A Skill-based Approach to Modeling the Attentional Blink

Corné Hoekstra (c.hoekstra@rug.nl)
Bernoulli Institute for Mathematics, Computer Science, and Artificial Intelligence
University of Groningen, Groningen, the Netherlands

Sander Martens (s.martens@umcg.nl)
Department of Biomedical Sciences of Cells and Systems
University Medical Center Groningen, University of Groningen, Groningen, the Netherlands

Niels A. Taatgen (n.a.taatgen@rug.nl)
Bernoulli Institute for Mathematics, Computer Science, and Artificial Intelligence
University of Groningen, Groningen, the Netherlands

Abstract

People can often learn a new task very quickly. This suggests that people are able to use skills that they have learned from a previous task, and apply them in the context of a novel task. In this paper we used a modeling approach based on this idea. We created a model of the attentional blink (AB) out of the general skills needed to perform an AB-task. The general skills were acquired from creating separate models of other tasks, in which these same basic skills are used. Those models showed a good fit with reported data, indicating that the basic skills we created are valid. Subsequently, we created the AB model by tying together the basic skills taken from the basic models. The AB model generated the same basic AB effects as reported in the literature. The models produced by the skill-based approach suggest that this is a feasible modeling method, which could lead to more generalizable models. Furthermore, it shed new light on previously difficult to explain findings in the AB literature.

Keywords: Attentional Blink, PRIMs, ACT-R, Skill-based modeling, Cognitive model, Instruction learning

Introduction

Humans have the impressive ability to learn certain relatively simple tasks with minimal instruction and in a very short period of time. The experimental tasks used in (cognitive) psychology are particularly good examples of these types of tasks. Participants have often never encountered these tasks before, yet they are quickly able to work out what to do. This quick learning suggests that people reuse previously learned skills and apply them to new contexts (Salvucci, 2013; Taatgen, Huss, Dickison, & Anderson, 2008). For example, if a task requires a stimulus to be remembered for later recall, people do not have to work out how to remember the stimulus, but they can simply reuse the already learned 'remembering-skill'. It would be unnecessary, in this case, to reinvent the wheel. Learning how to do a new task simply means selecting the appropriate skills, assuming all these skills have already been acquired.

Reusing skills speeds up learning, but it can also have negative side effects that lead to sub-optimal performance even though the cognitive system is, in principle, capable of optimal performance. That is, it is sub-optimal strategy that underlies the impaired performance, not a fundamental information processing limit (e.g., Taatgen, Juvina, Schipper, Borst, & Martens, 2009). One factor underlying the sub-optimal strategy-choice might be the selection of the wrong skills, either because the "right" skill is not available, or because the interpretation of the task cues the wrong skill. A well-known instance of this is the Stroop effect (Stroop, 1935). Because people are so used to reading words, this automatically triggered skill interferes with the task of naming the color of the word. In this case, selecting the 'reading-skill' leads to worse performance. Another instance where this can happen is the attentional blink (AB).

The AB is a well-studied phenomenon in cognitive psychology (Martens & Wyble, 2010). It refers to the finding that the second of two to-be reported targets in a stream of distractors presented at a rate of 100 ms per item is often missed when it is presented within an interval of 200-500 ms after the first target (T1) (Raymond, Shapiro, & Arnell, 1992). Interestingly, the second target (T2) is hardly ever missed if it is presented directly (i.e., within 100 ms) after the first target (lag-1 sparing). This suggests that the cognitive system does possess the processing capacity to identify both targets, but that the chosen strategy prevents the second target from being reported.

The crucial aspect of the strategy that most participants use can be the selection of a sub-optimal skill to consolidate the targets in memory. Many theories of the AB assume that consolidation of T1 into memory underlies the AB. Memory consolidation is thought to be a serial process, meaning that only one consolidation process can occur at a time and that the consolidation has to be completed before a new item can be consolidated. This means that T2 cannot always be consolidated straight away, but sometimes has to wait for the consolidation of T1 to be completed. This leads to the AB when consolidation of T1 has not yet been completed before T2 has disappeared from visual short-term memory. However, such theories all assume that targets are consolidated as separate memory items, whereas in other areas of memory research it is assumed that multiple items are consolidated in a single chunk.

The strongest indication that strategy underlies the AB phenomenon is an experiment by Ferlazzo and colleagues (Ferlazzo, Lucido, Di Nocera, Fagioli, & Sdoia, 2007). In their experiment, participants were instructed to report two target letters (which were always a vowel and a consonant) either separately or as a single syllable. In the latter condition participants did not exhibit an AB. A possible explanation is that the original instruction cues a strategy in which all targets are consolidated separately, while the syllable instruction encourages consolidation of both targets in a single chunk. We will explore this difference by creating two versions of an AB-model that only differ in their consolidation strategy.

To create the model, we have used a novel approach. Instead of creating the model specifically for the AB, we built a model from the general skills that we have constructed as parts of other models. In other words, the AB model only links together existing skills. We chose this approach because it mirrors how participants performing an AB-task work out what to do. They do not start from scratch, but they tie skills they already possess together in such a way that allows them to perform an AB-task.

We created this model in the cognitive architecture PRIMs (Taatgen, 2013, 2014). PRIMs is based on ACT-R (Anderson et al., 2004) and works in a highly comparable way. The architectures of both ACT-R and PRIMs consist of a 'central workspace' and a number of modules capable of performing specific cognitive functions. The modules can communicate (i.e., exchange the results of their cognitive operations) with each other through the central workspace, which is subdivided in buffers. This exchange of information between the modules in PRIMs is controlled in largely the same way as it is in ACT-R. In ACT-R this is done by productions, and in PRIMs it is done by operators, but they have similar functionalities. A crucial difference between ACT-R and PRIMs is that in PRIMs operators are further organized into skills. A skill is a collection of general operators capable of accomplishing a certain goal or processing step. The generalizability of skills makes it possible to use the same skills in models of different experimental tasks. The organization into skills thus allows us to employ a novel approach to constructing cognitive models, placing them in a context of related models, tasks, and skills.

Each skill has a number of variables that are instantiated when a skill is used in the context of a task. It is by this mechanism that we tie together tasks, but also fill in specific values.

We had two main goals in this project. Firstly, we wanted to investigate the feasibility of creating a cognitive model by tying together already existing skills. Secondly, we wanted to create a model of the AB which is capable of capturing most of the effects found in the AB-paradigm, including differences due to instruction.

Method

Instead of creating operators specifically for the attentional blink, we first considered which general skills are required to perform an AB-task and assembled the AB-model from these skills.

Based on previous work and other models of the attentional blink, we identified four basic skills (cognitive processing steps) which had to be performed by our model of the AB. We developed these four skills by first creating models of other tasks which share (some) of these same basic skills. This step was done to get a better idea of what these general skills should be capable of and to test the plausibility of these skills.

First, we will describe the three models that provided the building blocks for the AB-model. The three models are: (1) a visual search model, (2) a model of a simple working memory (SWM) task and (3) a model of a complex working memory (CWM) task. Not all parts of all three models will be used for the AB-model, but all three contain at least one of the four basic skills needed to perform an AB-task.

The first model, the visual search model, is very straightforward. The goal of this model is to find a vowel on a screen filled with other letters. It is composed of three skills. The main search skill processes the current visual item and determines its category through memory retrieval. If it does not match the target category (vowel in this case), it transfers control to another skill which focuses on the next search item. In visual search this is a shift of attention to another item. If it does match the target category, it transfers control to a third skill, in this case a skill that clicks on the target with the mouse. Finally, if it runs out of items to attend to, it transfers control to yet another skill, which is not instantiated in the visual search model. In the AB-model, we will reuse the search skill to find targets, but we will instantiate it differently.

To illustrate, here are the operators that make up the search skill, slightly abbreviated for clarity. In these operators `V`*x* refers to a slot in the visual buffer, `RT`*x* refers to a slot in the retrieval (declarative memory) buffer, and `G`*x* refers to a slot in the goal buffer.

```
operator look-for-target {
     V1 <> nil // if there is a visual input
==>
     *fact-type -> RT1 // build a
     V1 -> RT3 // retrieval request
     nil -> V1 // and clear the input
}

operator keep-looking {
     V1 = nil
     RT2 <> *target-type // if it is not a target
==>
```

```
      *next-stim -> G1  // change to the skill that
                        // selects the next stimulus
}

operator found-target {
      RT2 = *target-type  // if it is a target
==>
      RT3 -> G8 // Store the target in the goal
      *after-found-target -> G1  // and
               // switch to the skill to handle a target
}
```

In these operators, values that are preceded by an asterisk are variables that need to be instantiated for a particular task. For visual search, we instantiate `*fact-type` with `vowel`, `*next-stim` with the `attend-next` skill, and `*after-found-target` with the `click-item` skill.

The second and third basic model are strongly related and provide the final basic skills. Both models deal with working memory tasks which require the participants to remember presented items and, after presentation of the items, recall which items have been seen. Although they both include a consolidation step, they accomplish this step with a different skill. Both build a chunk in working memory, however they differ in the moment of consolidation. The "consolidate-separate" skill, used in the SWM-model, starts consolidation immediately after an item is encountered. In contrast, the "consolidate-chunk" skill, used in the CWM-model, only starts consolidation after all items have been presented. Using these two consolidation skills, we created two versions of the AB-model, a "consolidate-separate" version and a "consolidate-chunk" version.

Finally, these two working memory task models provide the "retrieve" skill and the "respond" skill. The "retrieve" skill retrieves the appropriate consolidated item from memory and the "response" skill gives the appropriate response based on the retrieved item.

The four skills described above form the basic building blocks of both versions of our attentional blink model. To finalize the AB-model, the basic skills were put together in one model and were instantiated to fit the context of an AB-trial. This procedure was the same for both versions of the AB-model. In the AB-model, after presentation of a stimulus, the "search" skill checks, whether this is a target or a distractor. In other words, the `*fact-type` variable is instantiated with `letter`. If the stimulus is a distractor, it is ignored and the model waits for the next stimulus (`*next-stim` is instantiated with `wait`). If the stimulus is a target it switches to the consolidate skill (by instantiating `*after-found-target` with that skill) that moves the stimulus into a working memory slot. The consolidate skill is the source of the attentional blink in our model. Depending on which skill is used to accomplish consolidation, the model either starts consolidating directly after encountering the first target or postpones consolidation until the second target is encountered. If the chunk is consolidated, no other operator can be executed for a period of, on average, 200 ms (the imaginal delay parameter in ACT-R), leading to a possible attentional blink. If consolidation is postponed until the arrival of the second target, no attentional blink will occur at this point and the model will keep performing the task normally. After all stimuli are presented, the model will retrieve the targets that were consolidated on this trial (the "retrieve" skill) and then, after the retrieval, responding to the retrieved items (the "respond" skill).

Results

We compared the behavior of the models with human performance. This was done in order to verify the feasibility of the basic models and to check how well the final AB-model could model the AB phenomenon. The comparisons were made with existing data from the literature.

We did not find suitable data to which we could compare our visual search model. This is likely due to the fact that our visual search model is very simple and does not have any other functionalities besides what is described in the method section. Furthermore, the visual search model was not our primary interest, as it is not responsible for creating the AB.

Firstly, we will discuss the comparison between the SWM-model and human performance. The specific task we modeled required participants to remember a certain number of digits and report them at the end of a stream (Anderson et al., 1998). The critical manipulation in this experiment was that the digits were presented in multiple groups. This grouping was thought to influence chunking of the digits, digits grouped together during presentation would also be grouped together in memory (i.e., chunked together). The findings supported this expectation, such that participants showed longer reaction times during recall for the first item of a group, indicating that the groups were remembered (and recalled) as one chunk. The data from the simple working memory model showed this same pattern in reaction times as reported in Anderson et al. (1998).

As can be seen in Figure 1, the reaction times produced by the model show the same typical pattern as the human participants. This reflects the strategy used by the model (and presumably the participants) of recalling the remembered digits. The digits are stored in chunks of three in memory and this influences how the recall occurs. Firstly, the full chunk containing all three digits is retrieved from memory and, subsequently, the three responses are given without any further memory retrieval. Note however that the model is unable to capture the extra-long reaction times at the start of the recall-phase. These increased reaction times are likely due to processes relating to getting started on a new task, an aspect of the task unrelated to working memory so we chose not to model it at this moment.

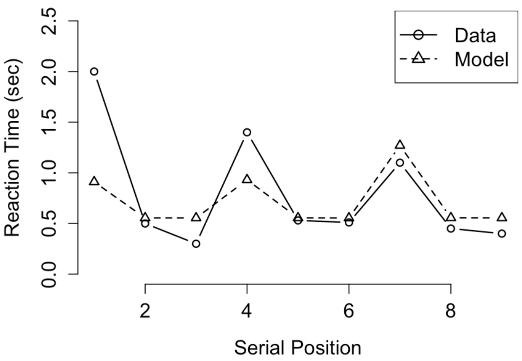

Figure 1: Model fit for reaction times in the SWM-task. Figure depicts the RTs produced by the model (dashed line) and human data (solid line).

Secondly, we will discuss the comparison between the CWM-model and human performance. In the task we modeled, a series of 3, 4, 5, or 6 digits were presented to the model. In between presentation of the digits, the model did a word-decision task in which it had to distinguish between nouns and adjectives. We compared the performance of our model on this task to a similar experimental task (Daily et al., 2001). In this task, participants were instructed to remember a series of digits (also 3 to 6), but here the digits were presented among letters which they were required to read aloud. Both of these tasks have in common that working memory is required to perform the interrupting task (either deciding between a noun or adjective or reading a letter aloud). This demand on working memory makes it impossible for the participants (and the model) to chunk the items in memory.

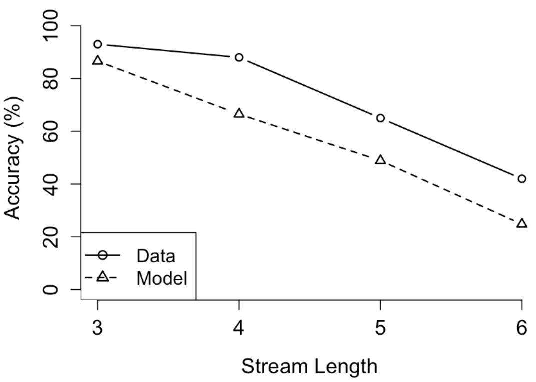

Figure 2: CWM-model fit for accuracy data. The average accuracy as a function of list length for the model (dashed line) and the human data (solid line).

We compared model performance with human performance with respect to accuracy (see Figure 2). Generally, the model shows a good fit to the human data reported by Daily et al. (2001). Both the model and the participants show decreased accuracy when the length of the presented list is longer. This decreased accuracy for longer lists occurs in the model because the presentation of the longer lists takes a longer time to be completed. The longer time required for presentation allows for additional item-decay in memory, leading to reduced accuracy for longer lists. The model, however, generally underestimates accuracy, this is probably due to the model being unable to capture the primacy effect (Murdock, 1962). The primacy effect is often modeled by including a rehearsal mechanism. The fact that we did not include such a mechanism to the model could thus explain the general underestimation of the accuracy.

Finally, we compared our AB-model (which resulted from the combination of the above discussed models) with human AB-performance (see Figure 3). The exact task we modeled is the classic AB-task reported in Chun & Potter, 1995. In this standard version of the AB, participants are instructed to identify two digits within a stream of distracting letters and, at the end of the stream, report which digits they have seen. We modeled this experiment with the version of the AB-model that used the "separate-consolidation" skill. The crucial effect in an AB-task is, unsurprisingly, the attentional blink itself. This refers to the strong performance decrement at lags 2 and 3, which our AB-model nicely captures. In the model, the AB occurs because consolidation of the first target (T1) is still in progress when the second target (T2) is presented. Therefore, T2 cannot be consolidated and will not be reported at the end of the stream. Our model also shows the typical lag-1 sparing effect. This is because consolidation of T1 often has not started at the moment that T2 is presented at lag 1. Therefore, they can both be consolidated into a single chunk and reported at the end of the stream. Finally, the model shows the slow performance increase for the later lags (lag 4 and higher). This is caused by the slow increase of the likelihood that T1 consolidation is finished by the time T2 is presented.

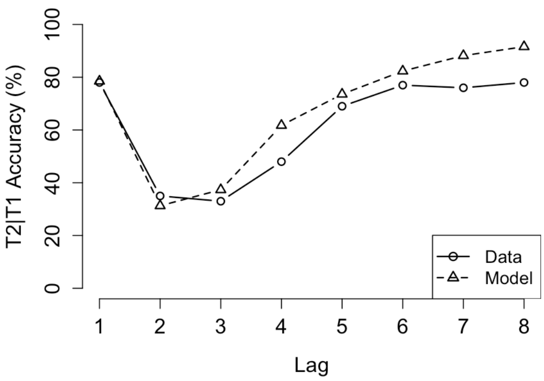

Figure 3: AB-model fit for T2 accuracy. Figure showing T2 accuracy in an AB-task for the model (dashed line) and human data (solid line).

Using the other version of the consolidation skill (the "consolidate-chunk" version) in the AB-model, however, will prompt the model to always try to consolidate both targets into a single chunk, thereby eliminating the AB all together. We compared the performance of the AB-model instantiated this way to the data from the study reporting a reduced AB when participants were instructed in a way that promoted chunking (Ferlazzo, Lucido, Di Nocera, Fagioli, & Sdoia, 2007) (see Figure 4). The model mirrored the general performance level and, crucially, showed no blink. The model, however, shows a slight performance decrease at lag 1. This is caused by our means of simulating noise in the visual system, which meant that occasionally T2 had already disappeared before it was processed fully and therefore it was missed. We do not consider this problematic, because in many AB experiments lag 1 performance is slightly lower than performance on long lags.

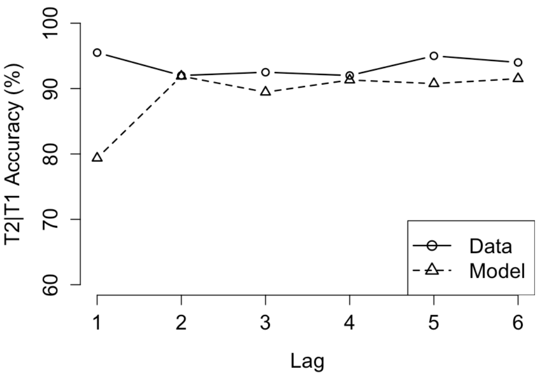

Figure 4: Model fit for the alternative AB model. Figure showing T2 accuracy for the alternative AB model (dashed line) and human data (solid line).

Discussion

Computational models of cognitive psychological phenomena are often able to accurately capture one specific phenomenon, however they are often hard to generalize to other tasks and cognition in general (Anderson et al., 2004). In this paper, we attempted to (partly) bridge this gap by employing a novel approach to building cognitive models, which mirrors the way people approach a new task. People do not consider every task in isolation but they use knowledge gained from the past. That is, they reuse skills learned from doing other tasks and apply them to the (new) task at hand (Salvucci, 2013; Taatgen, 2014). This paper describes our attempt to apply a similar approach. We created a new model of the attentional blink by reusing the models of other cognitive tasks. In short, we had two goals: (1) test the feasibility of the described approach and (2) create a model with the potential to shed new light on differences in AB due to instruction.

The comparisons between our models and human data show that our models are reasonably able to capture human performance. This result demonstrates the basic feasibility of the described modeling approach. It is possible to break a task down into a limited set of skills that are reusable in different tasks. This is an important first step towards creating more generalizable models, because it allows for a method of creating models that are built up from the same building blocks. Using existing building blocks when modeling a new task allows for much more integration of any new model into the already existing collection of models and, more importantly, might better reflect the way people approach a new task.

Note, however, that the devil is in the details. Building a model using this approach can be challenging, especially when it comes to determining how small differences between tasks should best be handled. Such differences make it difficult to use exactly the same operator (and therefore the same skill). Every operator has a condition-checking part (which checks whether this operator should be activated now) and an action-performance part (which actually executes the 'cognitive action' or PRIM). The action-performance part is relatively easy to generalize across tasks, but the condition-checking part is more challenging. Basically, the condition-checking part checks whether the situation matches the predefined situation in which this operator should be executed. This makes it difficult to generalize the condition-checking across tasks since a different task usually also means a different situation. We solved this problem in the models described in this paper by defining the conditions in such a way that they work for all the modeled tasks. This is a workable solution, but it is time-consuming and a better method for condition-checking is needed.

A further limitation to the models described here is that they did not perfectly capture all aspects of human performance. However, we do not see this as a major issue because we did not set out to create complete models of the described experimental paradigms. Instead we aimed to create models of the main findings only because we were merely interested in the skills that are important for the AB. Although there remain limitations and improvements to be made to the skill-based method, we consider it a feasible and promising approach to improve the generalizability of models.

The second goal we set out to achieve in this paper was to create a model of the AB that can account for differences due to instruction. The model described in this paper produces most of the basic effects from the classic AB-task, showing lag-1 sparing, the AB itself and the gradual improvement on later lags. Although there are many additional aspects of the AB reported in the extensive literature which we did not discuss, we believe that the model described here is an adequate first attempt that we will build on in future work.

For now, the fact that the model captured the basic AB-effects implies that these effects, at their core, may be caused by improper selection of skills. At the start of a new task, a participant has to figure out which skills to combine in order to be able to perform the new task. The models we created suggest that there are (at least) two different skills which can take care of the consolidation into working memory aspect of the task: (1) consolidate every presented target into working memory separately (as in the CWM-task) or (2) consolidate targets as larger chunks (as in the SWM-task). The chunk-consolidation skill as used in the SWM-task would be the optimal pick in this situation, two items can be consolidated into one chunk and there would be no negative unexpected effects. This approach is perhaps employed by participants after receiving the experimental instructions from the Ferlazzo et al. (2007) study. However, given that standard AB instructions consider targets as separate items probably prompts most participants to use the separate-consolidate skill from the CWM-task.

The emphasis put on strategy by our model could explain previous findings in the AB literature that have proven difficult to explain. This includes the effect of instructions as well as the existence of non-blinkers (individuals who do not show an AB) (Martens, Munneke, Smid, & Johnson, 2006), and the reduction of AB-magnitude because of training (Choi, Chang, Shibata, Sasaki, & Watanabe, 2012). All these effects could be explained by the type of consolidation strategy. Different instructions might cue the 'correct' consolidation skill, non-blinkers could be more naturally inclined to use the 'correct' chunking strategy compared to blinkers, and the training procedure by Choi and colleagues might nudge participants toward using the same optimal strategy.

To summarize, our novel skill-based approach to cognitive modeling resulted in valid models, created using a more natural and human-like method. In addition, we believe it shows great potential to generate more generalizable and thus more flexible models. Furthermore, it can lead to interesting new perspectives on well-established cognitive phenomena such as the AB. The choice of consolidation-strategy may play an important role in the AB, explaining individual differences as well as instruction and training effects of the AB.

References

Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review*, *111*(4), 1036–1060. https://doi.org/10.1037/0033-295X.111.4.1036

Anderson, J. R., Bothell, D., Lebiere, C., & Matessa, M. (1998). An Integrated Theory of List Memory. *Journal of Memory and Language*, *38*(4), 341–380. https://doi.org/10.1006/jmla.1997.2553

Choi, H., Chang, L.-H., Shibata, K., Sasaki, Y., & Watanabe, T. (2012). Resetting capacity limitations revealed by long-lasting elimination of attentional blink through training. *Proceedings of the National Academy of Sciences of the United States of America*, *109*(30), 12242–12247. https://doi.org/10.1073/pnas.1203972109

Chun, M. M., & Potter, M. C. (1995). A two-stage model for multiple target detection in rapid serial visual presentation. *Journal of Experimental Psychology: Human Perception and Performance*. US: American Psychological Association.

Daily, L. Z., Lovett, M. C., & Reder, L. M. (2001). Modeling individual differences in working memory performance: a source activation account. *Cognitive Science*, *25*(3), 315–353.

Ferlazzo, F., Lucido, S., Di Nocera, F., Fagioli, S., & Sdoia, S. (2007). Switching between goals mediates the attentional blink effect. *Experimental Psychology*, *54*(2), 89–98. https://doi.org/10.1027/1618-3169.54.2.89

Martens, S., Munneke, J., Smid, H., & Johnson, A. (2006). Quick minds don't blink: electrophysiological correlates of individual differences in attentional selection. *Journal of Cognitive Neuroscience*, *18*(9), 1423–1438. https://doi.org/10.1162/jocn.2006.18.9.1423

Martens, S., & Wyble, B. (2010). The attentional blink: past, present, and future of a blind spot in perceptual awareness. *Neuroscience and Biobehavioral Reviews*, *34*(6), 947–957. https://doi.org/10.1016/j.neubiorev.2009.12.005

Murdock, B. B. (1962). The serial position effect of free recall. *Journal of Experimental Psychology*, *64*(5), 482.

Raymond, J. E., Shapiro, K. L., & Arnell, K. M. (1992). Temporary suppression of visual processing in an RSVP task: An attentional blink? *Journal of Experimental Psychology: Human Perception and Performance*. US: American Psychological Association.

Salvucci, D. D. (2013). Integration and reuse in cognitive skill acquisition. *Cognitive Science*, *37*(5), 829–860. https://doi.org/10.1111/cogs.12032

Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology*. US: Psychological Review Company.

Taatgen, N. A. (2013). The nature and transfer of cognitive skills. *Psychological Review*, *120*(3), 439–471. https://doi.org/10.1037/a0033138

Taatgen, N. A. (2014). Between architecture and model: Strategies for cognitive control. *Biologically Inspired Cognitive Architectures*, *8*, 132–139. https://doi.org/10.1016/j.bica.2014.03.010

Taatgen, N. A., Huss, D., Dickison, D., & Anderson, J. R. (2008). The acquisition of robust and flexible cognitive skills. *Journal of Experimental Psychology. General*, *137*(3), 548–565. https://doi.org/10.1037/0096-3445.137.3.548

Taatgen, N. A., Juvina, I., Schipper, M., Borst, J., & Martens, S. (2009). Too much control can hurt: A threaded cognition model of the Attentional Blink. *Cognitive psychology* (Vol. 59).