

University of Memphis

University of Memphis Digital Commons

CCRG Papers

Cognitive Computing Research Group

2012

A LIDA-based Model of the Attentional Blink

T. Madl

S. Franklin

Follow this and additional works at: https://digitalcommons.memphis.edu/ccrg_papers

Recommended Citation

Madl, T., & Franklin, S. (2012). A LIDA-based Model of the Attentional Blink. Retrieved from https://digitalcommons.memphis.edu/ccrg_papers/47

This Document is brought to you for free and open access by the Cognitive Computing Research Group at University of Memphis Digital Commons. It has been accepted for inclusion in CCRG Papers by an authorized administrator of University of Memphis Digital Commons. For more information, please contact khggerty@memphis.edu.

A LIDA-based Model of the Attentional Blink

Tamas Madl (tamas.madl@gmail.com)

Department of Philosophy (Cognitive Science), University of Vienna

Dr. Karl Lueger Ring 1, Vienna, 1010 Austria

Stan Franklin (stan.franklin@memphis.edu)

Institute for Intelligent Systems, University of Memphis

Memphis, TN 38152 USA

Abstract

The attentional blink (AB) refers to the impairment in consciously perceiving the second of two targets presented in close temporal proximity (200 – 500ms) in a rapid serial visual presentation paradigm. The present paper is a preliminary report describing a conceptual and partially computational model of the AB based on the LIDA (Learning Intelligent Distribution Agent) cognitive architecture. The model aims to provide a biologically plausible explanation of the AB, explaining a wide range of AB-related phenomena, among other mental phenomena accounted for by LIDA. Computational results in a basic visual AB paradigm are presented and compared to human data.

The Attentional Blink

When subjects are asked to identify two targets separated by a short time (200-500ms) in a stream of distractors, an Attentional Blink (AB) occurs - subjects often fail to report the second target (see Fig 1). In this paper we will focus on the AB in rapid serial visual presentations (RSVP) of pictures (however, the AB has been shown to occur across a wide range of stimuli types and modalities – see (Martens & Wyble 2010)).




Fig 1. The visual attentional blink paradigm (from Potter et al. 2010). Subjects who have to identify targets in a stream of images often fail to report the second target (T2) if it is presented shortly after the first target (T1).

Brain related evidence has shown that during an AB task, both targets are processed at least perceptually, regardless of conscious reportability - at least the first 150ms of neural activity exhibits a normal pattern (Martens & Wyble 2010). An fMRI study conducted by (Marois et

al. 2004) showed parahippocampal place area activations (associated with high-level scene representations) even in non-conscious T2 targets. However, EEG studies have revealed the electrophysiological activity that correlates with the AB – the N2pc ERP¹ component, occurring about 200ms poststimulus and associated with the allocation of attention to targets - is suppressed at short temporal distances between T2 and T1. Also, in trials where T2 cannot be perceived because it is presented shortly after T1, the P3 component - associated with working memory consolidation - is not elicited (Martens & Wyble 2010), (Dux & Marois 2009). The above evidence implies that the AB has to occur at a later stage of processing (later than perceptual recognition, and after 150ms).

Apart from this finding, a number of attentional blink related phenomena have been found, some of which have proven hard to explain – no complete, formal account for all of these has been found yet (Dux & Marois 2009). Elaborating on all the AB-related effects that have been identified would exceed the scope of this paper (See Martens & Wyble 2010 and Dux & Marois 2009 for more phenomena). The following phenomena have been chosen to highlight the ABs main properties, and to show that while only a simulation of the basic AB paradigm is presented in this preliminary paper, our LIDA-based model can provide much wider explanations. Future work will be required to computationally simulate and verify them.

1) Lag-1 sparing. Paradoxically, T2 can be reported with high accuracy if presented shortly after T1 (about 100ms after T1; “lag n” describes the temporal distance between the targets) (Martens & Wyble 2010).

2) Spread lag-1 sparing. Multiple targets can be reported as long as they are presented in immediate succession – it has been observed that target reports were accurate even for four successive targets (Olivers et al. 2007).

3) Posttarget intrusion. Varying the experimental conditions revealed that the AB only occurs if T2 is backward masked (Giesbrecht & Di Lollo 1998). Often, this mask or distractor succeeding T2 can be reported even if T2 cannot, implying that the distractor somehow interferes with the reporting of the target (Chun 1997).

4) Whole report attenuates the AB. The accuracy of reporting stimuli is high when subjects are asked to report all stimuli (whole report). However, a significant accuracy

¹ Event-Related Potential, brain activity directly resulting from and time locked to a stimulus

drop at lags 2 – 4 (an AB) can be observed for the same stimuli sequence if subjects are required to report only two targets in the sequence (Potter et al. 2008).

5) Increasing T2 salience/arousal attenuates the AB. If the salience of the second target is increased, it can be reported more accurately, although an AB effect can still be observed (Martens & Wyble 2010). Emotional arousal (but not valence) also alleviates the AB (Anderson 2005).

6) Task-irrelevant cognitive load attenuates the AB. If the stimuli are presented together with a background field of moving or flickering dots, much smaller drops in accuracy are observed at AB-relevant lags (Arend et al. 2006). The AB is also attenuated if subjects are asked to listen to task-irrelevant music or think about their holiday (Olivers & Nieuwenhuis 2005).

7) Target Confusion. The order targets were presented in is often confused for temporally adjacent targets (i.e. during lag-1 or spread lag-1 sparing) (Dux & Marois 2009), (Chun 1997).

8) AB without T1 masking. Although T2 masking is necessary to obtain an attentional blink, recent studies found that there is an (attenuated) blink even if T1 is unmasked (i.e. if there is no distractor between T1 and T2) (Nieuwenstein et al. 2009).

As will be described below, the LIDA model's attention mechanism is capable of explaining these phenomena, and is detailed enough to make a computational implementation reproducing an actual AB experiment feasible.

Attention in the Brain

The following brain areas play a role in top-down attentional control. The visually selective regions of the posterior parietal cortex – the intraparietal sulcus in humans and lateral intraparietal area in primates -, which contain coarse representations of spatial topography, and are also involved in controlling eye movements (saccades) and directing them towards targets. In this area, neuronal activity correlated with the voluntary allocation of attention can be observed, and leads to greater target stimulus selectivity based on spatial location and/or salience (Knudsen 2007), (Serences & Yantis 2006). The frontal eye fields (FEF) in the pre-frontal cortex are involved in saccade control as well, but have also been shown to play a role in representing the current locus of attention (Serences & Yantis 2006) – FEF neurons can be covertly selective to targets, without shifting gaze (Thompson et al. 2005). Also, when a particular stimulus is attended to and conscious, neurons representing the target in sensory areas, in the PPC and in the PFC exhibit synchronized discharges in the gamma band (Knudsen 2007), (Doesburg et al. 2009).

The superior colliculus (SC), like FEF, mediates both overt saccades and covert shifts of attention (Serences & Yantis 2006). Corollary discharges associated with eye saccades occur in the SC and propagate to the FEF, via the mediodorsal thalamic nucleus, and shift the locations of visual receptive fields in FEF before each saccade (Knudsen 2007).

Recently, the locus coeruleus (LC) – norepinephrine (NE) – system has been shown to influence top-down attentional selection (Nieuwenhuis et al. 2005)(Sara 2009), (Warren et al. 2009). LC neurons exhibit increased activity during the processing of motivationally salient targets, leading to the release of NE in widespread cortical projection areas, among others in the forebrain. This increased NE presence increases the responsibility of target neurons, enhancing signal detection and stabilizing a neural representation against noise or interference. Thus, LC activity can facilitate the processing of a target stimulus (Nieuwenhuis et al. 2005), (Warren et al. 2009), (Aston-Jones & Cohen 2005).

Shortly after target processing, there is an increase of LC neuronal activity (Fig 2A). The LC is autoinhibitory – increased activity during target processing is followed by a




Fig 2. Peri-event time histograms (PETHs) for a typical individual monkey LC neuron in response to various events during performance of the signal detection task, accumulated over 100 sweeps of activity. Note the increased activity during target processing (panel A). Adapted from (Aston-Jones & Cohen 2005).

period of reduced activity and thus reduced NE release. Non-target stimuli do not elicit significant changes in LC activity (Fig 2B). Interestingly, tonic (regular spiking) LC activity levels were significantly higher when the animal was less focused on the task (Fig 2C), but there was also a much smaller target-locked phasic response in this case. Behaviorally, these elevated LC activity levels have led to more frequent false alarm errors (Aston-Jones et al. 1999).

These properties of the LC-NE system play a key role in LIDA's attentional blink model.

A LIDA-based Attentional Blink Model

The LIDA cognitive architecture is based on prevalent neuroscience and cognitive science theories (Global Workspace Theory, situated cognition, perceptual symbol systems, working memory, memory by affordances, transient episodic memory, etc. - (Baars & Franklin 2009)). It has been implemented computationally, and has been shown to perform simple psychological tasks with mechanisms similar to humans (Madl et al. 2011). LIDA's cognitive cycle has the purpose of selecting an appropriate action based on the perceived current situation, and has the following components (Madl et al. 2011):

- 1) *Perception*. Sensory stimuli are received and stored in a sensory buffer in the Sensory Memory. Percepts, emotions, and concepts are represented by nodes in the Perceptual Associative Memory (PAM). These are based on perceptual symbols (Barsalou 1999); their activations reflect recognition confidence as well as bottom-up salience.
- 2) *Percept to preconscious buffer*. Recognized percepts are stored in the preconscious buffers of LIDA's long-term working memory (Workspace).
- 3) *Local associations*. Local associations are automatically retrieved from the Transient Episodic and Declarative Memory using the Workspace contents.
- 4) *Competition for consciousness*. Attention codelets² (AC) in the Attention Codelet Module (ACM) view long-term working memory, and compete to bring novel, relevant, important, urgent, or insistent events to consciousness.
- 5) *Conscious broadcast*. A coalition of codelets, typically an AC and its content of PAM nodes, gains access to the Global Workspace (GW) and has its content broadcast consciously.
- 6) *Recruitment of resources*. Relevant behavioral schemes in Procedural Memory respond to the conscious broadcast.
- 7) *Activation of schemes in the Procedural Memory*. Schemes are instantiated in the Action Selection module, and receive activation, based on the conscious contents.
- 8) *Action chosen*. The Action Selection module chooses a single scheme from the newly instantiated schemes and remaining previously active schemes.
- 9) *Action taken*. The execution of the action of a scheme results in external or internal consequences, or both.

The major components implementing top-down attention in the LIDA model are the GW module and the ACM. Feature detectors (corresponding to feature-sensitive neurons in the visual cortices) pass activation to their corresponding PAM nodes, which represent objects (or categories, concepts, ...) and could correspond to neuronal ensembles in the inferior temporal cortex, which contain object category information (Liu et al. 2009). The resulting activation of PAM nodes will depend on the number of relevant features, as well as the salience of those features.

The ACM contains ACs, which create coalitions from important or relevant percepts in the Workspace. The coalition with the highest activation will be broadcast consciously. Coalition activation depends on four factors: a) the activations of the percepts it contains, b) the base level activation of the AC, c) the modulatory activation of the ACM and d) a matching factor on how well the percept matches the pattern that the Codelet is looking out for. The computational implementation of the LIDA AB Agent also contains a fifth factor, e) stochastic noise, which is added to account for extraneous, uncorrelated afferent activity (Knudsen 2007), (Nieuwenhuis et al. 2005).

² The term codelet refers generally to any small, special purpose processor or running piece of computer code. The concept is essentially the same as Baars' (1988) processors or Minsky's (1988) agents. The term was borrowed from (Hofstadter & Mitchell 1994).

The first factor a) corresponds to bottom-up salience in the brain, as described above. The second, b), the base level activation, depends on how useful the AC has been in the past and facilitates attentional learning.

The third factor, c), is the modulatory activation of the ACM. It has been proposed many times in attention literature that human attentional processing is limited for targets presented in short succession - observable, among others, in an AB paradigm -, presumably because of a suppression of attentional enhancement of subsequent stimuli during the processing of a target (Nieuwenhuis et al. 2005), (Wyble et al. 2009), (Olivers & Meeter 2008). The modulatory activation reflects this mechanism, and regulates attentional enhancement of stimuli by increasing or decreasing the activation of coalitions in the Global Workspace. The most probable neural counterpart of this regulatory activity is the LC, which can enhance target processing through the release of NE in the forebrain (LC activity was proposed to play a role in the attentional blink by (Nieuwenhuis et al. 2005)). Similarly to LC neuron activity, the ACM activation at first increases upon processing a relevant or important target, followed by a period of low activation which is similar to the posttarget refractory-like autoinhibition exhibited by the LC (Fig 3 bottom). The ACM activation is governed by a function derived from interpolating LC PETH data (Aston-Jones & Cohen 2005).

The fourth parameter d) influencing coalition activation is a matching factor that is based on how well the percept in a coalition matches the pattern sought by the AC that creates the coalition. This accounts for the finding that in some cases, nontargets are attended to and reported instead of the targets if they are similar or share a common salient feature (Martens & Wyble 2010), (Lavie & Cox 1997), (Bichot & Schall 1999) although with less probability and less selective neuronal activation (Duncan et al. 1997).

LIDA's attentional mechanism can provide a computational explanation for the attentional blink and related findings. Two major reasons are proposed to account for the performance drop at intervals of 200ms – 500ms between the two targets (see Fig 3A bottom): a) the posttarget refractory-like period of the ACM activation, which leads to reduced target activations after ~200ms, and b) the discrete, competitive conscious broadcast mechanism (Baars & Franklin 2009).

For the current description, an RSVP attentional blink paradigm with images is assumed (see Fig 1). Stimuli are presented to the LIDA agent at a rate of one image every 107ms. The agent's task is to report target images pertaining to a specific target (in this case, vehicles), which means that there are at least two ACs, looking out for targets (vehicles) and distractors, respectively. This is also the paradigm used for the implementation of the LIDA Attentional Blink agent. If only a single target is presented, that target is added to a coalition by the Target Attention Codelet (TAC), will win the competition for consciousness since there is nothing that could compete with it, and can be consciously reported.




Fig 3. A) The attentional blink at lag 2. T_n and D_n refer to targets and distractors, respectively. The vertical black lines intersecting with the timeline on top represent the approximate borders of LIDA cognitive cycles. AC1 is looking out for targets, and AC2 for distractors, adding them to Coalitions in the Global Workspace. The coalitions have to compete for consciousness, and the one with the highest activation is broadcast consciously. The reason the agent fails to report T2 is that in the second cognitive cycle, Coalition 2 (containing the distractors) wins the competition for consciousness. B) LC activity - PETH of a monkey LC neuron during target processing.

If two targets are presented in an RSVP of images at lag-1, without a distractor, both targets are perceived in the first 200ms – before the refractory-like period of the ACM – and they are both added to a target coalition by an AC looking out for targets. This TAC has higher base level activation than the Distractor Attention Codelet (DAC). Thus the targets will win the conscious broadcast and can be reported consciously. Possible subsequent targets are also added to the target coalition by the same AC, which adjusts the coalition activation based on the factors described above and on the previous coalition activation – this accounts for the spread lag-1 sparing effect.

At lags 2 and 3, the second target sometimes cannot be reported consciously because a coalition containing

distractors wins the competition for consciousness instead of T2 (see Fig 3 and 4). The reason for the low activation of the target coalition is the low ACM activation at this point in time (due to the refractory-like period, see Fig 3 bottom). D2 is added to the distractor coalition by the DAC, and the coalition activation is updated. The distractor coalition is also modulated with a lower ACM activation, but will come out with a higher activation because a) depending on the timing of the presentation, the ACM activation might be higher at the point the distractor is perceived than at the point when the target is perceived, and b) since the distractor coalition was created upon perceiving D1, at which point the coalition activation was higher (0.4 in Fig 3, due to the high ACM activation at that point).

At lag 4, the ACM activation has regenerated to its initial level of activation, and T2 can be reported with a high level of accuracy again (the T2 accuracy at lag 4 approximately equals T2 accuracy at lag 1 in this paradigm, see Potter et al. 2010).

Results

LIDA's attentional mechanism conceptually accounts for all of the AB-related phenomena described above:

1-2) Lag-1 and spread lag-1 sparing. See above.

3) Posttarget intrusion. During the blink, the distractor succeeding T2 often can be consciously reported even if T2 itself cannot (see Fig 3).

4) Whole report attenuates the AB. In case of an instruction to report the entire RSVP sequence, a different Attentional Codelet would be required, which would move every presented image into the Global Workspace and into the same coalition – every image would be a target. Thus, for short RSVP sequences, every image could be reported and no AB could be observed (if the sequence is too long, activation decay could lead to “forgetting” of the first images. There is also a limit on how much information the Workspace and the Global Workspace can hold, although this limit has not been quantitatively determined yet).

5) Increasing T2 salience/arousal attenuates the AB. Increased T2 bottom-up salience leads to a higher activation of the PAM node representing T2 and thus to a target coalition with a higher activation, which increases the probability that T2 wins the competition for consciousness. In the case of emotional content with high arousal (Anderson 2005), a PAM node representing this emotion (with an activation value corresponding to the arousal) would be included in the coalition along with the target representation (Franklin et al., in press), increasing its activation and the probability of its conscious broadcast.

6) Task-irrelevant cognitive load attenuates the AB. Subject less focused on a task exhibit higher levels of tonic LC activity (see lowest panel of Fig 2), which can explain this phenomenon. In experimental conditions in which moving dots are presented around the target, and in conditions where the subject is instructed to think about something else, subjects are less focused on the AB task – therefore their AttentionCodelet Module Activation

(modeling LC activity) is higher before a target, and there is a much smaller post-target activation drop, which makes it possible to almost always report T2 accurately – in accordance with behavioral AB experiments with distracted subjects (Olivers & Nieuwenhuis 2005), (Arend et al. 2006), (Nieuwenhuis et al. 2005).

7) Target Confusion. Targets presented temporally adjacently in the same cognitive cycle (e.g. during lag-1 sparing) land in the same coalition. Since coalitions do not contain ordering information, the temporal order of the two targets is unknown to the agent. This could account for the target confusion effect. However, no exact computational mechanism has been implemented yet for reproducing how human subjects “guess” (often incorrectly, see Dux & Marois 2009, Chun 1997) the first target in such a case.

8) AB without T1 masking. There is an AB effect without the T1 mask, since the AB in this model is due to a) the refractory-like period of the ACM and b) intrusion of the post-T2 distractor, and doesn’t depend on the post-T1 distractor.

The above explanations show that our model of the AB is capable of explaining more than just the basic AB paradigm simulated in this preliminary report. To the authors’ knowledge, no model provides detailed explanations for every AB effect described above. Our LIDA based model could be able to do so, provided further computational simulations that can verify these results; and is more general than most AB models since it grows out of a universal model of cognition instead of being specific to the AB like most AB models (except for the Threaded Cognition model, which is based on ACT-R). Of the models published in the last five years, the following have the most explanatory power (Dux & Marois 2009): eSTST (Bowman & Wyble 2007), Attention Cascade (Shih 2008), Threaded Cognition (Taatgen et al. 2009), and Boost and Bounce (Olivers & Meeter 2008). Table 1 shows a comparison of these models and the proposed LIDA-based model. They are similar to our model in that they also rely on a suppressed/delayed attentional enhancement of T2; the major differences, apart from the LIDA architecture’s generality and plausibility (Baars & Franklin 2009), are the following. First, apart from a depleted attentional resource (ACM activation), which our model shares with these models (except for the Threaded Cognition model, which instead of a depleted resource relies on an unnecessary consolidation protection rule) there is also competition with the post-T2 distractor in the Global Workspace. Second, the postulated theoretical reason for T2 not being reported in some short lag trials is that the T2 percept does not win the competition for consciousness, i.e. the gamma-coherent neuronal ensembles representing T2 do not become part of the large-scale theta-gamma synchronized network representing conscious contents, which is consistent with the recently implicated importance of oscillatory activity in the AB (Janson & Kranczioch 2011). This theta-gamma oscillatory synchrony is proposed to be the neuronal basis of functional consciousness, and of the global broadcast in the LIDA model (Madl et al. 2011).

↓ AB-related phenomenon	LIDA	eSTST Bowman & Wyble 2007	Attention Cascade Shih 2008	Threaded cognition Taatgen et al. 2009	Boost and Bounce Olivers & Meeter 2008
Lag-1 sparing and spread lag-1 sparing	✓	✓	✓	✓	✓
Posttarget intrusion	✓	✓	✓	✓	✓
Whole report attenuates the AB	✓	✓	✗	✓	✗
Increased T2 salience attenuates the AB	✓	✓	✓	✓	✓
Increased T2 arousal attenuates the AB	✓	✗	✗	✗	✗
Task-irrelevant cognitive load attenuates the AB	✓	✗	✗	✓	✗
Target confusion	✓	✓	✓	✓*	✓
AB without T1 masking	✓	✓	✓	✓	✗

Table 1. A comparison of the LIDA-based AB model, and other models, conceptually accounting for AB phenomena. *: The original ACT-R based implementation does not account for target confusion, but the authors did include a simulation of target confusion using a custom visual module

The model has been implemented computationally to reproduce an experiment similar to (Potter et al. 2010) (Fig 1) to show that it can model human behavior. Every 107ms an image is presented to the agent, and it has to report two targets (images of vehicles) in a stream of distractors (faces, for easy discrimination; see Fig 3A top. The images were taken from the Caltech image database. The second target succeeds the first either immediately after 107ms (lag 1), or after a distractor (213ms, lag 2), or after 3 distractors (427ms, lag 4). Human reporting accuracies in such a setting are displayed in Fig 4 (Potter et al. 2010).

The LIDA Attentional Blink Agent is based on the LIDA computational framework (Snider et al. 2011). Its environment consists of the screen displaying the images, and of three buttons for each possible response (first target, second target, distractor). Images are recognized using a number of feature detectors looking out for scale and rotation invariant features. These implemented feature detectors are based on Speeded Up Robust Features (Bay et al. 2008). The task instructions were pre-defined in the form of a TAC and DAC bringing relevant images to consciousness, and schemes in the Procedural Memory for reporting targets by pressing buttons. Fig 4 shows the LIDA AB Agent’s performance in this task, compared to human data. These result were obtained using the same framework, and the same parameters, as previous LIDA agents (Madl et al. 2011), except for the addition of the ACM activity modeling LC activation.




Fig 4. Comparison of human and the LIDA AB Agent’s performance in a visual AB task. Left panel adapted from (Potter et al. 2010).

References

- Anderson, A.K., 2005. Affective influences on the attentional dynamics supporting awareness. *J Exp Psychol Gen*, 134(2), p.258.
- Arend, I., Johnston, S. & Shapiro, K., 2006. Task-irrelevant visual motion and flicker attenuate the attentional blink. *Psychon Bul Rev*, 13(4), pp.600–607.
- Aston-Jones, G. & Cohen, J.D., 2005. An integrative theory of locus coeruleus-norepinephrine function: adaptive gain and optimal performance. *Annu Rev Neurosci*, 28, pp.403–450.
- Aston-Jones, G., Rajkowski, J. & Cohen, J., 1999. Role of locus coeruleus in attention and behavioral flexibility. *Biol Psych*, 46(9), pp.1309–1320.
- Baars, B.J., 1988. *A cognitive theory of consciousness*, Cambridge Univ Pr.
- Baars, B.J. & Franklin, S., 2009. Consciousness is computational: The LIDA model of global workspace theory.
- Barsalou, L.W., 1999. Perceptual symbol systems. *Behav Brain Sci*, 22(04), pp.577–660.
- Bay, H. et al., 2008. Speeded-up robust features (SURF). *Comput Vis Image Und*, 110(3), pp.346–359.
- Bichot, N.P. & Schall, J.D., 1999. Effects of similarity and history on neural mechanisms of visual selection. *Nat Neurosci*, 2, pp.549–554.
- Bowman, H. & Wyble, B., 2007. The simultaneous type, serial token model of temporal attention and working memory. *Psychol Rev*, 114(1), p.38.
- Chun, M.M., 1997. Temporal binding errors are redistributed by the attentional blink. *Atten Percept Psychophys* 59(8), pp.1191–1199.
- Doesburg, S.M. et al., 2009. Rhythms of Consciousness: Binocular Rivalry Reveals Large-Scale Oscillatory Network Dynamics.
- Duncan, J., Humphreys, G. & Ward, R., 1997. Competitive brain activity in visual attention. *Curr Opin Neurobiol*, 7(2), pp.255–261.
- Dux, P.E. & Marois, R., 2009. The attentional blink: A review of data and theory. *Atten Percept Psychophys*, 71(8), pp.1683–1700.
- Franklin, S. et al., in press. LIDA: A Systems-level Architecture for Cognition, Emotion, and Learning. *BICA*.
- Giesbrecht, B. & Di Lollo, V., 1998. Beyond the attentional blink: Visual masking by object substitution. *J Exp Psychol Human*, 24(5), p.1454.
- Hofstadter, D.R. & Mitchell, M., 1994. The Copycat Project: A model of mental fluidity and analogy-making. *Adv. in connectionist and neural comp theory*, 2, pp.31–112.
- Janson, J. & Kranczioch, C., 2011. Good vibrations, bad vibrations: Oscillatory brain activity in the attentional blink. *Adv Cog Psych*, 7(-1), pp.92–107.
- Knudsen, E.I., 2007. Fundamental components of attention. *Annu. Rev. Neurosci.*, 30, pp.57–78.
- Lavie, N. & Cox, S., 1997. On the efficiency of visual selective attention: Efficient visual search leads to inefficient distractor rejection. *Psychol Sci*, 8(5), p.395.
- Liu, H. et al., 2009. Timing, timing, timing: Fast decoding of object information from intracranial field potentials in human visual cortex. *Neuron*, 62(2), pp.281–290.
- Madl, T., Baars, B.J. & Franklin, S., 2011. The Timing of the Cognitive Cycle. *PLoS ONE*, 6(4), p.e14803.
- Marois, R., Yi, D.J. & Chun, M.M., 2004. The neural fate of consciously perceived and missed events in the attentional blink. *Neuron*, 41(3), pp.465–472.
- Martens, S. & Wyble, B., 2010. The attentional blink: past, present, and future of a blind spot in perceptual awareness. *Neurosci Biobehav R*, 34(6), pp.947–957.
- Minsky, M., 1988. *The society of mind*, Simon and Schuster.
- Nieuwenhuis, S. et al., 2005. The role of the locus coeruleus in mediating the attentional blink: A neurocomputational theory. *J Exp Psychol G*, 134(3), p.291.
- Nieuwenstein, M.R., Potter, M.C. & Theeuwes, J., 2009. Unmasking the attentional blink. *J Exp Psychol Human* 35(1), p.159.
- Olivers, C.N. & Meeter, M., 2008. A boost and bounce theory of temporal attention. *Psychol Rev*, 115(4), p.836.
- Olivers, C.N. & Nieuwenhuis, S., 2005. The beneficial effect of concurrent task-irrelevant mental activity on temporal attention. *Psychol Sci*, 16(4), p.265.
- Olivers, C.N., van der Stigchel, S. & Hulleman, J., 2007. Spreading the sparing: Against a limited-capacity account of the attentional blink. *Psychol Res*, 71(2), pp.126–139.
- Potter, M.C. et al., 2010. Picture detection in rapid serial visual presentation: Features or identity? *J Exp Psych Human*, 36(6), p.1486.
- Potter, M.C., Nieuwenstein, M. & Strohminger, N., 2008. Whole report versus partial report in RSVP sentences. *J Mem Lang*, 58(4), pp.907–915.
- Sara, S.J., 2009. The locus coeruleus and noradrenergic modulation of cognition. *Nat Rev Neurosci*, 10(3), pp.211–223.
- Serences, J.T. & Yantis, S., 2006. Selective visual attention and perceptual coherence. *Trends in Cognitive Sciences*, 10(1), pp.38–45.
- Shih, S.I., 2008. The attention cascade model and attentional blink. *Cog Psychol*, 56(3), pp.210–236.
- Snider, J., McCall, R. & Franklin, S., 2011. The LIDA framework as a general tool for AGI. *Artificial General Intelligence*, pp.133–142.
- Taatgen, N.A. et al., 2009. Too much control can hurt: A threaded cognition model of the attentional blink. *Cog Psychol*, 59(1), pp.1–29.
- Thompson, K.G., Biscoe, K.L. & Sato, T.R., 2005. Neuronal Basis of Covert Spatial Attention in the Frontal Eye Field. *J Neurosci*, 25(41), pp.9479–9487.
- Warren, C.M. et al., 2009. Target—distractor interference in the attentional blink implicates the locus coeruleus—norepinephrine system. *Psychon Bul Rev*, 16(6), pp.1106–1111.
- Wyble, B., Bowman, H. & Nieuwenstein, M., 2009. The attentional blink provides episodic distinctiveness: Sparing at a cost. *J Exp Psych Human*, 35(3), p.787.