“Date: 6/19/21: This paper has not been published yet, please do not copy without author’s permission”

How to interpret attentional blink findings? A practical MCMC tool to assess the attentional blink with the eSTST model

Authors:

Shekoofeh Hedayati *¹, Brad Wyble¹

¹Department of Psychology, The Pennsylvania State University

Corresponding author:

***Email: Shokoufeh.hed@gmail.com**

Phone: (+1) 484-767-0102

Data and materials are publicly available at: <https://osf.io/yk3dq/>

Abstract

Previous research has shown that Attentional blink (AB) data can differ between tasks, or subjects and it can be challenging to interpret these differences. In this paper, we provided a ready-to-use tool that allows researchers to map their data onto the *episodic Simultaneous Type, Serial Token (eSTST)* model. This tool uses the Markov Chain Monte Carlo algorithm to find the best set of 3 model parameters to simulate a given AB pattern. These 3 parameters have cognitive interpretations, such that differences in these parameters between different paradigms can be used for inferences about the timing of attentional deployment or the encoding of memory. Additionally, our tool allows for a combination of quantitative fitting against the overall pattern of data points, and qualitative fitting for theoretically important features. We demonstrate the algorithm using several data sets, showing that it can find cognitively interpretable parameter sets for some of them, but fails to find a good fit for one data set. This indicates an explanatory boundary of the eSTST model. Finally, we provide a feature to avoid overfitting of individual data points with high uncertainty, such as in the case of individual participant data.

Keywords: *attentional blink, computational models, MCMC, model fitting, qualitative, eSTST*

Introduction

The attentional blink (AB) has been the subject of intensive study for decades, because it provides important clues about the timing of cognitive processes when attention and memory work together to create memories of multiple memoranda (Akyürek, & Hommel, 2005; Chun & Potter, 1995; Raymond et.al.,1992; Wyble et al. 2009). In AB paradigms, a sequence of items is presented rapidly, and participants are asked to pick out multiple targets presented at varying temporal lags. The AB occurs when two targets are presented sequentially in time, and the second target is missed if it is presented within 200-400 ms after the first target onset. There were hints of this phenomenon in findings nearly 50 years ago (Broadbent & Broadbent, 1987; Schneider & Shiffrin, 1977; Weichselgartner & Sperling, 1987) and it was characterized with a now-classic two-target paradigm in the well-known Raymond Shapiro & Arnell (1992) study. Figure 1 shows an example of an AB paradigm, and its behavioral correlate as a function of relative temporal lag to the first target.

Many studies have explored variations in the shape of the AB by for example, showing how eliminating the mask following the T1 largely reduces the effect (Nieuwenstein, Potter & Theeuwes, 2009), that changing the spatial location of the two targets eliminates sparing (i.e., sparing is higher T2 accuracy at lag 1 compared to the subsequent lags; Visser, Bischof & Di Lollo, 1999) or that there are individual differences in the depth of the AB (Dale & Arnell, 2010; Martens& Valchev, 2009)

These effects are typically described in terms of the data pattern, such as the AB being either deeper, or shallower. This descriptive analysis is useful but leaves the cognitive interpretation ambiguous (Guest & Martin, 2020). What would be more useful from a cognitive theory

perspective is to project these effects onto parameters that are thought to reflect interpretable cognitive variables.

For example, it has been shown that accuracy at lag 8 is greatly reduced for schizophrenic patients, which could be described as a “shallower blink” compared to healthy control groups (Su et al., 2015). This is true in a literal sense but perhaps not helpful from interpreting the underlying cognitive impact of schizophrenia. In other words, observing the shallow blink per se does not reveal what cognitive aspects are involved in the attentional system of schizophrenic patients.

Studying the AB for different types of stimuli and populations is still an ongoing research in the field of psychology and neuroscience (Olatunji, 2021; Zhou et al., 2020). Yet, there has been no attempt to provide a tool that systematically maps different behavioral data of the AB onto cognitively relevant constructs. Having such a tool can reveal, for instance, how individual difference in cognitive variables causes different patterns of AB, or how a task manipulation mediates the AB pattern by tuning related mental constructs. Providing tools of this sort is helpful for connecting empirical investigations to theoretical understandings of cognitive mechanisms and this connection helps to increase the diversity of approaches to cognitive science, which has been shown to be generally beneficial to the general pursuit of good explanations (Devezer, et al., 2019).

In this study, we are providing a tool that uses an existing theory of the AB called eSTST (episodic Simultaneous Type, Serial Token; Wyble, Bowman & Nieuwenstein, 2009) as a generative model to translate behavioral patterns of AB data into a set of cognitively useful parameters. For example, one could provide a set of data points (i.e., accuracy of T1 and T2

across lags) and the tool will fit three key parameters from the model so as to recreate the data pattern in the model's output.

Figure 1. The attentional blink paradigm and its behavioral correlate taken from Chun & Potter (1995).

Note. Each lag represents 100ms of stimulus onset asynchrony (SOA). The accuracy represents the percentage correct of second target in trials where the first target was reported correctly.

The tool described here uses an automated algorithm based on the Markov chain Monte Carlo (MCMC; Gilks & Roberts, 1996) technique to fit AB data to the eSTST model. This fitting process results in a set of parameters each representing a cognitive variable (e.g., target perceivability, the time it takes to encode the information and the latency of attentional deployment).

This is useful for several reasons. First, these three parameters have explanatory cognitive explanations and can be compared across different paradigms or individuals. Second, these parameters can decouple effects that might otherwise be mistakenly confounded (e.g., a shallow blink may arise from either less memory encoding, or a reduced baseline and this model will separate those examples). Finally, a failure to find a good fit is informative about the explanatory limits of eSTST, by showing what patterns can or cannot be explained by it.

Quantitative vs Qualitative Features

The typical approach to fitting data is to run simulations and modify the parameters so as to reduce the average difference between the real data and the simulated data. The error value would reach 0 when the simulation matched the real data exactly. MCMC works to get the error value as close to zero as possible.

One obstacle in using quantitative fitting is that simply minimizing error can lead to solutions that may not capture the quality (or overall shape) of the AB data. In other words, error minimization is generally insensitive to individual qualitative features of the data, as all it does is minimize the error between two curves (i.e., model's output and the behavioral data). For instance, the AB data for a given task might demonstrate a feature of lag-1 sparing (i.e., sparing is higher accuracy at lag 1 compared to lag 2), but a given fit fails to capture that pattern because it finds a fairly good solution that matches very closely on the other data points.

Figure 2 illustrates this example by showing two hypothetical model fits (fit 1 and fit 2) on an AB data. Although fit-1 results in a lower amount of overall error in a quantitative sense, it fails to replicate the lag-1 sparing. On the other hand, fit-2 generates more error in total, but better simulates the overall pattern of data and therefore might be preferred over fit-1. To help ensure that our tool does not miss the forest for the trees in this way, we buttress our MCMC algorithm with a novel approach to encourage solutions that match on particular qualitative patterns that are considered important. For example, if one wishes to find the best fit of the parameters for an AB curve while still ensuring that the deepest point of the blink is at lag 3, this constraint could be included as a *qualitative feature* in the MCMC search (See Pitt et al., (2006) for more discussion on the importance of qualitative methods).

Figure 2. Two hypothetical model fits (Fit 1 and Fit 2) against a behavioral pattern.

Note. For theoretical purposes, fit 2 could be preferred over Fit 1 even though it yields a greater amount of error because it exhibits the quality of lag-1 sparing.

Background

Attentional blink and its qualitative features: The primary characteristic of the AB is missing the second target based on its temporal position relative to the first target. However, examining this phenomenon more closely revealed that the AB often demonstrates ancillary features dependent on the first target accuracy and the memory of the targets' temporal orders (Bowman & Wyble, 2007; Chun & Potter, 1995; Hommel & Akyürek, 2005; Wyble, Bowman & Nieuwenstein, 2009). Figure 3 shows the AB data obtained from Chun & Potter (1995), and its ancillary features. The features we consider important are as follows although other researchers may have different perspectives:

Blink: The lowest accuracy in reporting T2 occurs within 200-400ms after the first target presentation (i.e., normally at lags 2-5 with SOA=100ms; Figure 3, middle panel).

Lag-1 sparing: The blink can often be prevented by presenting the second target within 100 ms after the first target onset. As a result, T2 is reported more accurately at lag 1 compared to the subsequent lags (Figure 3, middle panel)

Swapping: The encoding of both targets at lag 1 often comes at a cost of reporting the targets in the wrong order (Figure 3, right panel). According to the eSTST this occurs because both targets are represented in one episodic memory at lag 1.

Lower T1 accuracy at lag 1: Another cost of encoding two targets at lag 1 is often a small impairment in reporting T1 compared to the subsequent lags (Figure 3, left panel). According to the eSTST this occurs because both targets are represented in one episodic memory.

Figure 3. The attentional blink and its qualitative features

Note. The data were taken from Chun & Potter (1995) with each lag representing 100 ms (i.e., SOA = 100ms). Accuracy of reporting T1 (left), accuracy of reporting T2 given T1 was reported correctly (middle), and the number of times T1 and T2 are reported correctly but in wrong orders (right).

Whether the AB pattern for a given dataset contains any particular feature can be diagnostic as to the internal mechanisms of attention. If an experiment shows a blink, but no lag-1 sparing or swapping there might be some mechanistic explanation behind it in terms of how attention was

deployed during the information processing. Thus, testing the data against the eSTST model can provide further explanation of the attentional dynamic in specific tasks.

It should also be noted that the ancillary features described here are not required in a given fit. A researcher may consider some other features of an AB shape to be theoretically important. On the other hand, she may decide to discard all or some of these ancillary features and assume that the best fitting curve from an error minimization approach is sufficient to explain a given result.

eSTST: The approach described here uses the Episodic Simultaneous Type-Sequential Token model to generate simulations of the T2 and T1. This computational model is a neural network that takes a sequence of inputs indicating when targets appear, for how long and whether they are masked. The model returns an output in terms of which targets were stored in memory and what their remembered temporal order was. This model has been successful in explaining numerous ancillary features of the AB such as reduced accuracy of T1 at lag 1, the time course of swaps, spreading of sparing, repetition blindness (Wyble, Bowman & Nieuwenstein, 2009), and the fact that unmasked targets can cause an AB (Nieuwenstein et al. 2009; Wyble, et al. 2011; Lagroix et al., 2012).

The model simulates the AB as the consequence of a mechanism that encourages the formation of discrete representations in working memory by accumulating neural activity over time. While the encoding of any target is underway, an inhibitory projection (i.e., transient attention) puts pressure on the entry of new targets into the encoding stage of working memory (Figure 4; see Wyble et al., 2009 for more details).

Figure 4. A simplified schematic of the eSTST mechanism adopted from Wyble et al., 2009.

Note. The first target triggers attention and that amplifies processing of the target in working memory. While any target is actively being encoded into working memory, a top-down inhibitory projection from working memory to attention inhibits processing of additional targets, thus causing an attentional blink.

The model contains three critical parameters that mediate the pattern of data, primarily for T2 accuracy, but these also affect T1 accuracy, and order swaps. These parameters are: *Blaster delay*, *Binding threshold* and *Base* and they represent distinct aspects of how attention and working memory interact to guide input into working memory for a sequence of targets. Here is a description of each of the parameters.

Blaster delay: This parameter represents the delay of the onset of attention (i.e., blaster) relative to when it is triggered. A higher value of this parameter would correspond to a stimulus that is only detectable via higher order representations in deeper levels of the visual pathway (e.g., selecting a picture according to its conceptual meaning). This parameter is an integer that can range from 1-10 corresponding to delay values of 20ms- 110ms ($\text{delay (ms)} = (\text{blaster delay}) * 10 + 10$). Increasing this value decreases the ability to report a briefly presented T1 because it means that the T1's duration is nearly expired by the time that attention begins. On the other hand, if a T2 occurs right after the T1, it benefits from a slower attentional latency, since the T1

has already “turned on” attention prior to its onset. In terms of the shape of AB, increasing this parameter from 40ms to 90ms, causes the AB to occur in later lags) and increases the sparing window (Figure 5; bottom-left).

Binding threshold: This parameter represents the encoding threshold in working memory and can range from 1-20. It represents the amount of information that must be encoded about a given stimulus for the memory to be complete. For example, correctly encoding a more complex or more unfamiliar stimulus would correspond with this parameter being larger than for a simpler or highly familiar stimulus. Increasing the binding threshold causes the encoding time of the targets to increase and thereby increases the depth and width of the AB (Figure 5; bottom-middle).

Base: This parameter represents the sensory strength of the target. For example, increasing the contrast of a target would correspond to a larger Base parameter. Generally, increasing the base value elevates the baseline accuracy of T1 and T2 (Figure 5; bottom-right).

These parameters interact with other as well, such the simulated pattern from a given triad of parameter values is not just the linear combination of their individual effects. Thus, one must actually run the simulations at many locations in the parameter space to determine how they affect the AB rather than using simpler techniques such as regression. MCMC provides a tool to explore the space in an efficient way, by moving towards the more effective parameter values for a given set of data.

Figure 5. eSTST simulation output for 121 trials as the three parameter changes.

Note. The top panel shows the original parameters of the Wyble et al., 2009 paper. In the bottom three panels we illustrate the effect of varying the three cognitive parameters that are the subject of the MCMC fitting process.

Parameter Fitting: Our approach uses Monte Carlo Markov Chain technique (MCMC; Gilks & Roberts, 1996) to explore the space of parameters in search of the best simulation of the data using a Metropolis-Hasting process (Hasting, 1970; Metropolis, et al., 1953). To understand how MCMC works, imagine that there is a three-dimensional space (with axes of Blaster delay, Binding Threshold and Base) which the MCMC algorithm explores by trying different combinations to see which ones best simulate the data. The space is too big to try all possible combinations of these parameters, so MCMC describes a process for exploring the space systematically but in a probabilistic manner. There is a series of steps through the space called a “chain” and at each step (or iteration) in this chain, MCMC will generate a new random location which is a short distance away from its current location. The algorithm contrasts how well the simulated data at the new location matches the real data relative to its current location. Then, it will “flip a weighted coin”, deciding whether to jump to the new location or not, where the coin

flip is weighted by the ratio of the quality of the fit (including both qualitative and quantitative aspects) between current and proposed simulation result. If the proposed location is better than the current location, the search process is more likely to jump to that location and vice versa.

In the end, the algorithm is expected to find the best location in the parameter space (i.e., the *best fit*), where best is defined as: the smallest difference between the real data and the simulated data that also satisfies the qualitative expectations. To determine whether we have found the best location(s) in the space, we run multiple copies of the search, with each chain starting at a random different location.

Quantitative fitting: Any fitting process brings the question of *what constitutes a good quantitative fit*. The goodness of fit is typically defined as how well a model simulates a given data with minimal error. For the purposes here, we compute this simply as the average of the squares of the difference between each data point, and the simulation results (mean squared error; MSE). Thus, if the simulation values are exactly equal to the real data, this number is zero, which might seem to be the best possible outcome. There is a subtle issue here which is a trade-off between goodness of fit and the model's generalizability, meaning that more accurate simulation is probably less generalizable to new datasets (Pitt, et al., 2002). Therefore, one should tolerate some error in the final value. There are random variations in a given data, and expecting the model to fit all of those micro-fluctuations precisely is counterproductive. For example, in the AB if Lag 8 accuracy is 2% lower than Lag 7, asking the model to reproduce that exact effect (which eSTST cannot really achieve, because there is no mechanism in that model for making T2 accuracy worse at Lag 8 compared to Lag 7), is going to result in a worse fit to other aspects of the data that are probably more relevant to the AB.

Ultimately there is no single decision rule for determining when a model fit is quantitatively good enough. However qualitative features of an AB curve provide an additional information of approaching the goodness-of-fit dilemma.

Qualitative fitting: Minimizing the mean squared error is a good way to achieve an overall match between the data and the simulation, but sometimes there is a specific theoretical reason underlying a particular feature of the data. The importance of evaluating such features can potentially be problematic when used as a simple checklist for testing a model (see Lee et al. 2019 for discussion), but here we describe a technique for linking such features to theoretically relevant cognitive variables and provide a means to evaluate tradeoffs between the presence of such features vs. overall error within a fitting process. As described above, finding reduced accuracy of T1 at Lag 1 could be important and we would like a model fitting that has this feature, even if that comes at the cost of a slightly worse overall quantitative fit. Whether such qualitative features matter depends on the theoretical question that is being addressed and so we cannot recommend the use of specific features. However, the tool provides a mechanism for encouraging the MCMC algorithm to search for locations that match particular qualitative features like the ones in Figure 3.

In our modified version of MCMC that tries to fit qualitative features (QMCMC), the likelihood of hopping from one location to another in the chain is likely constrained if that step would reduce the number of qualitative features that the model is simulating at the proposed location. In other words, in QMCMC the chain will likely not jump to a spot that hits a lower number of qualitative features even if that location is a better quantitative fit (refer back to Figure 2 for an example of fit that is better quantitatively but worse qualitatively).

The extent to which we would like to enhance the qualitative features over qualitative error can be specified via a weighting parameter (`qual_ratio`), so it is not necessary to use qualitative fitting, and it can be switched off in the simulation by setting this parameter to zero.

Model Convergence: An important aspect of an MCMC search is convergence, which indicates the region(s) of the space that one or more chains have moved towards and then remained at. If a chain converges at a given location, it means that the parameter values remain fairly similar for several hundred iterations, and this means that the algorithm has found a “local minimum” in parameter space. If you imagine the parameter space like a bumpy surface, a local minimum is akin to a bowl, and the chain, once it gets into the bowl, tends to stay there.

In general, convergence is a good thing because it means that the search has found one of the best solutions in the entire parameter space. However, there could be other good solutions as well. This is why running multiple chains is helpful, because if they all converge at the same location, it means that there is probably only one good location in the entire space of parameter values. If different chains converge to different locations, and they’re similar in terms of error, it means that there are multiple parameter values that can simulate the data. If parameters of the chains fail to converge, and yet the fit of the simulation to the data is really good, it means that a large region of the parameter space is similarly good at simulating your data.

However, if a chain fails to converge and the fit is not so good, it means that there are probably no good parameter sets for simulating that data. As a result, convergence of all chains at one location is helpful for making stronger inferences about the importance of particular parameter values on fitting the data.

How to use this tool

We are providing a set of MATLAB scripts that allows one to configure an MCMC search to find the best parameters that will provide a fit to a particular AB dataset using the eSTST model. For example, if one has AB curves from two different conditions, the fitting process could indicate that condition X is a better fit by a longer blaster delay than condition Y and this would allow inference as to what specific cognitive mechanisms are being affected by the experimental manipulation. The tool could also be used to project individual AB curves into the three-parameter space, thus allowing inferences about variation in internal cognitive variables (i.e., interindividual differences in attention latency between two conditions) and those parameters could be regressors for other variables of interest.

Method

We selected four examples of the AB to illustrate fits of the eSTST model; three studies drawn from Nieuwenstein, Potter & Theeuwes (2009) and the classic result by Chun & Potter (1995) that the eSTST was parametrized on (originally by hand, not using MCMC; so, this was a chance to compare human vs machine parameter adjustment). All four data sets are shown in Figure 6. Note that experimental results in Nieuwenstein et al., (2009) did not include the swapping data, so the fitting process involves only T2 and T1 accuracies. All the data and codes for the simulations are publicly available at OSF: <https://osf.io/yk3dq/>

Figure 6. AB data from behavioral tasks to test on the eSTST model

Note. An AB study conducted by Chun & Potter (1995) on top consists of T1 and T2|T1 accuracies and the percentage of swapping. Three AB conditions conducted by Nieuwenstein et al., (2009) on the bottom consists of T1 and T2|T1 accuracy. In T1 masked experiments, T1 was followed by distractors, whereas in T1 unmasked condition, T1 was followed by a blank display until the T2 onset. SOA is the stimulus onset asynchrony.

Task description of each Experiment

Chun & Potter (1995): the RSVP consisted of 16 items, 14 black digit distractors and 2 black letter targets. T2 could appear anywhere from lag 1 to lag 8 from the T1 position with each lag corresponding to 100 ms (SOA=100ms).

Nieuwenstein et al., (2009): In their paper, the authors conducted 4 different experiments to test whether distractors are necessary to elicit the AB. We selected the three conditions presented in the first experiment as they elicited different AB shapes. One condition was a replication of Chun & Potter (1995) with distractors being digits (2-9) and targets being uppercase letters excluding (I, O, W and M) that were presented with SOA=100 (Figure 6, bottom left). The second and third conditions were using a similar paradigm, but interleaving blanks between T1

and T2 (i.e., unmasking T1) instead of digit distractors with SOA=100ms and SOA=58ms (Figure 6, middle and right bottom respectively).

Parameter selection

We selected the fitting parameters in eSTST to be blaster-delay (an integer ranging from 1-10; each unit represent 10ms) binding threshold (ranges from 1-20) and base (ranges from .5 to 1), for they correspond to specific working memory constructs and each could mediate the AB pattern differently as shown in Figure 5.

The experimental conditions such as SOA, blank or distractor and the number of lags were manually entered for each experiment and are fixed by the design of the experiment.

MCMC

For each experiment we ran 4 independent MCMC chains, each performing 1200 iterations. The programming was carried out in MATLAB R2018a on 64 bits Windows 10 operation system. The required toolboxes to run the program have been provided in the instruction sheet located within the MATLAB directory at this link: <https://osf.io/yk3dq/>

Simulations for each experiment search for parameter sets that best simulate particular data patterns for a set of three dependent variables (i.e., T2 accuracy, T1 accuracy and swapping if it is reported). In doing so, each MCMC chain starts from a random location in the parameter space and the behavioral data is the reference that the model should ideally simulate. Then, given the parameter set (x,y,z) corresponding to *blaster delay*, *binding threshold* and *base*, the model generates an output and computes the MSE ($MSE_{current}$) against the behavioral data. A proposed location (x', y', z') drawn from normal distributions centered at (x,y,z) with predefined standard deviations (σ_x , σ_y and σ_z) is proposed and the MSE between the model's output for

the new parameter values and the behavioral reference data is computed ($MSE_{proposed}$). The ratio of error ($\alpha = MSE_{proposed} / MSE_{current}$) is compared to a randomly generated number U that we adjusted via formula 1 (below). In the formula, αSD is the standard deviation of the last 50 α values over the course of the chain. The moving decision (whether the algorithm should accept the proposed location or not) is defined by formula (2).

$$U_{adjusted} = (random\ number\ between\ 0\ and\ 1) (\alpha SD) + (1 - \alpha SD) \quad (1)$$

$$\alpha > U_{adjusted} \quad (2)$$

The algorithm moves the parameter values to the proposed location if the condition stated in formula 2 is satisfied, otherwise the algorithm rejects the proposed location. Using this method, we intended to find parameter sets that best describe each experimental data from a strictly quantitative perspective. This formula is only used for the quantitative version of MCMC, but formula 3 (below) incorporates the qualitative aspects

The first 600 samples at the beginning of each chain (i.e., first half of the chain) were considered as the burn-in period. We found the jumping sizes of $\sigma_{bd} = 0.9$, $\sigma_{bt} = .38$ and $\sigma_b = .01$ for blaster delay, binding threshold and base respectively could achieve an overall average acceptance rate of 35% for the last 600 samples.

Qualitative MCMC (QMCMC)

In some cases, the quantitative fit (i.e., minimizing error) would bring the qualitative features for free, but as we will see this is not always the case. In those cases, we make the jumping less likely to a proposed location in parameters space (x', y', z') that reduces the number of qualitative features relative to the current location (x, y, z). An example of a qualitative feature could be that T1 accuracy should be lower at lag 1 than lag 2. The user is able to specify any number of

qualitative features according to match any expected result. Qualitative features are defined within the MATLAB code using simple logical statements and any number of them can be added. The general moving decision is described in formula (3). Where $sum_proposedQuals$ is the sum of qualitative features at the proposed location and $sum_currentQuals$ is the sum of qualitative features at the current location. The $qual_ratio$ is to enhance the qualitative features. If it is set to zero, the algorithm would perform only MCMC (and formula 3 would essentially be the same as formula 2). By increasing this number, say to 0.1, locations with more qualitative features are encouraged.

$$alpha + (sum_proposedQuals * qual_ratio) > U_{adjusted} + (sum_currentQuals * qual_ratio) \quad (3)$$

During a hypothetical MCMC chain, if there are 5 distinct qualitative features described for a given experimental result, and the current step in the chain hits 4 of those features, the QMCMC process will likely decline to jump to a proposed location that has 3 or fewer features if $qual_ratio$ is greater than zero. Thus, the search process is constrained to gradually increase the number of qualitative features that the simulation satisfies over the course of the chain.

To account for the qualitative constraint, we decreased the jumping sizes to $Sigma_{bd} = 0.6$, $Sigma_{bt} = 0.24$ and $Sigma_b = 0.006$ for blaster delay, binding threshold and base respectively and achieved an acceptance rate of 31% on average for the last 600 samples.

Here we describe an example of how three qualitative features can be defined in the QMCMC search.

Blink feature: This feature is set to 1 if T2|T1 accuracy at last lag is 5% greater than T2|T1 accuracy at lag 3, otherwise it is set to 0.

Lag-1 sparing feature: This feature is set to 1 if T2|T1 accuracy at lag 1 is 5% greater than T2|T1 accuracy at lag 3, otherwise it is set to 0.

Lower T1 accuracy at lag 1 feature: This feature is set to 1 if T1 accuracy at lag 3 is 5% greater than T1 accuracy at lag 1, otherwise it is set to 0.

Swapping feature: This feature is set to 1 if swapping at lag 1 is 5% greater than swapping at lag 3, otherwise it is set to 0.

Note that we used an arbitrary threshold of 5%, even though the actual percentage difference might be higher in some data. We find that choosing this difference to be small gives the algorithm more flexibility to explore within the desired region.

MCMC Results

These results are for the simulations where $\text{qual_ratio}=0$ and we enforced no qualitative feature to guide the MCMC search.

Simulation 1 (Chun & Potter, 1995): Figure 7 shows the simulation result for MCMC samples of the first chain. Each of the outputs is the average of 24 (number of iterations/50) samples drawn from the beginning (bottom-left), middle (bottom-middle) and end of the chain (bottom-right).

As the MCMC chains progress, the average error value of the 4 chains ($M = .06$, $SD = .002$, for the last 600 samples), and also the error within each chain declines (Figure 8).

Figure 9 demonstrates the trajectory of sampled values for each parameter with an acceptance rate of 33.24% ($SD = 1.2\%$) for the last 600 samples for four independent chains. Note that although the proposed values for blaster delay could be non-integers, the eSTST only accepted integers for this parameter.

This trajectory can also be shown via the histograms of MCMC samples for each chain for the last 600 iterations (Figure 10). The final central tendency for each parameter was computed by taking the average of accumulated last 600 samples across all MCMC chains (Figure 10, bottom panel), which was 4 (50 ms; $SD = .21$) for blaster delay, 7.70 ($SD = 1.12$) for binding threshold and .81 ($SD = .02$) for base. The summary of the results for all the simulations is shown in Table 1.

Figure 7. Behavioral data vs. the model's outputs

Note. Average of 24 samples from the beginning, middle and end of the first chain (bottom panel) against the behavioral data (top panel). Each lag corresponds to 100 ms in the experiment and the simulation results. A better fit to the behavioral data is observed as the chain progresses.

Figure 8. Errors within and across the chains at each iteration

Note. Error for each chain as a function of iterations (left panel) and the average error across 4 chains. The error was computed as the mean squared error (MSE) between the model's output and the behavioral data at each iteration.

Figure 9. Trajectory of sampled parameter values.

Note. Sampled parameter values for binding threshold, base and blaster delay at each iteration for simulation 1. Each chain is denoted by a different color. Each parameter starts from a random value within the pre-defined range.

Figure 10. Histogram of the last 600 samples for the chains

Note. Histograms of the last 600 MCMC samples for Blaster-delay, Binding Threshold and Base for each chain for simulation 1. Vertical axis indicates the number of samples and is scaled differently for each parameter. Horizontal axis indicates the range of parameter values. The last row represents the accumulated data for all chains (i.e., sum of chain1-4 for the last 600 samples).

Simulation 2 (Nieuwenstein et al., 2009): Figure 11 (top panel) shows the average of 24 samples from the end of the first chain. Average of the error across all chains is also shown in Figure 11 (bottom panel) with the mean error for the last 600 samples to be .048 ($SD = .01$). The simulations resulted in an acceptance rate of 31.24% ($SD=12.17$) for the last 600 samples. Average of accumulated last 600 samples resulted in the blaster delay of 4 (50; $SD= 1.28$), binding threshold of 3.19 ($SD=1.01$) and base value of .81 ($SD=.058$). The histogram of the last 600 samples for each chain is shown in Figure 12. See the supplementals (Figure S1) for sampled parameters.

Figure 11. Behavioral data vs. the model output and the mean error of the 4 chains

Note. Average of 24 samples from end of the first chain (top-right panel) against the behavioral data (top-left panel) for Simulation 2. Each lag corresponds to 100 ms in the experiment and the simulation results. The bottom panel shows the mean MSE across the 4 chains at each iteration.

Figure 12. Histogram of the last 600 samples for the chains

Note. Histograms of the last 600 MCMC samples for Blaster-delay, Binding Threshold and Base for each chain for simulation 2. Vertical axis indicates the number of samples and is scaled differently for each parameter. Horizontal axis indicates the range of parameter values. The last row represents the accumulated data for all chains (i.e., sum of chain1-4 for the last 600 samples).

Simulation 3 (Nieuwenstein et al., 2009; SOA = 100, unmasked): Figure 13 (top panel) contrasts the behavioral data with the average of 24 samples from the end of the first chain. The average error declines (Figure 13, bottom panel), however the end-of-chain simulation does not resemble the behavioral pattern. The mean error for the last 600 samples was .054 ($SD = .006$). The simulation resulted in an average acceptance rate of 32.24% ($SD = 12.49$) for the last 600 samples. The accumulated samples for the 4 chains are shown in the histogram (Figure 14). The average of the last 600 accepted samples accumulated across chains for blaster delay was 3 (40ms, $SD = 1.24$), for binding threshold was 4.00 ($SD = 2.32$) and for base was .81 ($SD = .058$). (See Supplementals (Figure S2) for parameter samples).

Figure 13. Behavioral data vs. the model output and the mean error of the 4 chains

Note. Average of 24 samples from end of the first chain (top-right panel) against the behavioral data (top-left panel) for Simulation 3. Each lag corresponds to 100 ms in the experiment and the simulation results. The bottom panel shows average error (MSE) across the 4 chains at each iteration.

Figure 14. Histogram of the last 600 samples for the chains

Note. Histograms of the last 600 MCMC samples for Blaster-delay, Binding Threshold and Base for each chain for Simulation 3. Vertical axis indicates the number of samples and is scaled differently for each parameter. Horizontal axis indicates the range of parameter values. The last row represents the accumulated data for all chains (i.e., sum of chain1-4 for the last 600 samples).

Simulation 4 (Nieuwenstein et al., 2009; SOA= 58, unmasked): The behavioral data for this experiment and the simulation from the first chain is shown in Figure 15 (top panel) along with the average error per simulations for the chains (bottom panel). The mean error for the last 600 samples was .14 ($SD = .002$). The results indicated the average acceptance rate of 45.88% ($SD = 4.66$) for the last 600 samples. The chains resulted in the blaster delay of 1 (20ms; $SD = .067$), binding threshold of 6.00 ($SD = 1.34$), and base of .70 ($SD = .034$). See the Supplementals (Figures S3 and S4) for parameter samples and histograms.

Figure 15. Behavioral data vs. the model's output and the mean error of the 4 chains

Note. Average of 24 samples from end of the first chain (top-right panel) against the behavioral data (top-left panel) for Simulation 4. Each lag corresponds to 58 ms in the experiment and the simulation results. The bottom panel shows average error (MSE) across the 4 chains at each iteration.

Discussion

Simulations 1 and 2 generated a reasonable representation of the data by the end of the chains, including the qualitative features. The features that we considered theoretically important are

blink, lag-1 sparing and lower T1 accuracy at lag 1 in simulations 1 and 2, and swapping in simulation 1. The results also showed an average MSE that stabilized at a level below 0.1 (.06, and .05 for simulations 1 and 2 respectively). However, Simulations 3 and 4 outputs do not seem to be a good fit to the behavioral data. In the Nieuwenstein et al., 2009 study with SOA=100 ms and unmasked condition, there is a blink in the data shown to be significant in their paper, however the MCMC failed to simulate the blink as it tries to minimize the error, although the average MSE for the 4 chains was no more than simulation 1 (error = .05). This is an example where MCMC minimizes the error, but does not capture a key qualitative feature in the data. To investigate whether eSTST is capable of matching this feature we carried out a qualitative MCMC (i.e., QMCMC: simulation 3b).

Likewise, simulation 4 does not capture the overall shape of the data of Nieuwenstein et al., 2009 study with SOA=58 ms and unmasked condition. In this case, the aforementioned qualitative features exist in both the data and the simulation's output, however, the temporal lags of the features mismatch. For example, the blink starts earlier in the data than in the simulation's output. To explore whether eSTST is capable of matching these aspects of the data we created a set of qualitative features to constrain the AB to occur at the appropriate time. (i.e., QMCMC: simulation 4b & 4c).

Parameters' interpretation: Accepting simulations 1 and 2 as successful simulations of the empirical data, the resulting parameters can be discussed and interpreted according to the related task. In simulations 1 and 2 the changing parameters resulted in blaster delay = 4, 4; binding threshold = 7.7, 3.19 and base = .81, .81 respectively. It is notable that these two paradigms (Chun & Potter 1995 experiment 1, and Nieuwenstein et al. 2009; experiment 1) were methodologically similar and yet provided AB curves that clearly differed in shape. While it is

not clear what is the cause of the different AB curves between these two ostensibly similar paradigms, the fits to the model provide a cognitive explanation for the largest parameter difference, which is in the binding threshold. The larger binding threshold for Chun & Potter (1995) suggests that subjects in this experiment may have needed to encode more information about the targets than the subjects in Nieuwenstein et al. 2009. There is a large list of minor methodological differences between the two paradigms, including the fonts, the display size, the font size, and the subjects (Cambridge, USA community resident's vs Vrije University students) and a difference of about 15 years. Any of these could have caused the difference. See Grassi et al., (2021) for a demonstration of the difference in quality of a cognitive phenomenon with two different groups of subjects separated by 30 years. Unwinding the specific cause of the difference is difficult, but the MCMC fitting process provides a starting point in mapping such differences to cognitive explanations.

QMCMC results

Simulation 3b (qualitative; Nieuwenstein et al., 2009 with SOA=100, unmasked): We used QMCMC to provide a combined qualitative and quantitative fit to the data to see if we could improve on the results of the purely quantitative fit. The missing feature in the simulation 3 result was the *blink* itself. We included other features as well to prevent the model from generating completely different shapes to the AB curves: *T1 accuracy at lag-1* and the *lag-1 sparing* feature with the 5% threshold we defined earlier. As a result, we prioritized 3 features in the QMCMC algorithm. Figure 16 demonstrates the behavioral data and the average of 24 samples at the end of the chain for the first chain. Average MSE for all the chains ($M = 0.058$, $SD = .003$ for the last 600 samples) and the number of qualitative features met at each step is shown in Figure 17. The simulations resulted in an acceptance rate of 33.98% ($SD = 7.27$) for the last

600 samples. Furthermore, the QMCMC resulted in the blaster delay of 4 (50ms, SD = .79), and base of .86 (SD=.04). Binding threshold, though more variable resulted in 6.50 (SD = 2.33) (See Figure 18 for histograms). See supplementals (Figure S5) for parameter samples.

Figure 16. Behavioral data vs. the model's output

Note. Average of 24 samples from end of the first chain (right panel) against the behavioral data (left panel) for simulation 3b. Each lag corresponds to 100 ms in the experiment and the simulation.

Figure 17. The mean number of qualitative features and the MSE error

Note. The average MSE per iteration across the 4 chains (bottom) and the mean number of qualitative features that are satisfied for each chain (top) for simulation 3b. Note that the QMCMC encouraged the algorithm to go to locations with higher number of qualitative features.

Figure 18. Histogram of the last 600 samples for the chains

Figure 17. Histograms of the last 600 MCMC samples for Blaster-delay, Binding Threshold and Base for each chain for Simulation 3b. Vertical axis indicates the number of samples and is scaled differently for each parameter. Horizontal axis indicates the range of parameter values. The last row represents the accumulated data for all chains (i.e., sum of chain1-4 for the last 600 samples).

Simulation 4b (qualitative; Nieuwenstein et al., 2009 with SOA=58, unmasked): The data and the MCMC simulation 4 both show the qualitative features. However, the mean error showed to be relatively high, and the simulation's output did not quite capture the overall pattern of the data. Looking at the data and the simulation's output more closely indicates that the blink is represented as early as lag 2 in the behavioral data, but is represented at lag 4 in the simulation's output. To see if the model can simulate a blink as early as lag 2, we used QMCMC with customized features such that the *blink* and the *lag-1 sparing* features compared the lag-1 with the lag 2 instead of lag 3. In other words, the *blink* feature is set to 1 if T2/T1 accuracy at *lag 7 (last lag)* is 5% greater than T2/T1 accuracy at *lag 2*, otherwise it is set to 0. Likewise, the *lag 1 sparing* feature is set to 1 if T2/T1 accuracy at *lag 1* is 5% greater than T2/T1 accuracy at *lag 2*,

otherwise it is set to 0. The lower T1 accuracy feature was kept the same (i.e., similar to the previous QMCMC). As a result, there were three features included in the QMCMC. The results show that the simulation output is still quite different from the behavioral data (Figure 19). The average acceptance rate for the last 600 samples was 29.70% ($SD = 10.47$). Also, the error still remains high across the chains ($M = .16$, $SD = .03$) and the three qualitative features were met in regions of exploration towards the end of the chains (Figure 20). See table 1 for the final parameter values.

Figure 19. Behavioral data vs. the model's output

Note. Average of 24 samples from end of the first chain (right panel) against the behavioral data (left panel) for simulation 4b. Each lag corresponds to 58 ms in the experiment and the simulation.

Figure 20. The mean number of qualitative features and the MSE error

Note. The average error (MSE) per iteration across the four chains (bottom) and the mean number of qualitative features that are satisfied for each chain (top) for simulation 4b. Note that the QMCMC encouraged the algorithm to go to locations with higher number of qualitative features.

As it is evident from the simulation result of Figure 18 the final output is still different from the actual data in terms of shape of the blink. In this case, lag 2 for T2 accuracy is 5% lower than lag

1, but maybe this amount is not sufficient to capture the depth of the blink. As a result, we increased the 5% threshold for the qualitative features comparison to 20% with the hope of generating a deeper blink at lag 2 in simulation 4c.

Simulation 4c (qualitative; Nieuwenstein et al., 2009 with SOA=58, unmasked): This simulation was similar to simulation 4b except that we changed the qualitative features threshold to be 20% instead of 5% for the T2 (blink and lag-1 sparing). The output again failed to capture the data pattern (Figure 21). This time we can see that the qualitative features maxed at 1 feature on average with a mean error of .14 (SD = .02) (Figure 22). The acceptance rate was 61.43% (SD = 1.74) for the last 600 samples. By inspecting the features of each chain, we found that there was no region in the space that could meet the blink feature where T2 at lag 2 is 20% smaller than T2 at lag 7. See table 1 for the final parameter values.

Figure 21. Behavioral data vs. the model's output

Note. Average of 24 samples from end of the first chain (right panel) against the behavioral data (left panel) for simulation 4c. Each lag corresponds to 58 ms in the experiment and the simulation.

Figure 22. The mean number of qualitative features and the MSE error

Note. The average error (MSE) per iteration across the 4 chains (bottom) and the mean number of qualitative features that are satisfied for each chain (top) for simulation 4c. Note that the QMCMC encouraged the algorithm to go to locations with higher number of qualitative features.

Discussion

Limiting the search with the qualitative features enabled us to successfully simulate the AB task with 100ms SOA in the blank condition. We also realized that the eSTST model is likely not capable of simulating the AB condition of SOA=58ms with blanks using just these three parameters. In other words, the early onset of blink was not achieved when T1 was followed by a blank with 58ms SOAs. However, enforcing the features conferred us a greater confidence that the model is unable to simulate this particular data. Comparing simulation 3b with simulation 2 (unmasked vs. masked), the greatest difference between the AB curves can be attributed to higher binding threshold in simulation 3b when T1 is unmasked. This means that the rapid processing of the unmasked T1 was compensated by the higher binding threshold to trigger the blink. It also suggests that more information of an unmasked stimulus is encoded as representations were not interrupted by the mask.

Table 1. The final parameter values for the accumulated last 600 samples averaged across the 4 chains

Simulation number	Blaster delay	Binding Threshold	Base	Simulates the data?
1	4 (.21)	7.7 (1.12)	.81 (.02)	Yes
2	4(1.28)	3.19 (1.01)	.81 (.06)	Yes
3	3 (1.24)	4.00 (2.32)	.81 (.06)	No
3b	4 (.79)	6.50 (2.33)	.86 (.04)	Yes
4	1 (.00)	6.00 (1.34)	.70 (.03)	No
4b	1 (.00)	7.47 (1.65)	.79(.08)	No
4c	1 (.00)	6.30 (1.53)	.72 (.04)	No

Note. The numbers in parentheses indicate the standard deviation. Simulations 3b, 4b and 4c all include the qualitative features (i.e., QMCMC), whereas simulations 1,2,3 and 4 are the MCMC with no qualitative features.

One additional tool: probabilistic estimate fits

Accommodating data with high uncertainty: So far, we consider only the mean of the accuracy at each point, placing equal emphasis on each. However, there are some cases in which, due to limitations in sample size, some data points have greater uncertainty than others. In such cases the MCMC algorithm should place less emphasis on fitting data points with higher uncertainty. To accomplish this, instead of comparing the proposed value with the mean values, we modified the algorithm to sample from a distribution centered at the mean for a given data point with a standard deviation equal to the standard deviation of that point. For instance, an individual AB pattern is shown in figure 23 with high variability¹. To perform fitting while allowing for variable uncertainty, the algorithm computes the error between the proposed location and a randomly selected point from a distribution centered at the mean of the point estimate with the standard deviation at that point. As a result, the MCMC search is less constrained by data points with higher uncertainty.

¹ The data was acquired from our collaborators based in China, conducting AB studies. We obtained their consent to show one of the subject's data in this manuscript as an example of a data with high uncertainty.

Figure 23. Behavioral data vs. the model's output and the mean error of the 4 chains

Note. The behavioral data of one subject (top left panel) and the average of 24 samples from end of the 1st chain (top right panel). The bars depicted in the figure represents standard errors. The bottom panel indicates the mean error across 4 chains during the course of the simulation.

General Discussion

Decades of research on AB using various paradigms or populations generated different AB patterns (Dale & Arnell, 2010; Nieuwenstein, Potter & Theeuwes, 2009; Visser, Bischof & Di Lollo, 1999). Inferring what differences in cognitive mechanisms are involved in generating these patterns is critical for interpreting the effect of experimental manipulations. Using the eSTST as a baseline theory of the AB, we provide a tool that enables us to map behavioral data onto meaningful parameters that are associated with cognitive variables of working memory and attention. These parameters can add a cognitive interpretation to the shape of any particular AB.

Additionally, we showed that when key qualitative features of the data are not captured by the simulation output, the QMCMC approach can drive the simulation chain towards solutions that match those results. For instance, we showed that MCMC cannot simulate the blink pattern in the data in simulation 3, but including the blink as a qualitative feature in simulation 3b guided the algorithm to find a solution. We also showed that there might be some behavioral patterns that could not be simulated by the eSTST (simulations 4, 4b and 4c) even when simulations are optimized for qualitative features. In that case we are unable to extract meaningful parameter estimations for the data, and, this informs us about the explanatory limits of the eSTST.

Hence, when a data pattern is simulated by the eSTST model, it shows that not only the data can be explained by the model theory, but also provides useful information about the underlying cognitive mechanisms involved in that particular experimental condition. On the other hand, when qualitative features cannot be simulated by eSTST, despite an effort to match them, we can identify the explanatory limits of the eSTST model. Another potential use for the proposed tool is to investigate different shapes of AB at the individual difference level. We provided a separate MCMC MATLAB file that can be used for data with heterogeneous uncertainty per datapoints,

which is more likely to be observed in data from individuals. The final output parameters can be used to compare subjects behaviorally or be used as regressors in other analyses.

Conclusions

How a certain behavior is generated is the central question at the heart of cognitive psychology, and computational models, with all the limitations they have, are great tools that allow us to look into the internal machinery generating a specific behavior (Guest & Martin, 2020). However, the complexity inherent in these models leaves some of their potential explainability unexplored. By developing this tool, we would like to encourage researchers using both computational and experimental means to work together to test existing models against various kinds of empirical data. This approach provides more insight about how a behavior was mediated in relation to cognitive constructs, and reveals the boundary of these models when they fail to replicate the data.

References

- Akyürek, E. G., & Hommel, B. (2005). Target integration and the attentional blink. *Acta Psychologica*, 119(3), 305-314.
- Broadbent, D. E., & Broadbent, M. H. (1987). From detection to identification: Response to multiple targets in rapid serial visual presentation. *Perception & psychophysics*, 42(2), 105-113.
- 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*, 21(1), 109.
- Coull, J. T., & Nobre, A. C. (1998). Where and when to pay attention: the neural systems for directing attention to spatial locations and to time intervals as revealed by both PET and fMRI. *Journal of Neuroscience*, 18(18), 7426-7435.
- Dale, G., & Arnell, K. M. (2010). Individual differences in dispositional focus of attention predict attentional blink magnitude. *Attention, Perception, & Psychophysics*, 72(3), 602-606.
- Devezer, B., Nardin, L. G., Baumgaertner, B., & Buzbas, E. O. (2019). Scientific discovery in a model-centric framework: Reproducibility, innovation, and epistemic diversity. *PloS one*, 14(5), e0216125.
- Gilks, W. R., & Roberts, G. O. (1996). Strategies for improving MCMC. *Markov chain Monte Carlo in practice*, 6, 89-114.

- Grassi, M., Crotti, C., Giofrè, D., Boedker, I., & Toffalini, E. (2021). Two replications of Raymond, Shapiro, and Arnell (1992), The Attentional Blink. *Behavior research methods*, 53(2), 656-668.
- Guest, O., Martin, A. E. (2020). How computational modeling can force theory building in psychological science. *Perspectives on Psychological Science*. Advance online publication. <https://10.1177/1745691620970585>
- Hastings, W. K. (1970). Monte Carlo sampling methods using Markov chains and their applications.
- Hommel, B., & Akyürek, E. G. (2005). Lag-1 sparing in the attentional blink: Benefits and costs of integrating two events into a single episode. *The Quarterly Journal of Experimental Psychology Section A*, 58(8), 1415-1433.
- Lagroix, H. E., Spalek, T. M., Wyble, B., Jannati, A., & Di Lollo, V. (2012). The root cause of the attentional blink: First-target processing or disruption of input control?. *Attention, Perception, & Psychophysics*, 74(8), 1606-1622.
- Lee, M. D., Criss, A. H., Devezer, B., Donkin, C., Etz, A., Leite, F. P., ... & Vandekerckhove, J. (2019). Robust modeling in cognitive science. *Computational Brain & Behavior*, 2(3), 141-153.
- Martens, S., & Valchev, N. (2009). Individual differences in the attentional blink: The important role of irrelevant information. *Experimental Psychology*, 56(1), 18-26.

Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., & Teller, E. (1953).

Equation of state calculations by fast computing machines. *The journal of chemical physics*, 21(6), 1087-1092.

Olatunji, B. O. (2021). Emotional induced attentional blink in obsessive-compulsive disorder. *Journal of Affective Disorders*, 283, 101-107.

Pitt, M. A., Kim, W., Navarro, D. J., & Myung, J. I. (2006). Global model analysis by parameter space partitioning. *Psychological Review*, 113(1), 57.

Posner, M. I., Snyder, C. R., & Davidson, B. J. (1980). Attention and the detection of signals. *Journal of experimental psychology: General*, 109(2), 160.

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*, 18(3), 849.

Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending and a general theory. *Psychological review*, 84(2), 127.

Su, L., Wyble, B., Zhou, Lq. *et al.* Temporal perception deficits in schizophrenia: integration is the problem, not deployment of attentions. *Sci Rep* 5, 9745 (2015).

Visser, T. A., Bischof, W. F., & Di Lollo, V. (1999). Attentional switching in spatial and nonspatial domains: Evidence from the attentional blink. *Psychological Bulletin*, 125(4), 458.

Weichselgartner, E., & Sperling, G. (1987). Dynamics of automatic and controlled visual attention. *Science*, 238(4828), 778-780.

Wyble, B., Bowman, H., & Nieuwenstein, M. (2009). The attentional blink provides episodic distinctiveness: sparing at a cost. *Journal of experimental psychology: Human perception and performance*, 35(3), 787.

Zhou, L., Zhen, Z., Liu, J., & Zhou, K. (2020). Brain Structure and Functional Connectivity Associated with Individual Differences in the Attentional Blink. *Cerebral Cortex*, 30(12), 6224-6237.