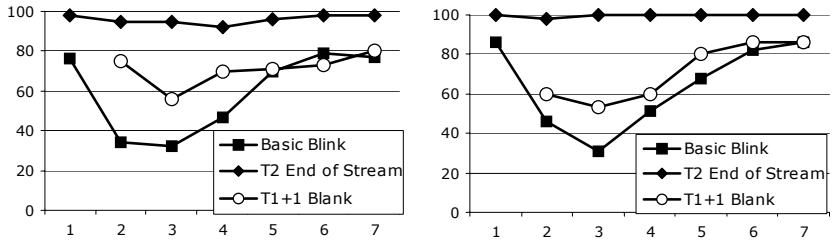

(a) Human (b) Model

Fig. 2. Comparison of the ST^2 model to behavioural data

physiological domain. Although successful replication of a broad spectrum of behavioural data relating to the AB is convincing validation of the ST^2 model, being able to relate ERPs to specific parts of the model allows for a more fine-grained verification. To make these correlations, we build upon the correspondence between specific parts of the model and successive stages of temporal visual processing that generate ERP components. Though the approach used in this paper to connect modelling and electrophysiology is exploratory, it allows for the use of another dimension of experimental evidence for validating specific processing stages proposed in the model.

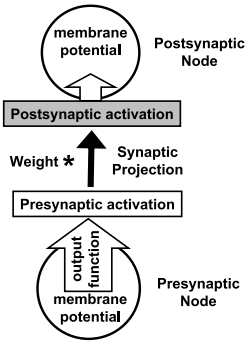

Fig. 3. A pair of model neurons

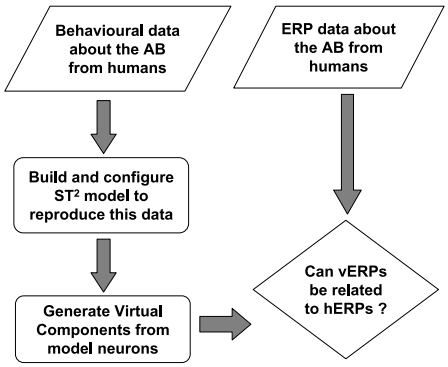

Fig. 4. Can vERPs be related to hERPs?

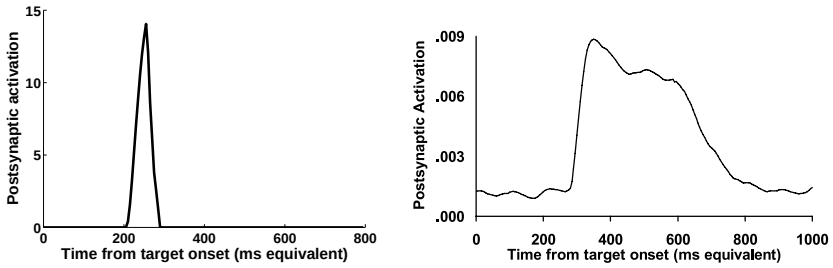

(a) Blaster Output (b) Virtual P3

Fig. 5. Virtual ERPs from the the ST^2 model

Virtual ERPs from ST^2 In order to compare activation dynamics of the neural-ST^2 model to human ERPs, we generate *virtual ERPs* from traces produced by specific processing layers recorded during simulations. Virtual ERPs are thus defined as grand averages of neural activation, summing over all possible combinations of target strengths, and represent the typical pattern of activity set up in the model when a particular experimental condition is simulated. Figure 3 depicts a pair of nodes in the model, connected by a weight that represents an excitatory synaptic projection. The membrane potential of the presynaptic node feeds into its output function and the resulting presynaptic activation contributes to the postsynaptic membrane potential after being multiplied by the intervening weight. This method to generate virtual ERPs is intended to best approximate the mechanism assumed to be responsible for the generation of human ERPs. Electrical activity observed at the scalp is thought to reflect the summation of the postsynaptic potentials generated at a large number of spatially aligned pyramidal neurons in the cortex,[8] which are known to release the excitatory neurotransmitter glutamate. Analogously, a virtual ERP comparable to a human ERP is the average excitatory postsynaptic potential recorded across functionally equivalent layers in the model. The functional role of the layers chosen to generate a particular virtual ERP is hypothesized to correspond to the neural processing that generates the corresponding human ERP. Though external factors like scalp distortion have not been considered in generating virtual ERPs, we think that being able to make qualitative comparisons and predictions about human ERPs using virtual ERPs can contribute to the process of cognitive modelling.

Figure 5 shows sample virtual ERPs generated from the ST^2 model. Figure 5(a) depicts the burst of activation generated by the blaster in response

to a single target. Figure 5(b) depicts the virtual P3 generated during the encoding of that target, calculated by averaging activation across those layers in the model that simulate target selection and WM encoding by associating the type of the target with a token.

Given this definition of virtual ERPs, the question we pose is summed up in the diagram depicted in Figure 4. Empirical research on the AB has produced a substantial dataset of behavioural and ERP data. The ST^2 model is used to generate virtual ERPs from model layers that simulate distinct cognitive functions. In juxtaposing human ERP data with these virtual ERPs, we are interested in investigating whether being able to find correlations between them leads to useful insights into both the architecture of the model and the human ERPs themselves.

3.1. *The Blaster and the N2pc*

We now state the key hypothesis of this paper; that there is a correlation between the blaster component in the ST^2 model and the N2pc ERP component. The blaster, as has been described thus far, is responsible for providing the transient attentional burst necessary for encoding targets into WM. The N2pc ERP component, on the other hand, is a well-studied negative deflection occurring in the ERP waveform 200–300 ms after the onset of a salient stimulus. The N2pc is a lateralized component, in that it is larger at parietal electrodes contralateral to the stimulus position, and is usually plotted as a difference waveform obtained by subtracting the ipsilateral electrode from the contralateral one. The N2pc has been observed in spatial visual search experiments[4,5,9] and in RSVP paradigms.[10–12] It is thought to reflect the locus of visual spatial attention, and the tracking of the instantaneous deployment of attention.

This definition of the N2pc has similarities with the functional effect of the blaster. The preliminary hypothesis that we put forward is that the N2pc corresponds to the firing of the blaster in response to targets. In order to verify this hypothesis, we attempt to connect the firing pattern of the blaster to the relative strength of the N2pc in different experimental conditions. A key prediction from this hypothesis is that the blaster firing pattern, and consequently the N2pc amplitude, are fundamentally different for targets "inside" (T2 at lag 3) and "outside" (T1) the blink window. As will be discussed further in Section 4.2, all T1s get the benefit of the blaster and elicit a similar N2pc, irrespective of whether or not they are consciously reported. This is in contrast to T2s occurring during the blink, which are available for conscious report only if they get the benefit of the blaster.

Hence, seen T2s elicit a larger N2pc than missed ones. In order to test this hypothesis, we conducted an EEG experiment employing a dual stream AB paradigm designed to record the N2pc. The details of the experiment and the results obtained are the topic of the next section.

4. Dual Stream Experiment

We now describe an experiment employing a dual stream RSVP paradigm designed to produce the attentional blink and simultaneously record EEG activity.

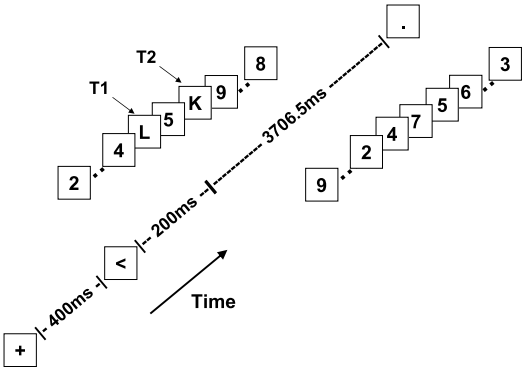

Fig. 6. A sample trial in the dual stream RSVP experiment

4.1. *Design*

Our experimental paradigm employed a pair of laterally presented RSVP streams consisting of letters and digits, where letters were targets to be identified, and digits were distractor stimuli. Figure 6 depicts the timeline of a sample trial. In each trial, 35 stimuli at 105.9ms SOA were presented in each of two RSVP streams at either side of fixation. The actual streams were preceded by a central fixation cross. This cross was replaced by an arrow after 400ms, indicating one of the two RSVP streams in which two target letters would appear. Participants were instructed to direct their covert attention to the indicated stream while continuing to fixate their gaze centrally. The streams began 200ms after the arrow and lasted for 3.706s. Immediately after the end of the streams, the central arrow turned into either a dot or a comma, chosen randomly, which stayed on for 105.9ms.

After the end of the trial, participants reported the identity of any targets they saw, and a dot or a comma for the last item. This additional task was included to ensure the participants fixated centrally throughout the presentation of the streams. Each participant was presented with 4 blocks of 100 trials each. A given trial could contain either 0 or 2 randomly chosen letter targets, appearing equally randomly in one of the two streams. The second target (T2) appeared at lag positions 1, 3 or 8 after the first one (T1).

Participants Experimental data from 14 university students (mean age 22.4, SD 2.9; 6 female; 13 right-handed) were included in the analysis. All were free from neurological disorders and had normal or corrected-to-normal vision.

EEG Recording Scalp EEG was recorded at 1000Hz (bandpass filtered at 0.25Hz–80Hz during recording) while participants performed the task, from 19 electrodes placed at standard 10/20 locations (Fp1, Fp2, Fz, F3, F4, F7, F8, Cz, C3, C4, C7, C8, Pz, P3, P4, P7, P8, O1, O2, T7 and T8). In addition, a bipolar EOG channel below and to the left of the left eye recorded eye movements, used to reject trials with artefacts. Recorded EEG data was referenced to a common average online and re-referenced to linked earlobes offline. An electrode at the left mastoid acted as ground.

EEG Analysis For each subject, EEG segments for conditions of interest were time-locked to the onset of the target and extracted from -200ms to 800ms with respect to target onset. For each such segment, direct current drift artefacts were removed using a DC de-trend procedure employing the first and last 100ms of each segment. The segments were baseline corrected to the 0–150ms interval following target onset that does not contain any lateralized ERP components, and then averaged together for each condition. The average additional activity elicited by the targets for a particular subject was then calculated by subtracting the contralateral parietal electrode from the ipsilateral one, either P7 or P8, in order to isolate the N2pc. For plotting the waveforms, they were low pass filtered at 25Hz for visual clarity.

4.2. *Results and Discussion*

This section compares human ERPs from the dual stream experiment and virtual ERPs from the ST^2 model, with regard to the hypothesis put forward in Section 3.1

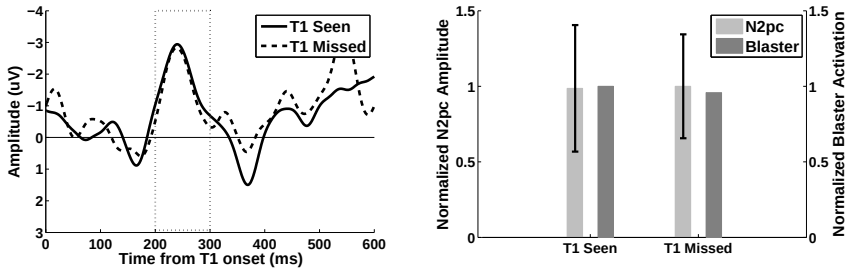

(a) N2pc for seen vs. missed T1s (b) Blaster for seen vs. missed T1s

Fig. 7. Comparing the N2pc and the blaster for targets outside the blink

Targets outside the blink Figure 7(a) depicts the N2pc elicited by seen vs. missed T1s in the 200–300ms window following T1 onset. As can be clearly seen, the N2pc has similar amplitude, onset and offset for both conditions. This fact is reflected in the results of a repeated measures ANOVA that shows no significant effect of condition ($F(1, 13) = 0.001, MSE = 1.466, p = 0.98$). This implies that T1s, i.e., targets presented outside the blink window, elicit a similar N2pc, irrespective of whether they are eventually seen or missed. To correlate this finding with the ST^2 model, Figure 7(b) compares the normalized postsynaptic activation of the blaster to the normalized N2pc amplitude for the same pair of conditions. The postsynaptic activation plotted here is averaged over all possible target strengths in the model. Two key observations can be made from this comparison: firstly, the amplitudes of the blaster for seen and missed T1s are similar. This can be understood by noting that in the ST^2 model, the blaster always fires for T1s as it is not inhibited by any ongoing tokenization. Hence T1s that are missed in the model are too weak to complete tokenization, despite the attentional enhancement provided by the blaster. Secondly, the average amplitude of blaster activation covaries with the N2pc for seen and missed T1s.

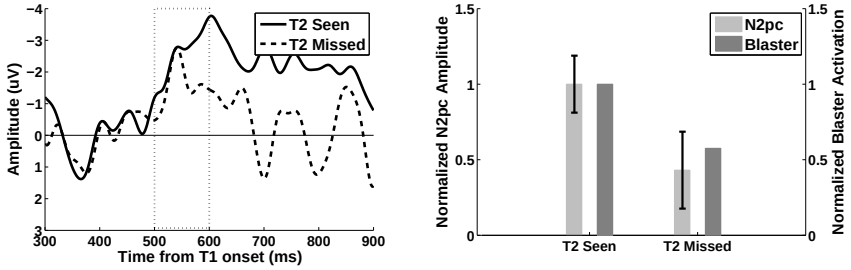

(a) N2pc for seen vs. missed T2s (b) Blaster for seen vs. missed T2s

Fig. 8. Comparing the N2pc and the blaster for targets inside the blink

Targets inside the blink Figure 8(a) depicts the N2pc elicited by seen vs. missed T2s at lag 3, i.e. well within the attentional blink, during the 200–300ms window following T2 onset. The human ERPs have been computed by averaging over only those trials in which T1 was seen. In contrast to the previous comparison, the N2pc for seen vs. missed T2s is markedly different. Specifically, it is larger in amplitude and area for seen T2s. An ANOVA supports these observations by suggesting a significant effect of condition ($F(1, 13) = 8.133, MSE = 2.699, p = 0.013$). This variation in the N2pc is mirrored by the postsynaptic activation of the blaster, as depicted in Figure 8(b). The blaster is inhibited during the blink window and can fire again only after T1 tokenization has completed. Only those T2s at lag 3 that have enough activation left at the end of T1 tokenization manage to refire the blaster. Of these, some are nevertheless missed due to insufficient bottom-up strength. As a result, the average blaster activation for missed T2s at lag 3 is significantly lesser than that for seen T2s.

5. Conclusions

The comparisons between the blaster and the N2pc emphasize a key property of the ST^2 model of the attentional blink: the unavailability of the transient attentional enhancement for targets during the blink window. This *temporal spotlight* of attention provided by the blaster, though necessary for tokenization, is not sufficient. The ERP data shows that both seen and missed T1s elicit a similar N2pc, but the N2pc for seen T2s at lag 3 is significantly larger than that for missed T2s. Together, the N2pc data and the blaster activation traces point toward a key observation: the deployment of transient attentional enhancement is fundamentally different for targets inside and outside the blink.

Implications for Modelling and ERPs This paper also highlights the more general idea of connecting cognitive modelling and ERPs. Though the attempt to do so in this paper is a qualitative one, the results are encouraging. To a large extent, the computational detail of the ST^2 model has enabled us to generate virtual ERPs to compare and contrast with human ERPs. But this level of detail comes with implementation costs. As with any model required to reproduce a broad spectrum of data, there is a trade-off between modelling capability and design complexity, since attempting to model a broad spectrum of data tends to increase the computational requirements of the model.

In spite of these challenges, being able to connect model dynamics to human ERP data has reciprocal benefits. Firstly, generating ERPs from computational models allows us to make predictions about the pattern of variation in human ERPs that we can expect to see across experimental conditions. Models inspired by neurophysiologically plausible architectures can be used to theorize about and direct the effort to better understand the neural sources of ERPs in the brain. Reciprocally, ERP data can be used to comparatively evaluate competing theories of psychological phenomena.

References

1. M. Chun and M. Potter, *Journal of Experimental Psychology: Human Perception and Performance* **21**, 109 (1995).
2. J. Raymond, K. Shapiro and K. Arnell, *Journal of Experimental Psychology: Human Perception and Performance* **18**, 849 (1992).
3. H. Bowman and B. Wyble, *Psychological Review* **114(1)**, 38 (2007).
4. M. Eimer, *Electroencephalography and Clinical Neurophysiology* **99**, 225 (1996).
5. G. F. Woodman and S. J. Luck, *Journal of Experimental Psychology: Human Perception and Performance* **29**, 121 (2003).
6. J. L. McClelland, *Psychological Review* **86**, 287 (1979).
7. N. Kanwisher, *Journal of Experimental Psychology: Human Perception and Performance* **17**, 404 (1991).
8. S. J. Luck, *An Introduction to the Event-Related Potential Technique* (MIT Press, Cambridge, MA, 2005).
9. S. J. Luck, M. Girelli, M. T. McDermott and M. A. Ford, *Cognitive Psychology* **33**, 64 (1997).
10. P. Jolicoeur, P. Sessa, R. Dell'Acqua and N. Robitaille, *European Journal of Cognitive Psychology* **18(4)**, 560 (2006).
11. P. Jolicoeur, P. Sessa, R. Dell'Acqua and N. Robitaille, *Psychological Research* **70(6)**, 414 (2006).
12. R. Dell'Acqua, P. Sessa, P. Jolicoeur and N. Robitaille, *Psychophysiology* **43**, p. 394 (2006).