

Trinity College

Trinity College Digital Repository

Senior Theses and Projects

Student Scholarship

Spring 2017

Speed-Accuracy Trade-off in Value-Driven Attentional Capture

Yuxuan Li

Trinity College, Hartford Connecticut, effieli1127@gmail.com

Follow this and additional works at: <https://digitalrepository.trincoll.edu/theses>



Part of the [Cognition and Perception Commons](#)

Recommended Citation

Li, Yuxuan, "Speed-Accuracy Trade-off in Value-Driven Attentional Capture". Senior Theses, Trinity College, Hartford, CT 2017.

Trinity College Digital Repository, <https://digitalrepository.trincoll.edu/theses/746>

Trinity College
HARTFORD CONNECTICUT

Speed-Accuracy Trade-off in Value-Driven Attentional Capture

Yuxuan Li

Trinity College

Spring 2017

Advisor: Professor Michael Grubb

A Thesis submitted in partial fulfillment for the Bachelor of Science Degree in Psychology

Acknowledgements

I would like to first thank Professor Grubb for taking me as a thesis advisee, for helping to shape this research project and guiding me through it. This paper would not have been possible without his support in every aspect. Thanks also to him for supporting me outside of thesis research throughout my senior year. He is my role model in conducting rigorous research. I owe much intellectual debt to him.

I would also like to thank Professor Casserly for being a great advisor during my time at Trinity. I have had countless enlightening conversations with her that shaped and inspired my interest in the study of the mind and the brain. Thanks to her for encouraging me to do a thesis and for the support and advise in completing a psychology major.

Thanks to the ISP students in the perception lab, Ben and Kalsang, for assisting in the follow-up study and continuing this line of research. Thanks to the thesis colloquium for sharing opinions and forming a sense of community. Huge thanks to all professors in the psychology department who put together great lectures and labs for students to dive into topics in psychology. Thanks to my friends who made the college journey a memorable one.

Last but not least, I would like to thank my parents and family for the love and encouragement. My years at Trinity would not have been possible without their support.

Abstract

Attention is traditionally divided into two types: voluntary, goal-directed attention and involuntary, stimulus-driven attention (Corbetta & Shulman, 2002; Theeuwes, 2010). However, seminal work on value-driven attentional capture (VDAC) has shown that stimuli associated with reward during a reward learning phase slowed reaction time (RT) in a test phase even when task-irrelevant and non-salient (Anderson, Laurent, & Yantis, 2011). Slower RT, however, can also indicate a preference for more accurate responses at the expense of speedy ones rather than an attentional shift (Wickelgren, 1977). Performance-contingent reward and a response deadline of 800ms impose additional constraints in the VDAC paradigm: responding too quickly decreases reward likelihood and responding too late drops the reward probability to zero. Thus, to maximize reward, participants must carefully decide when to respond, potentially altering the strategic balancing of speed and accuracy and confounding attentional effects with decisional ones. We replicated the VDAC paradigm to address the influence of different response strategies, directly comparing an experimental group (trial-by-trial reward, $n=24$) and a control group (flat reward, $n=24$). Analysis includes only the baseline trials without previously rewarded distractors in the test phase. Using maximum likelihood estimation, RT distributions were fitted with an exGaussian model containing three parameters: μ , mean of the Gaussian component; σ , standard deviation of the Gaussian component; and τ , rate of the exponential component. We found that RT variability (σ) was significantly greater in the experimental group ($p<0.05$), suggesting that reward learning produced a less stable strategy. Further, RT variability positively correlated with error rate ($r=0.51$, $p<0.001$), reflecting a behavioral cost with greater RT variability. These results

call into question the validity of the baseline trials used in the VDAC paradigm, as reward learning altered the response strategy even after the reward was removed.

Introduction

Consider a simple task. In the collection of numbers shown below, find the number 821.

812 333 775
 618 456 659 639 339 410
 802 529 155 566 676 903
 276 722 697 682 502 257 363
 362 299 973 477 140 354 696
 355 381 648 235 394 765
 570 185 957 329 821 126 741
 767 624 402 839 352 320 624
 579 820 485

That was not crazily hard. Perhaps you spotted it right away. Perhaps you had to spend several seconds before realizing that 821 lies in the lower right corner. However, what other numbers do you remember seeing? Was the number 156 in there?

You definitely “saw” all the numbers. All candidate numbers in the figure were well within your visual field provided that you are reading the current paper at a reasonable distance. Even if you only focused your search in the lower half of the figure - that you never really “looked” at any number in the upper half - the light from these numbers still hit your eyes. The only difference was that they lied in the peripheral, non-central part of your visual field, whereas the numbers that you actively looked at during the search were once in the foveal, central point of your visual field, where visual perception is of high resolution. If you did not spot 821 right away, there was a good chance that you visited some other numbers with your foveal vision. You would have to make a quick judgement that these numbers were not 821. What this indicates is that you had clearly seen, registered, and processed these numbers at some point. However, you

probably do not recall exactly what those numbers were. Now, if you search for 821 again, you would know the location immediately. This number-spotting task resonates with the famous Where's Waldo game, although much simpler.

What was making your success at finding the target number and your failure of forgetting other numbers that you saw was, among other things, attention. Attention is usually thought of as a spotlight among a sea of information, internally or externally. In this case, attention helped us focus on a small set of numbers during the visual search (or a single number, or even a single digit, depending on one's strategy), so that we are not overwhelmed by all the numbers. We direct our attention to a sequence of numbers that lie on the saccadic path, checking along the way whether the numbers we encounter happen to be the target number. If the target is found, we quickly register its location, and it becomes much easier to direct our attention to the location of the target number if we perform the same task a second time. If the target has not been found, we need not attend to what the current numbers really were. The important thing is to confirm that they were not the target number. Strategically speaking, we would not want to waste our cognitive resources on remembering the non-target numbers as they are irrelevant to the task at hand. Thus, attention plays quite a critical role in this process. In the following sections, we will unwrap attention and introduce the core attentional phenomenon relevant to the current study: value-driven attentional capture.

Attention

Before we discuss what attention is, let us consider a fundamental question: why do we need attention? Imagine all perceptual information that is hitting your body at this very moment, we cannot possibly handle all of the information at once. In the previous task, there was no way

for us to perceive all the numbers at one glance, we have to start somewhere. Attention modulates this process so that we receive and handle incoming information in a proper sequential manner. Attention not only operates on external information, but also comes into play when retrieving internal information. The spotlight analogy is also applicable in this case, for attention helps to look for the relevant piece of memory among all your memory storage.

Attention is a concept far too familiar to everyone. We pay attention to lecture slides. We are able to spot our best friend among other people. We fail to concentrate on doing work when music is loud next door. We effortlessly pickup our names among noises from cars or other conversations. In all these situations, attention is involved. As William James put it in his classic book *The Principles of Psychology* (1890):

“Everyone knows what attention is.”

However, the concept of attention opens up a door of countless questions. For example, the glossary of psychological terms on the American Psychological Association website lists one definition, that attention is “a state of focused awareness on a subset of the available perceptual information” (Gerrig, Zimbardo, Campbell, Cumming, & Wilkes, 2011). It captures the basic characteristics of attention. But if attention is viewed as a state of awareness, for how long does attention last? Is it static or dynamic? What is its relationship with consciousness? How much freedom does an agent possess to manipulate it? When does it take place along the process of sensation to perception? Is it part of perception, or a separate mechanism? How does it influence memory? Of all, how exactly does it work?

Indeed, as Tsotsos (2011) pointed out, to this date, we have not made much progress in defining attention. Broadly, attention modulates perception by allowing the agent to place priority on certain sensory information over others. Thus, in its simplest form, we could define

attention as selective prioritization. Attention not only affects perceptual performance, but also the state of the neurons in the brain (for a review, see Carrasco, 2011). To get a concrete sense of what attention is, it is helpful to visit some basic properties of attention and different types of attention. These distinct types of attention serve as a broad framework that relates subfields of attention research. In the sections below, a selection of important topics in the realm of attention is discussed. As the current study is concerned with visual attention, many examples and concepts involve visual attention. However, examples involving other sensory modalities are included when are applicable.

Selectivity and limited capacity

Two of the basic characteristics of attention, selectivity and limited capacity, are evident from the number-spotting task mentioned above. We selectively attend to a subset of the available information (numbers), and there is only a limited amount of information that we can attend to at a given time (judging a few numbers at once, but not all). In particular, the capacity of attention is restricted by the amount of high-quality sensory input, cognitive resources, and mental processing power. Selectivity and limited capacity are in a sense two sides of the same coin (Pashler, 1998). Selectivity emphasizes the constructive, positive nature of perceptual process, whereas limited capacity emphasizes the passive side and the incapability to absorb too much information at once. Selectivity in auditory perception was studied in the 1950s. Subjects repeated a spoken message as it was heard while simultaneously presented with another unrelated message. The task was found to be easy: participants successfully tuned their attention towards the target message, although in some cases meaningful contents of the unattended message such as their own names were also registered (Cherry, 1953; Moray, 1959). “Selective

listening” was coined to reflect that attention filters auditory information. In visual perception, a classic selective looking study showed that people were able to attend to one of the two distinct events shown in superimposed videos and played synchronously (Neisser & Becklen, 1975). Just as the back side of the coin, the well-known gorilla experiment showed that people failed to notice a fairly odd event (a woman in a gorilla costume walking by the scene) when they were attending to a task on other contents of the same video (counting the number of passes made by a team playing basketball, Simons & Chabris, 1999). From these experiments and results, it is evident that selectivity and limited capacity are the core properties of attention. When exactly does selectivity come into play during the process of perception is a central argument among theories of attention.

Overt and covert attention

In the number-spotting task, perhaps one of the most common strategies is to actively shift your gaze at different numbers or different areas in the figure. To illustrate with a simpler example, consider the following two numbers:

238

946

To perceive accurately what these two numbers are, you would probably need to shift your gaze and perceive them sequentially. It might be possible to “see” the two numbers at once, by placing your gaze at the mid point. However, it is extremely hard to perceive what the numbers really are, because they fall in your peripheral visual field where information is received at low quality. If the numbers were closer together:

238 946


then it is much easier to identify them at once without moving your eyes. You might place your gaze at any point near the two numbers, and perceive accurately every single digit of each number without moving your eyes. This point of gaze need not be the mid point.

When attention is paired with observable movement in the head or the eyes, it is called overt attention. In the first example where the two numbers are placed at the far ends of both sides of the page, we have to deploy overt attention to permit accurate perception. On the contrary, when attentional shift is not lined up with head or eye movements, covert attention is at work. In the second example where the two numbers are placed closer together, we are able to accurately perceive them by deploying covert attention. The distinction between overt and covert attention can be traced back to von Helmholtz (1867/2005), who noted that in some cases attention can be “entirely independent of the position and accommodation of the eyes...or on the organ of vision”. Both types of attention are widely used in daily activities. For example, when you search for keys, you are actively allocating your attention towards potential places in the external environment. However, you might notice that your cat ran by the room even when your eyes were still fixated on the table trying to find the keys. In some cases, covert attention elicit overt attention. When you noticed your cat running by, you might begin moving your head and your eyes, but your cat might be long gone before you fixate your gaze in that direction. As pointed out in Carrasco (2011), covert attention is critical in competitive situations such as sport activities.

Endogenous and exogenous attention

Another critical distinction of attention is based on the nature of the driving source that initialize the onset of attention. As early as in *The Principles of Psychology* (James, 1890), William James suggested that there are two types of attention: one is passive, reflexive, and involuntary, while the other is active and voluntary. The former is often termed involuntary, exogenous, transient, bottom-up, or stimulus-driven attention. As its names suggest, this type of attention is driven by properties of external stimuli. The latter is often termed (as opposite to the names of exogenous attention) voluntary, endogenous, sustained, top-down, or goal-directed attention. The driving force of this type of attention comes from the agent's own motivation, according to some goal, that directs attention in the search of the outside world. In this paper, the terms exogenous and endogenous attention will be used.

To better and more concretely understand the differences between exogenous and endogenous attention, we will revisit the number-spotting task. The original task clearly would elicit endogenous attention. Your goal was to search for the target number 821, and your search path reflected your attentional shift. When 821 was found, your search terminated, and hence the attention span relevant to the task terminated. If the task remained the same (finding 821), but the figure below is presented instead:



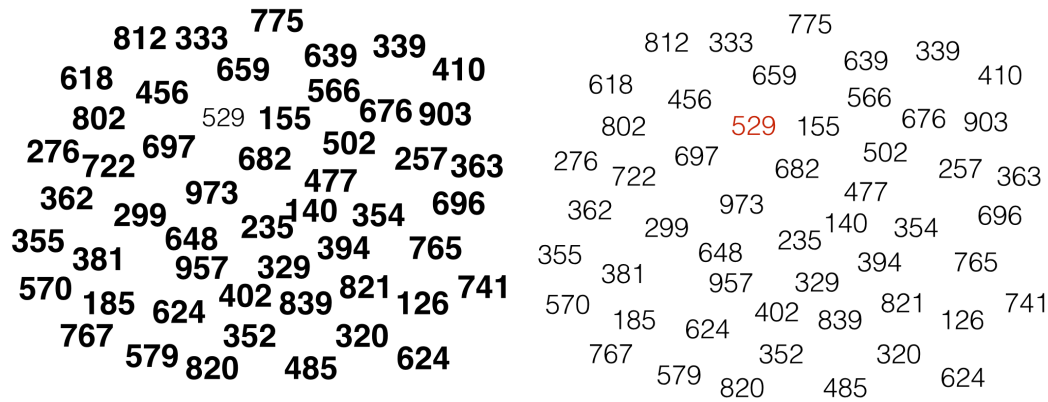
You probably could not neglect the number 529 that certainly stood out among the other numbers. The fact that the number 529 was big and boldfaced made it the most salient among all the other numbers, thus it would automatically capture our attention, or distract us. In this case, 529 would elicit exogenous attention. Further, we say that directing our attention to 529 is task-irrelevant as it is unnecessary nor contributing to our goal.

There are a number of other important differences between endogenous and exogenous attention (for review, see Carrasco, 2011). First, the duration of endogenous attention is, in most cases, longer than that of exogenous attention. Endogenous peaks at about 300ms after its onset, whereas exogenous attention peaks at about 100-120ms. The terms sustained attention and transient attention reflects this duration difference. Second, the degree of automaticity differs. Expectation can aid the allocation of endogenous attention, but no expectation is involved in exogenous attention by its definition. In addition, it is extremely hard to ignore exogenous cues even when we know they are irrelevant. This is perhaps also evident from our daily lives.


Sources of exogenous attention

Saliency

We briefly visited the concept of saliency in the last section. We know that if a stimulus is the biggest among the rest of the stimuli, it appears as the most salient and in turn draws our exogenous attention. However, saliency need not be based on the biggest stimulus. For example, it can be the smallest among all, or even of the same size, but in a different color:



It should be noted that for a particular stimulus to stand out from its environment, having distinct characteristics might not be sufficient. In the following figure, 529 was colored in gray, but it does not make 529 particularly salient among others in black:



Therefore, saliency is based on significantly distinct features that differentiate the stimuli from the rest of the objects in the environment. It should be noted that saliency in the real world is far more complicated than the simple example shown here, as the world is full of different objects. Any particular object will have more features that differ with the subsets of the objects in the neighboring environment. In psychophysics experiments, we are able to discover general principles of attention in visual search by manipulating the features of the target object and the distractor objects.

Learned associations

Another source of exogenous attention is learned associations. We can spot our friends at distance even though they are not physically more salient compared to other people (for instance, significantly taller, shorter, or wearing more colorful outfits). This is because we have interacted with our friends more than with other people in the scene, and our past experience made our friends “stand out” from the environment. A number-based example would be the following:


615 274 821

Because you have been exposed to 821 multiple times from the time you began reading this paper, 821 will most likely elicit your exogenous attention. This means that a stimulus need not be physically salient to draw our exogenous attention. In a sense, learned associations alter the saliency of the stimuli, so that although the stimuli are not physically the most salient, they are semantically more salient.

The idea of learned associations between stimuli stems from the famous example of Pavlov’s dog (Pavlov, 1927). Back in the time, researchers in Pavlov’s lab discovered that their experiment dog started to salivate at the sound of the bell before the food even came in. The dog had learned, after some repetitive experience, to associate the bell with food. In classical conditioning terms, we say that the food was the original unconditioned stimuli, and salivation was the original unconditioned response. After learned associations, the bell became a conditioned stimulus, whereas salivation upon hearing the bell became a conditioned response. Based on similar mechanisms, we acquire new associations between events and stimuli. These learned associations, or our past experience, will in turn modulate how we perceive the world in the future.


Value-driven attentional capture

Under exogenous attention driven by learned associations is a phenomenon called reward-related, or value-driven, attentional capture. The idea is that, a stimulus that was previously paired with value can capture our attention even when it is not otherwise more salient than other stimuli in the environment. This phenomenon was coined value-driven attentional capture (VDAC). Anderson, Laurent, & Yantis (2011) introduced a seminal experimental paradigm to show that this phenomenon manifests a distinct source of exogenous attention. In their study, 26 participants first completed a reward learning phase. They were shown visual search arrays of the following:



The targets were red or green circles. The task was to report whether the line segment inside the target circle was oriented horizontally or vertically. A response deadline of 800ms was imposed. Critically, one of the target colors was associated with high reward value, the other was associated with low reward value. For a given trial that contains a high value target, participants have a high probability of earning a high reward if the response was correct. For a given trial that contains a low value target, participants have high probability of earning a low reward. Through the reward learning phase, participants learned to associated reward with red or green circles.

Following the reward learning phase is a test phase, in which participants were shown visual search displays like the following:



The targets in the test phase are unique shapes among the rest of the objects. Specifically, the target might either be a diamond among circles or a circle among diamonds. The task was to report the line orientation inside the target shape within 1200ms. Critically, colors were no longer relevant in this phase. Half of the trials in this phase contain no previously rewarded distractors. The other half of the trials contain either a high value distractor or a low value distractor.

It turned out that in the test phase, average reaction time (RT) was the longest in trials with a high value distractor, followed by trials with a low value distractor, and trials with no valued distractors had the shortest RT. This slowing of RT seemed to show that the valued distractors had an effect on attentional shift. As a follow-up control experiment, 10 participants underwent the exact task without getting performance-dependent reward. The results suggested that when performance-dependent reward was removed, mean RT did not differ in the three types of test phase trials: no red/green distractor present, red distractor present, and green distractor present.

Speed-accuracy trade-off

Performance and reaction time in Anderson et. al (2011)

A particular interesting aspect of data in Anderson et. al (2011) concerns the dynamics between performance and RT. Recall that the first experiment involved 26 participants. Each participant was rewarded based on their trial-by-trial performance. On the contrary, the control experiment involved 10 participants completing the exact same experiment except using a flat-rate reward mechanism. The following tables from Anderson et. al (2011) show the difference between the data from these separate experiments (Figure 1):

Table 1. Mean response time (in milliseconds) and error rate, respectively, in the test phase of the experiments in which reward was delivered for each of three training conditions: long training (1,008 trials) with low and high reward of 1¢ or 5¢ per trial, brief training (240 trials) with rewards of 2¢ and 10¢ per trial, and brief training followed by a delay of 4–21 d

Training phase	Distractor condition in the test phase		
	None	Low value	High value
1,008 trials	665 (2.8)	673 (2.8)	681 (2.6)
	0.11 (0.004)	0.10 (0.004)	0.11 (0.004)
240 trials	667 (2.0)	675 (3.0)	682 (2.9)
	0.12 (0.005)	0.12 (0.006)	0.12 (0.006)
4–21 d ago	614 (1.8)	624 (2.7)	630 (3.3)
	0.06 (0.004)	0.07 (0.006)	0.08 (0.005)

The error terms, in parentheses, reflect the within-subjects SEM.

Table 2. Mean response time (in milliseconds) and error rate, respectively, in the test phase of the experiments in which no reward was delivered

Training phase	Distractor condition in the test phase		
	None	Red	Green
None	698 (4.1)	696 (4.7)	700 (3.4)
	0.13 (0.004)	0.13 (0.006)	0.14 (0.006)
1,008 trials (unrewarded)	602 (3.9)	606 (2.1)	593 (3.9)
	0.14 (0.004)	0.17 (0.006)	0.15 (0.005)

The error terms, in parentheses, reflect the within-subjects SEM.

Figure 1. Tables are from Anderson et. al (2011). Table 1, Experiment 1, 1008 trials. Table 2, Experiment 3 - control experiment, 1008 trials, unrewarded.

An interesting observation is that, the average RT of the test phase in the first experiment was longer than that of the control experiment, and the average error rate of the test phase in the first experiment was lower than that in the control experiment. In other words, participants who received performance-dependent reward in the training phase seemed to have respond slower but at the same time more accurately than participants who did not receive performance-dependent reward. This led to the idea of a speed-accuracy tradeoff account, which will be explained in greater details in the next section. We can imagine that, for any task, if we try to answer quickly and rush through it, we are more likely to make mistakes, and our accuracy would decline. If we slow down the response, we have more time to process the information and make a more

confident decision, thus yielding a higher accuracy. The dynamics between speed and accuracy have important implications on interpreting RT and performance results.

Speed-accuracy trade-off curve

Wickelgren once made a strong claim, that “speed-accuracy tradeoff method is so superior to the traditional reaction time method, that many psychologists...ought, in many instances, to do speed-accuracy tradeoff studies instead of reaction time studies” (Wickelgren, 1977). At the heart of the claim lies the fact that mean RT and mean accuracy together only represents a single point on what is called a speed-accuracy trade-off curve. A speed-accuracy trade-off curve captures significantly more information, and indeed also the dynamics of performance and RT, of a particular task. A speed-accuracy trade-off curve can be obtained by running the same task several times with multiple response deadlines and plotting performance against processing time. Figure 2 shows a typical speed-accuracy trade-off curve.




Figure 2. Speed-accuracy trade-off curve. Blue points represent hypothetical performance data collected at different time points. The dotted line is the best-fitting line.

There are four important characteristics of a speed-accuracy trade-off curve. First, responses immediately after the onset of the stimuli for a period of time will yield an accuracy at

chance (50%), as relaying the visual information in the brain and making a response require a processing overhead. In the hypothetical example shown in Figure 2, the processing overhead is about 0.4s. After this intercept between the speed-accuracy trade-off curve and the horizontal axis, we start to observe benefit from allowing more time to respond as the accuracy climbs up. This gain in performance with respect to time can be characterized by the rate of increase. Lastly, the curve asymptotes. The particular time point at which the curve asymptotes is considered to be long enough that the agent would have reached its maximum performance. If the task is simple, then the performance will likely asymptotes at 100% accuracy, or highest possible accuracy defined in alternative ways. If the task is relatively hard, then the asymptote might not be 100% correct responses or highest possible accuracy, as there will be a limit of maximum performance depending on the agent's ability even if no time restriction is imposed.

As noted in Wickelgren (1977), speed-accuracy trade-off curves follow an exponential form. One such model is characterized using the equation below;

$$Accuracy(t) = \lambda(1 - \exp[-\beta(t - \sigma)])$$

where t is the processing time, λ is the asymptote, β is the rate of increase, and σ is the intercept.

Reward and time constraint

Reward is a critical component in the VDAC experiment. In this section, we will disentangle the relationship between reward, time constraint, and speed-accuracy trade-off.

As mentioned earlier, speed-accuracy trade-off curve exists because accuracy tends to be higher given more processing time. Assume now that for each trial, a reward is given for a correct response. Naturally, higher accuracy would imply larger probability of obtaining the reward for the current trial. Thus, the corresponding probability of reward curve will exhibit

similar trend compared to the speed-accuracy trade-off curve (Figure 3). Note that, the base probability of obtaining the reward is also at chance level, with probability $P=0.5$. This directly follows the base chance-level accuracy.




Figure 3. Reward probability curve associated with a typical speed-accuracy trade-off curve.

The dynamics of reward probability become more complicated as we add time constraints. Specifically, responding too quickly will result in earning reward at chance level. Spending more time before responding will yield higher probability of reward, but any response that exceeds the given time constraint will immediately result in zero probability of reward even if the response is correct (Figure 4). Thus, for participants who receive performance-dependent reward, in order to maximize their reward, they face a fairly complicated decision problem. On the surface, they are judging the line orientation in the target shape. Implicitly, they must find the

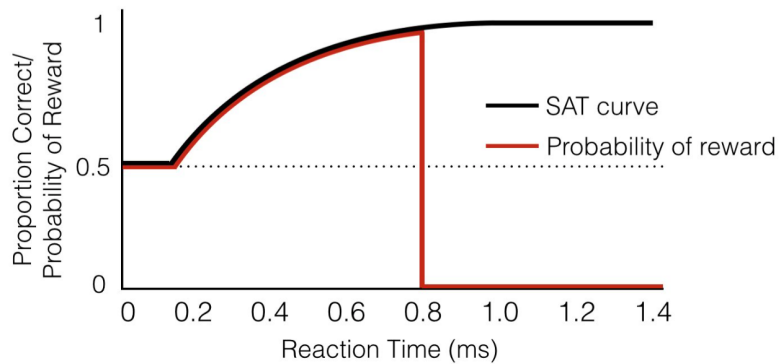


Figure 4. Reward probability curve and speed-accuracy trade-off curve in the VDAC paradigm.

best response time within the given time constraint. That is, they must find the best balancing point between speed and accuracy.

Clearly, these dynamics might interfere with RT in test phase trials. During the training phase, participants might develop different strategies depending on the presence of performance-related reward. It is possible that participants who receive performance-related reward tend to wait closer to the response deadline so that the probability of getting the reward is maximized. On the contrary, there is no incentive for participants who receive flat-rate reward to wait longer in the control experiment. If this strategic difference carries over to the test phase, then it confounds the effect of attentional shift following reward learning. We argue that it is crucial to rule out the possibility that slower RT in trials with previously defined targets in the test phase is driven by decisional factors.

The current study

In the studies that investigated value-driven attentional capture or related topics, the general consent is that top-down strategy is not likely involved in the process (Anderson et. al, 2011; Theeuwes & Belopolsky, 2012). However, the presence of reward learning might interfere with RT and performance through a strategic balance. As discussed in the last section, despite the explicit decision of what to respond, an implicit decision concerning when to respond must be considered. This means that slower RT might not fully reflect attentional capture, but to an extent different strategies involved. It is possible that the presence of reward influenced the strategies developed during the training phase. Because the effect of speed-accuracy trade-off directly impacts performance and RT, it is necessary to rule out the possibility that any difference in RT in Anderson et. al (2011) was due to the underlying dynamics during decision making.

This was not discussed in Anderson et. al (2011). In addition, the control experiment and the main experiment in Anderson et. al (2011) did not have a matching sample size, and were not conducted at the same time as a direct comparison. Therefore, two main goals of the current study are: 1) replicate the value-driven attentional capture paradigm; 2) investigate whether reward learning alters the strategic balancing of speed against accuracy in the subsequent, unrewarded test phase.

To our knowledge, we are the first study to investigate the effect of speed-accuracy trade-off in the phenomenon of value-driven attentional capture, with a direct comparison between an experimental group and a control group.

Methods

Overview

We based the experimental design on Experiment 3 introduced in Anderson (2011), with the following exceptions. We excluded the questionnaire and the evaluation of change detection. There was no follow-up experimental session aside from the main experiment. We added a control group to directly compare the results of the experimental group and the control group. We used a Logitech F310 gaming controller to collect responses instead of using keys on a keyboard. We also collected eye movement data. The experimental procedure was approved by the Institutional Review Board at Trinity College.

Participants

48 adults (18-24 years) participated in the experiment. 71% of the participants were female. The eligibility criteria included: 1) over 18-years of age; 2) normal vision or corrected-to-normal vision, and 3) not colorblind. All participants confirmed eligibility before the psychophysical sessions.

Procedure

Participants were randomly assigned to the experimental group or the control group. Each participant completed a single session of two parts in the Perception lab at Trinity College. The experiment was programmed using PsychoPy running on a 3.0GHz Dual-Core Intel Core i7 Mac Mini (Peirce, 2007). The stimuli were shown on 27.0" LED-Lit Dell Gaming Monitor (S2716DG), with dim environmental lighting. Participants were approximately 96 cm from the

screen. Experiment 1 served as the training phase and Experiment 2 served as the test phase (see Figure 5 for an overview of the experimental procedure). Both experiments started with instruction slides. Participants completed 2 blocks of practice trials (24 trials per block for the training phase, 10 trials per block for the test phase), and chose when to begin a block of trials. Following the practice trials, the experimenter set up the eye tracking device. A single session lasted approximately an hour, and participants were able choose to take a break between the two experiments. At the end of the session, participants were reimbursed with the corresponding reward they earned.

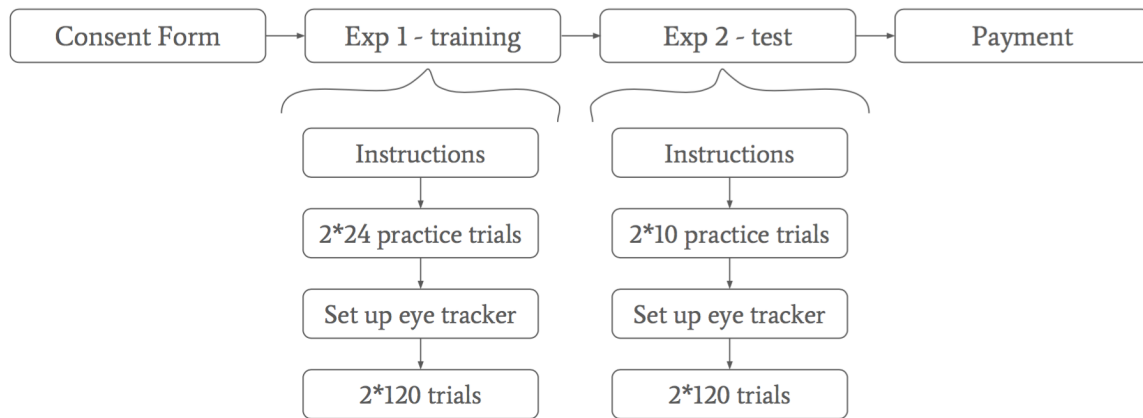


Figure 5. Flow of the experimental procedure.

Training phase

Each trial in the training phase began with a period of eye fixation randomly selected from 400ms, 500ms, or 600ms (Figure 6A). A white cross at the center of the screen served as the fixation cue (vertical segment, 1° visual angle; horizontal segment, 1° visual angle). The cross remained in the center during the subsequent visual search array. The visual search array consisted of six circles positioned as the vertices of an imaginary hexagon (Figure 6C). Each

circle had a radius of 1.15° and was placed 5° from the center of the screen. The target colors were defined as red and green. The distractor colors were cyan, blue, pink, orange, white and yellow. In each trial, all six circles were assigned different colors, and exactly one circle was assigned the target color red or green. This target circle appeared randomly in the six possible locations. Within each circle, there was a white line of length 1.75° . In the target circle, the line segment was either vertical or horizontal. In the distractor circles, the line segment was rotated randomly by 45° clockwise or counter-clockwise. Participants were asked to judge the orientation of the line segment in the target circles. Responses were collected using the L1 and R1 buttons on the controller, which indicated vertical orientation and horizontal orientation, respectively. Participants were allowed an 800ms time window to respond. Responses made after the time constraint were marked missed responses.

In the training phase, the experimental group received rewards based on their trial-by-trial performance. Monetary reward associated with correct responses on trials with a red target and a green target differed in the training phase. One of the target colors was defined as the high-value target and the other served as the low-value target. For trials that contained the high-value target, there was a 0.8 probability of receiving a high reward (\$0.10) and a 0.2 probability of receiving a low reward (\$0.02), if the response was correct. For trials that contained the low-value target, there was a 0.8 probability of receiving a low reward (\$0.02) and a 0.2 probability of receiving a high reward (\$0.10), if the response was correct. This color-value assignment was counter-balanced across participants in the experimental group.

On each trial, if a response was made or if the time constraint was reached, the search array was removed and visual feedback on the performance of the current trial was displayed

(Figure 6E). If a correct response was made before the deadline, participants were shown their reward of the current trial along with the accrued reward up until the current trial. If the response was made before the deadline but was incorrect, participants were informed that the response was incorrect. If no response was made before the deadline, participants were informed that their response was too slow.

The control group received \$13.25 (reported average in Anderson et. al, 2011) for completion of the first experimental phase. Their visual feedback indicated only correct, incorrect, or missed responses.

Test phase

The experimental procedure of the test phase resembled that of the training phase (Figure 6B). Trials in the test phase differed from that of the training phase in following ways. The search array consisted of a unique shape among the five other shapes. There were two such cases: a diamond among five circles, or a circle among five diamonds (Figure 6D). The target in the test phase was defined as the unique shape. Participants were instructed that colors are irrelevant in the test phase. On half of the trials, one of the non-target shapes appeared in red or green randomly among the six possible locations. The task was to report the orientation of the line inside the unique shape. Participants were allowed a time window of 1200ms to respond.

Both the experimental group and the control group were reimbursed \$5 for completing the test phase. Visual feedback for both groups only indicated correct, incorrect, or missed responses.

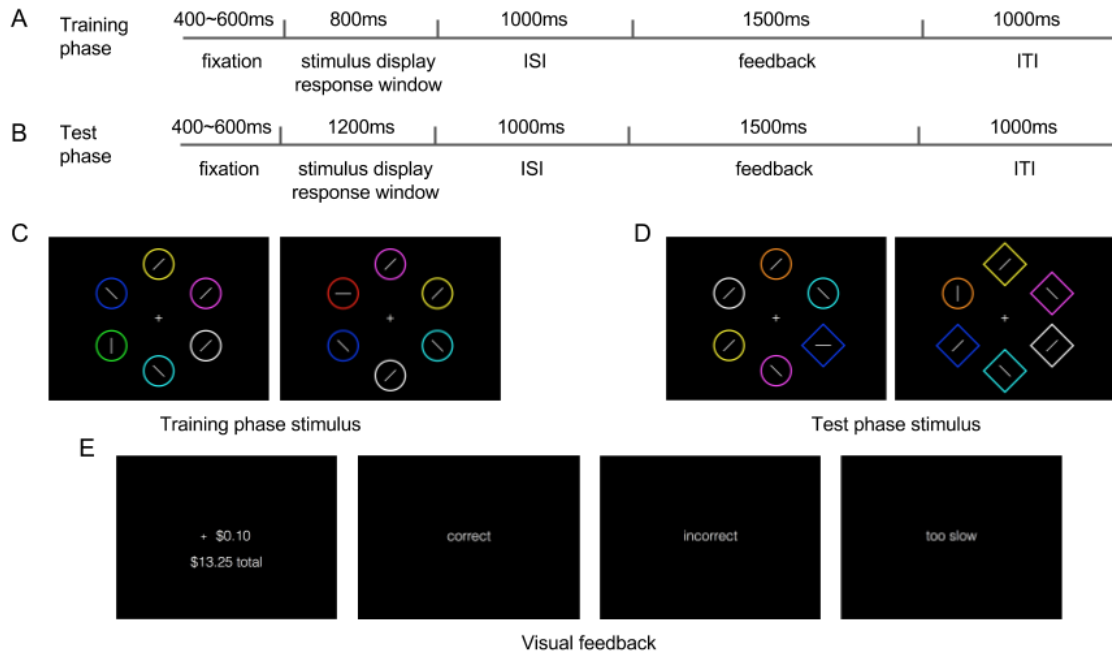


Figure 6. Experimental procedure and example visual displays. (A) Trial sequence in training phase. ISI, inter-stimulus interval. ITI, inter-trial interval. (B) Trial sequence in test phase. (C) Example visual stimuli in training phase. (D) Example visual stimuli in test phase. (E) Example visual feedback.

Eye tracking

Eye movement during trials in the training phase and the test phase was recorded with an Eyelink 1000 infrared eye tracker. The sampling rate was 500Hz. We used the remote viewing mode which allowed for unconstrained viewing. The eye tracker was calibrated using a 9-point configuration before Block 1 of each phase.

Data analysis

The goal of the current study is to investigate whether previous reward learning alternates the strategic balancing of speed and accuracy even when reward was removed. In addition to the baseline strategy, trials that contained a red or green distractor in the test phase introduced the possibility of attentional shift due to exposure of these value-associated distractors in the first

phase. For this reason, only RTs of trials in the test phase that did not contain red or green distractors were included in subsequent analysis. Trials in which the RT was less than 200ms were excluded in the analysis (Palmer, Horowitz, Torralba, & Wolfe, 2011). In general, the RT distributions presented a skewed tail to the right end for this task, which is not uncommon in psychophysics tasks (for an example RT distribution of a participant, see Figure 7).




Figure 7. A typical example of RT distribution and estimated exGaussian distribution fitting.

To better account for the entire distribution at the individual level, we fitted the RT distributions with an exponentially modified Gaussian distribution (exGaussian distribution; Lacouture & Cousineau, 2008). We used the open-source exGaussian model fitting Matlab package provided by the Visual Attention Lab at Harvard University (Wolfe et al., 2011). The exGaussian distribution consists of three parameters: 1) μ , mean of the Gaussian component; 2) σ , standard deviation of the Gaussian component; and 3) τ , rate of the exponential component. The probability density function of the exGaussian distribution is specified as:

$$f(x|\mu, \sigma, \tau) = \frac{1}{\tau} \exp\left(\frac{\mu}{\tau} + \frac{\sigma^2}{2\tau^2} - \frac{x}{\tau}\right) \Phi\left(\frac{x - \mu - \frac{\sigma^2}{\tau}}{\sigma}\right)$$

where Φ represents the cumulative density of the Gaussian distribution. The parameters were estimated through maximum likelihood estimation. Parameters were compared at the group level.




Figure 8. Example exGaussian distributions.

The effects of parameter values on the exGaussian distribution are reflected in Figure 8, where the statistical meaning of each parameter can be observed. In particular, change of μ shifts the distribution horizontally, change of σ alternates how wide the distribution is, and τ reflects the rate of decreasing in the tail (exponential component).

Excluded participants

Data from two participants, one in the experimental group and the other in the control group, were excluded from subsequent analyses due to failure of convergence during exGaussian parameter estimation.

Results

Overall, group average of mean RT ($M = 0.70$) was significantly greater than group average of median RT ($M = 0.68$) for all individuals (t -test, $p < 0.001$) (Figure 9). This strongly indicated that positive skewness is a common trend in RT distributions for this particular task.

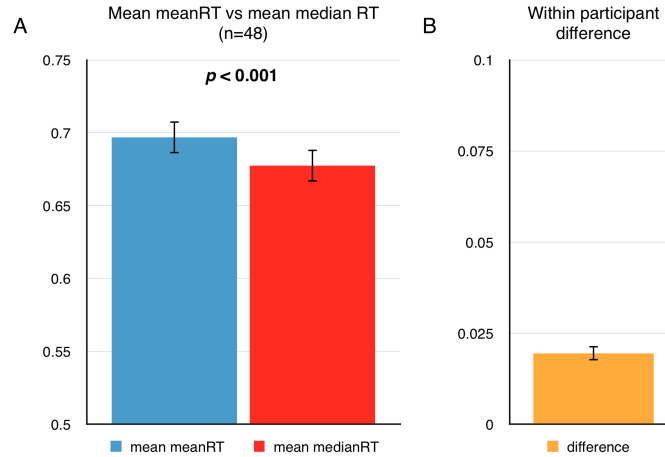


Figure 9. Average mean RT and average median RT for all participants, $n=48$. (A) Blue bar, average mean RT. Red bar, average median RT. (B) Yellow bar, average within-participant difference between mean RT and median RT. Error bars, standard error of the mean.

Effect of reward learning on RT variability

The results for the comparison between group averages of all three exGaussian parameters are presented in Figure 10. Our finding showed that the experimental group exhibited a wider RT distribution than the control group did, but not a significantly different mean of the normal component or rate of the right tail in the RT distribution. Specifically, we found that performance-dependent reward did not result in significantly different estimates of μ (t -test, $p = 0.57$). The group average of τ also showed no significant difference across two groups (t -test, $p = 0.21$). The group average of σ was significantly greater in the experimental group than in the

control group (t -test, $p = 0.03$). This evidence suggested that performance-dependent reward alternated RT variability in the test phase when reward was absent.

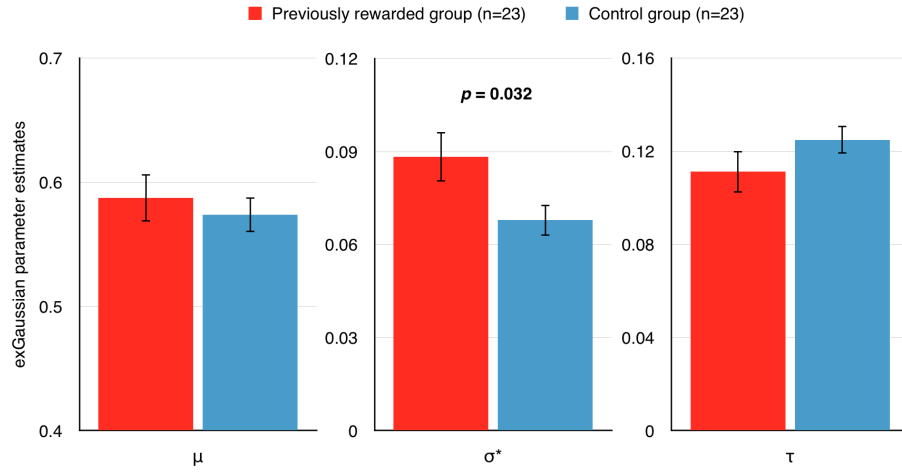


Figure 10. Comparison of group averages of μ , σ , τ , between the experimental group and the control group. Red bars, experimental group. Blue bars, control group. Error bars, standard error of the mean across participants. *, significant between group difference.

Correlation between RT variability and error rate

There was also a strong and positive correlation between σ and error rate for both groups on trials in the test phase with no red or green distractors ($r = 0.51$, $p < 0.001$; Figure 11). This suggested that a wider RT distribution is associated with higher error rate overall, reflecting a behavioral cost to increased variability.

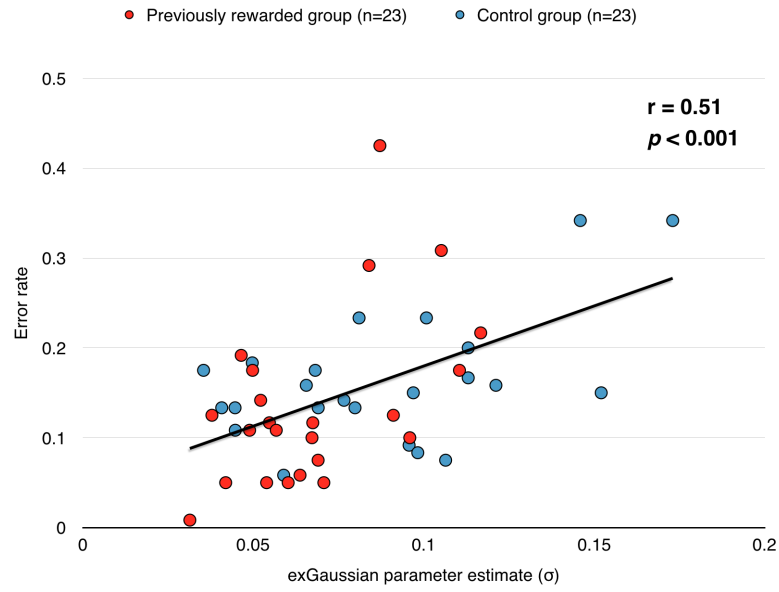


Figure 11. Error rate as a function of σ . Red instances, experimental group. Blue instances, control Group. Black line, linear trend line of all participants.

Discussion

The current study found that performance-dependent reward learning altered the response strategy in later stages. In particular, we showed that receiving performance-dependent reward in the training phase led to more variable RT across trials in the test phase. Further, a larger RT variability was associated with higher error rate, reflecting a critical behavioral cost that was associated with an unstable response strategy.

It is important to note that the observed RT distribution for this task was skewed. Therefore, using a distribution-fitting approach accounted for more characteristics of the RT distribution than simply using central tendency measures. The RT distribution appeared to be one of exponentially modified Gaussian shape, which is not uncommon among RT studies (Palmer, Horowitz, Torralba, & Wolfe, 2011). Additionally, our study provided support for the benefit of investigating speed-accuracy trade-off in tasks with time constraints. In tasks where participants were required to respond within a certain time constraint, an implicit decision of when to respond must be taken into account. If participants chose to ponder longer before responding, they would have a better chance of yielding a high performance, but also be subject to missing the response deadline. We showed that investigating this strategic balancing between speed and accuracy in RT studies provides insights to the phenomenon of value-driven attentional capture.

In contrast to the original VDAC study, our study provided a direct comparison between an experimental group with performance-dependent reward and a control group with flat-rate reward (Anderson et. al, 2011). Based on the significant group-level difference in RT variability, we found that performance in the experimental paradigm involved nuanced decision making factors that interfered with baseline trials. This led to a significant concern: using trials without

previously rewarded distractors might not serve as an ideal baseline to determine the effect of attentional shift. In particular, Anderson et. al (2011) interpreted the RT difference between trials with previously rewarded distractors and trials without previously rewarded distractors as an effect of attentional shift. However, group-level differences in RT variability revealed that RT can be influenced by reward learning even in trials without previously rewarded distractors. Since receiving reward learning altered the response strategy which carried over to the test phase, the slowing of RT might be partially driven by decisional factors other than an effect of attentional shift.

An increasing volume of studies have adapted this paradigm or employed a similar paradigm to investigate the phenomenon of value-driven attentional capture. Whether reward learning directed eye movement in accordance to attentional shift was extensively studied (Anderson & Yantis, 2012; Theeuwes & Belopolsky, 2012; Le Pelley, Pearson, Griffiths, & Beesley, 2015). For example, Anderson & Yantis (2012) showed that in the training phase, the first saccade after the search display was more likely to land on the side containing a rewarded target than on the side without a rewarded target. In the test phase, the first saccade tended to land on the side containing a previously rewarded distractor even if the target was on the opposite side. The effect of rewarded distractor on eye movement capture also persisted until the last block of trials. Theeuwes & Belopolsky (2012) also investigated oculomotor capture following reward learning. This study extended Anderson & Yantis (2012)'s result, confirming that distractors previously associated with high reward captured eye movements more often than distractors previously associated with low reward. However, this effect extinguished by the second half of the test phase. In addition, oculomotor capture under constrained viewing has

been assessed (Le Pelley et. al, 2015). This study examined value-driven oculomotor capture in an experiment that combined color saliency and reward magnitude. In addition, a customized time window was constructed for each participant through practice trials. The results were consistent with other findings, showing that high-valued distractors capture oculomotor movement more so than low-valued distractors did. Studies also showed that this particular type of attentional capture generalized to stimuli that share similar features with reward-related targets and its effect tended to persist over time (Anderson, Laurent, & Yantis, 2012; Anderson & Yantis, 2013).

Theories also centered on how exactly reward shifted attention. For instance, reward-based stimuli can capture attention by strengthening object saliency. One study demonstrated that when reward was associated with spatial location rather than stimuli features, it altered the representation of spatial priority map (Chelazzi, Eštočinová, Calletti, Gerfo, Sani, Della Libera, & Santandrea, 2014). Spatial priority map refers to a representation of the priority of spatial locations in the visual field that dynamically directs visual attention. In this study, reward was shown to have an influence on locations which were not subject to the stimuli features. Theeuwes & Belopolsky (2012) argued that the slowing of RT might have resulted from attentional holding rather than attentional capture, where the lengthened portion of slowed RT was the lengthened disengagement from the distractor.

Despite multiple accounts of why RT lengthened upon the presence of the distractors, it is generally consented that there was no top-down strategy involved in these tasks. It is true that paying attention to previously rewarded distractors contradicted with the task in hand, and would thus be counterproductive with respect to the explicit top-down strategy. However, because there

were time constraints imposed on the task, the implicit decision of when to respond within the time window could influence the response strategy that takes form in the training phase. If the strategies carried over to the test phase, they would interfere with RT observed in the test phase. Our study added an additional decision dimension to the value-driven attentional capture phenomenon. Our results suggested that longer RT might not only have been a result of attentional capture, but reflected the tension between speed and accuracy decisions relevant to the task. RT variability was further associated with behavioral cost. The inability to converge to a stable response pattern might have resulted from a more frequent distracted state, in that participants who received reward learning needed to frequently adjust their response patterns based on the performance on previous trials to ensure better performance in the future, especially following trials that contained previously valued distractors. The results provided a more nuanced account of how exactly value-driven attentional capture is manifested in RT and performance.

There are important timing factors to be considered in relevant studies. The current study, along with most of the studies mentioned above, were conducted with mostly participants in college populations aged around 20. Previous work has shown that the effect of value-based attentional capture persisted longer in adolescents than in adults, as the magnitude of slowed RT persisted over time for adolescents (Roper, Vecera, & Vaidya, 2014). This suggested that the effect of value-driven attentional capture might be sensitive to developmental stages. In addition, studies reported inconsistent results on the duration of the attentional capture effect. For instance, Anderson et. al (2011) and Anderson & Yantis (2012) found that the effect persisted over time,

while in Theeuwes & Belopolsky (2012) the effect diminished in the second half of the test phase. Future studies are needed to elucidate the VDAC effect across different time scales.

Two other potential confounding factors need to be noted. The first is inter-trial effect. Inter-trial effect refers to an elongated effect of stimuli in the current trial to subsequent trials. We analyzed trials in the test phase that did not contain previously defined targets to account for changes only due to reward learning (rather than attentional shift). However, because trials of all three conditions (previous high-valued distractor, previous low-valued distractor, no distractor) were randomized in the test phase, the collection of RT among trials without previously-rewarded distractors were intermingled with trials that contained previously-rewarded distractors. This means that although there were no previously rewarded distractors present spatially on the trials included in the analysis, valued distractors were present temporally. Trials that contained a valued distractor might have a remaining distracting effect in subsequent trials without previously-rewarded distractors. In this case, RT and performance on the current trial would influence the RT and performance of the next trial, which would suggest that variability in RT might not be entirely due to reward learning. A potential interpretation of such an effect is if the participants were put into a distracted state which does not diminish until after a certain period of time. Inter-trial effect can be addressed with a slight variation of the current experimental paradigm. One can include only trials without previously-rewarded distractors in the test phase. If we observe the same effect of contrasting RT variability between the experimental group and the control group, then we have strong evidence for the influence of reward learning on response strategy. This evidence would in turn lead to the decisional confounding of the effect of reward learning on attentional shift. If, on the other hand, no effect

of contrasting RT variability was observed, then inter-trial effect was at play, and reward learning did not alter response strategy. In this case, we can rule out the decisional confounds from response strategy in the effect of value-driven attentional capture.

The second potential confounding factor is different time constraints specified for the training phase and the test phase. We fully replicated the time constraints in Anderson et. al (2011) in our experiment, namely, 800ms for the training phase, and 1200ms for the testing phase. However, there exists potential confounding effects from this inconsistency of time constraint in the two phases, especially in terms of strategy formation. Even if participants actually formed different response strategies depending on the condition, they might further adapt the response strategy when they notice that they were given more time to respond for each trial. Thus it is not clear whether the contrasting variability in RT observed in two groups was entirely due to reward learning, or at least partially influenced by the loosening of time constraint. To investigate this further, a straightforward way would be to adapt the current experiment and administer consistent time constraints, making time constraints for both phases consistent. There are various options to further investigate the effect of varying time constraint on this particular task. For example, one can place participants in conditions with time constraint at several time windows, such as 200ms, 500ms, 800ms, 1200ms, 1500ms, 2000ms, and also a condition with no time limit.

Conclusions

We replicated the VDAC paradigm and found that reward learning altered the response strategy in baseline trials in the subsequent, unrewarded test phase. Further, a less stable response strategy is associated with a behavioral cost. These results suggested that value-driven attentional capture might be partially driven by decisional factors. We are following up on this study to investigate whether inter-trial effect interfered with changes in response strategy.

References

- Anderson, B. A., Laurent, P. A., & Yantis, S. (2011). Value-driven attentional capture. *Proceedings of the National Academy of Sciences*, 108(25), 10367-10371.
- Anderson, B. A., Laurent, P. A., & Yantis, S. (2012). Generalization of value-based attentional priority. *Visual Cognition*, 20(6), 647-658.
- Anderson, B. A., & Yantis, S. (2012). Value-driven attentional and oculomotor capture during goal-directed, unconstrained viewing. *Attention, Perception, & Psychophysics*, 74(8), 1644-1653.
- Anderson, B. A., & Yantis, S. (2013). Persistence of value-driven attentional capture. *Journal of Experimental Psychology: Human Perception and Performance*, 39(1), 6.
- Carrasco, M. (2011). Visual attention: The past 25 years. *Vision research*, 51(13), 1484-1525.
- Chelazzi, L., Eštočinová, J., Calletti, R., Gerfo, E. L., Sani, I., Della Libera, C., & Santandrea, E. (2014). Altering spatial priority maps via reward-based learning. *Journal of Neuroscience*, 34(25), 8594-8604.
- Corbetta, M., & Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature reviews neuroscience*, 3(3), 201-215.
- Gerrig, R. J., Zimbardo, P. G., Campbell, A. J., Cumming, S. R., & Wilkes, F. J. (2011). *Psychology and life*. Pearson Higher Education AU.
- James, W. (1890). *Principles of Psychology*, Vol. 1. New York: Holt
- Lacouture, Y., & Cousineau, D. (2008). How to use MATLAB to fit the ex-Gaussian and other probability functions to a distribution of response times. *Tutorials in Quantitative Methods for Psychology*, 4(1), 35-45.

- Le Pelley, M. E., Pearson, D., Griffiths, O., & Beesley, T. (2015). When goals conflict with values: Counterproductive attentional and oculomotor capture by reward-related stimuli. *Journal of Experimental Psychology: General*, 144(1), 158.
- Neisser, U., & Becklen, R. (1975). Selective looking: Attending to visually specified events. *Cognitive Psychology*, 7(4), 480-494. doi:10.1016/0010-0285(75)90019-5
- Simons, D. J., & Chabris, C. F. (1999). Gorillas in our midst: Sustained inattention blindness for dynamic events. *Perception*, 28(9), 1059-1074. doi:10.1068/p2952
- Palmer, E. M., Horowitz, T. S., Torralba, A., & Wolfe, J. M. (2011). What are the shapes of response time distributions in visual search?. *Journal of Experimental Psychology: Human Perception and Performance*, 37(1), 58.
- Pashler, H. E. (1998). *Psychology of attention*. MIT Press.
- Pavlov, I.P. (1927). *Conditional reflexes: an investigation of the physiological activity of the cerebral cortex*. Oxford, England: Oxford Univ. Press.
- Peirce, J. W. (2007). PsychoPy—psychophysics software in Python. *Journal of neuroscience methods*, 162(1), 8-13.
- Roper, Z. J., Vecera, S. P., & Vaidya, J. G. (2014). Value-driven attentional capture in adolescence. *Psychological science*, 25(11), 1987-1993.
- Theeuwes, J. (2010). Top-down and bottom-up control of visual selection. *Acta psychologica*, 135(2), 77-99.
- Theeuwes, J., & Belopolsky, A. V. (2012). Reward grabs the eye: oculomotor capture by rewarding stimuli. *Vision research*, 74, 80-85.
- Tsotsos, J. K. (2011). *A computational perspective on visual attention*. MIT Press.

von Helmholtz, H. (1867/2005). *Treatise on physiological optics* (Vol. 3). J. P. C. Southall (Ed.).

Courier Corporation..

Wickelgren, W. A. (1977). Speed-accuracy tradeoff and information processing dynamics. *Acta psychologica*, 41(1), 67-85.