INDUCING SLIPS OF ACTION:

CREATING A WINDOW INTO ATTENTION FAILURES:

by

Amanda Jane Clark

A thesis

presented to the University of Waterloo

in fulfillment of the

thesis requirements for the degree of

Master of Arts

in

Psychology

Waterloo, Ontario, Canada, 2007

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

ABSTRACT

Many of our daily activities are achieved through goal-oriented routines which illustrates the adaptability and efficiency of information processing. Nevertheless, slips of action do occur. This study was designed to determine if slips of action can be *induced* in a well learned task and if so, how these slips affect specific indicators of task performance. Thirty (12 male) right-handed undergraduate participants were taught, with arrow cues, a sequence of dominant hand movements. Following this learning phase, a portion of the sequences were altered by either changing the spatial location of the arrow cue or by changing the actual movement goal. Results revealed that participants made numerous action slips which were most prevalent when the movement goal was altered. This suggests that participants were unable to disengage their expected movement plan and thus were vulnerable to errors. In addition to exploring the frequency of action slips we also looked at participants' reaction and movement times on trials that preceded and followed errors and found that a speed-accuracy trade-off could not account for the slips. We also showed that frequency of slips on our task could reliably predict performance on the SART, a measure of inhibitory control, and the frequency of attention failures in daily living on the ARCES. Overall, the results of this study reveal that action slips can be induced by manipulating a well learned action routine and that the frequency of these induced slips reflects a participant's tendency to commit action slips in everyday life.

ACKNOWLEDGEMENTS

I would first like to thank my supervisor, Dr. Eric Roy, for his guidance, recommendations and friendship over the past couple of years while working on this thesis project. In addition, I wish to acknowledge the Natural Sciences and Engineering Research Council of Canada for research grants (to me and Dr. Roy) to support this work. Also, to my husband, Lucas, I thank you so much for your constant support, encouragement and loving nudges which have helped me to attain this goal. Finally, I want to thank my family for persistently believing in me and my abilities from the very beginning.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Introduction

In almost every hour of every day of our lives we are able to perform, with extreme ease, a number of activities that we once thought were virtually impossible. Consider for a moment your first time sitting in front of the steering wheel of a car. Quite possibly this moment was filled with excitement as it signaled your first foray into adulthood but chances are it was also accompanied by a sharp twinge of apprehension. How will I coordinate myself to be able to manage all of the controls? Gas and brake with one foot, clutch with the other; one hand on the stick shift while also trying to keep the machine on the road? At the time these tasks all seemed to require such a keen sense of attention and often mistakes were made as a result of pure inexperience. Now however, with a few years or decades of exposure to driving a vehicle, it should be amazing to reflect on how seamlessly all of the required actions are coordinated. Regardless of experience though, errors in the actions associated with operating a motor vehicle still occur. Now, while the errors made by skilled drivers are much different and perhaps less frequent that those made by novices, unfortunately, the errors that experienced drivers commit are often just as life-altering, if not more so.

Failing to check one's blind spot when changing lanes or not noticing a new stop light at an intersection on one's way home from work are dangerous yet common mistakes made everyday by many experienced drivers. While it is easy to appreciate the costs that could arise as a result of these mistakes it is also important to recognize the costs associated with some of the less dangerous but still annoying everyday errors that we all make. Meaning to turn left at an intersection to pick up some milk but instead turning right toward home like every other day is a common example of an annoying everyday error that results in lost time and efficiency and unnecessary frustration.

These everyday errors are certainly not unique to driving, instead, at one point or another they permeate almost every simple activity that we perform. In fact, on almost a daily basis we catch ourselves saying "what was I thinking…?!" after making a silly error during the execution of a simple task. Errors of this sort, or slips of action, are thought to indicate points in time when consciousness is either absent, disengaged or insufficient which allows automatic, unintended action sequences to be triggered inappropriately (Robertson, Manly, Andrade, Baddeley & Yiend, 1997). Typically, we are prone to these unintended action sequences when we are in familiar situations performing well-learned and overly practiced tasks (Reason & Mycielska, 1982). This type of environment makes very few demands on our conscious attention and as such, frees us up to think about or do a number of other things, leaving us vulnerable to distractions, boredom and slips (Broadbent, Cooper, FitzGerald & Parkes, 1982).

Late in the nineteenth century William James wrote extensively about errors that humans make in everyday life. He quite succinctly stated, "… habit diminishes the conscious attention with which our acts are performed" (James, 1890, pp. 114) and it is this and other statements made in his seminal writings that have caused many of his contemporaries to continue to pursue explanations about how and why these errors in everyday life take place. Central to this inquiry however is also an investigation into what role attention plays in managing the occurrence of everyday errors and for that matter, at least a primitive description of what attention is in the first place is also required.

In his book, *The Principles of Psychology* (1890), James devoted an entire chapter to his study of attention and this chapter is now used as a Bible of sorts that researchers today still use to describe the intricacies of the construct. While James' contemporaries continue to debate and feud over what attention is and is not, no other researcher yet has more concisely nor more elegantly provided a broad definition of attention than James himself. He wrote, "[attention]… is the taking possession of the mind in clear and vivid form, of one of what seems several simultaneously possible objects or trains of thought. Focalization, or concentration, of consciousness are of its essence. It implies withdrawal from some things in order to deal more effectively with others" (James, 1890, pp. 403-404). In this one brief definition, James clearly communicates that attention is, necessarily, a limited commodity and as such, it must be selective in order to achieve its goals. Inevitably then, while goal A is the object of one's attention, all other goals will need to be put off unless one is so effective at executing the task that it requires minimal conscious attention. Therefore, as James (1890, pp. 113) states, "… if practice did not make perfect, not habit economize the expense of energy … we would be in a sorry plight…" However, it is this exact fact that makes humans vulnerable to everyday errors as "these mistakes are the price we pay for being able to carry out so many complex activities with only a small investment of conscious attention" (Reason & Mycielska, 1982, pp. 243).

<u>Actions Not as Planned</u>

Considering the prevalence of errors in everyday functioning and the relevance of looking at failures of attention to provide insight into how attention operates, it might be surprising to find that little work has been done in this area. One of the main contributors to the line of inquiry however is James Reason. Through a number of studies either conducted by or inspired by Reason, he categorized errors of everyday life, or slips of action, as being either the result of inadequate planning or the result of unintended problems during the execution of the plan (Reason, 1977, 1979, 1984 and Reason & Mycielska, 1982). Everyday errors of the second type were termed 'actions-not-as-planned' by Reason and he suggested that they were of interest not only because of their potentially life-threatening consequences but also because they provided insight into the construct of attention and even more broadly, the overall functioning of the mind. Importantly, while these errors were not always the result of attention failures (perception and memory issues were also often implicated) his work on the subject has made large strides in bridging the gap between laboratory studies of attention and how attention actually functions in daily life.

Reason's Diary Studies.

In his writings, Reason makes clear that while measuring psychological constructs in real life situations is messy, for this topic of everyday errors, objective approaches are also riddled with difficulties. While everyone at one point or another has experienced one of these errors, they still occur relatively infrequently. In addition, as a result of the close relationship between reduced attention and error occurrence, the study of 'actions-not-as-planned' in a laboratory setting is highly susceptible to the artificial and intrusive methods of investigation. As such, Reason and his colleagues used two main methods of collecting information about actions-not-as-planned, participant diaries and self-report questionnaires.

In his first diary study, Reason had thirty-five volunteers keep a diary that detailed situations in which the actions that they performed differed from what they had intended. In addition to recording the action intention and what was actually done, the participants were also asked to keep track of the date and time of the error as well as the circumstances surrounding the error. This preliminary study resulted in a total of 433 incidents where a participant's action deviated from his/her intentions and it also caused Reason to wonder about other factors that may have been contributing to the errors. Consequently, he followed this first diary study with another, more elaborate, diary study in which sixty-three volunteers answered questions about the nature of their intended actions, the nature of their erroneous actions, their mental and physical states at the time of the errors and the prevailing environmental conditions when the errors were committed.

The results of these two studies revealed that slips were most frequent when executing activities that were highly practiced. In addition, Reason also reported that many participants indicated that when errors were committed they were carrying out activities in a "largely automatic way" (Reason & Mycielska, 1982). Interestingly, Reason also found large effects of time of day and the frequency of slips. Specifically, he reported that errors were most highly associated with periods of maximal activity since slips were most frequent just before noon and between five and seven o'clock.

Characterization of Errors. In addition to investigating the impact of time of day on the frequency of errors, Reason was also interested in characterizing and categorizing the types of errors that participants made. In the process of doing this, Reason found there were four main types of slips: actions that were repeating unnecessarily (i.e., putting an ingredient into a recipe twice), actions directed towards inappropriate objects (i.e., trying to use a comb to brush one's teeth), actions that are inappropriately incorporated into an action sequence (i.e., inappropriately putting vinegar in a recipe) and omitting required actions (i.e., neglecting to put cream into one's coffee). Importantly, with the exception of omission errors, Reason noted that all of the types of slips at least resembled normal actions that one could have performed quite acceptably under different circumstances. For example, if one prefers coffee with cream and sugar, putting cream in a cup of coffee, would perfectly acceptable. However, when preparing a cup of coffee for a guest who drinks their coffee black, this action would be considered erroneous.

Critical Decision Points. After closely analyzing the sections of the participant diaries that detailed the prevailing circumstances associated with the action slips, Reason also noted that the majority of reported slips occurred at points in an action sequence where a decision had to be made about how to proceed. Reason postulated that these decision points are critical in predicting whether an error would be made. Specifically he suggested that it was at these junctures where participants had to access information about the goals of the action sequence and what actions had been done up to that point to achieve those goals (Reason, 1979). For example, after adding cream to a cup of coffee one must consider whether one wants sugar in his/her coffee as well and if so, whether or not sugar has already been added. If one does not ask himself/herself these questions, he is liable to either omit the sugar step or repeat it. While Reason has suggested that everyday errors occur because of insufficient conscious attention and a failure to actually check on task goals, others have theorized that errors are the result of a degraded online representation of the overall task (Botvinick & Bylsma, 2005).

Like Reason's diary studies Botvinick and Bylsma (2005) were interested in looking at slips of action in everyday tasks but they attempted to do this within a laboratory setting. In their experiment they asked normal participants to make fifty cups of coffee and they intermittently interrupted the sequence of coffee making either in the middle of one of the actions (i.e. adding sugar) or toward the end of one of the actions. Their results indicated that while Reason and his colleagues might expect more errors after an interruption toward the end of an action (nearer to a critical decision point) accurate coffee making was more affected by disruptions that occurred mid-task. Botvinick and Bylsma (2005) have explained these results by theorizing that context information about a task is represented actively online and as such is vulnerable to disruption at any point in the action sequence.

Automaticity and A Model of Action.

Despite any disagreements in the literature about the specifics of action slip production, even Botvinick and Bylsma (2005) have conceded that checking and monitoring the progress of routine actions in everyday life is essential. In addition, few would argue that slips of action are the penalties that we pay for automization as the likelihood of slips is strongly correlated with one's experience with a task. Within psychology today, the term automatic, as it relates to cognitive processes, is often used to describe situations in which one task can be performed with ease and without interfering with other simultaneous tasks as it does not require many attentional resources (Norman and Shallice, 2000). Inherent in this working definition is the implication that attention is of limited supply at any one time and as such, when the tasks at hand stretch it beyond its boundaries, impaired performance is likely to result.

Despite the limitations of conscious attention humans are still able to perform, with very few errors, a number of actions simultaneously. This feat, in and of itself, is evidence of the strength of human information processing but taking that ability for granted provides little information about how the phenomenon actually occurs. However, Shallice and Burgess (1993) and later Norman and Shallice (2000) have suggested a model of action that seeks to explain the attentional mechanisms through which actions are carried out, either automatically or with deliberate conscious control. Their theory postulates two control systems, each based on different underlying neural networks that work separately at times and together at other times to accomplish action goals. One of the control systems, the supervisory attention system, works to allow for the conscious or controlled attention that is required when executing novel or dangerous tasks. Norman and Shallice (2000) as well as Reason (1984) also suggest that this system is responsible for exerting inhibitory action over habitual responses when they are inappropriate.

In order to allow for simultaneous action production, their model of action also includes a second system which they term, contention scheduling. This specific system, whose responsibility is the control of well-learning actions, has been suggested to operate primarily outside of conscious control except at decision points where one must consider the overall action goals (Schwartz, 1995). As a result, actions operated under the control of the contention scheduling system demand few cognitive resources. But while this may make sense from an efficiency point of view, actions controlled by this ‘auto-pilot’ like system are also more vulnerable to errors as they are not monitored closely to ensure their accurate performance.

Within the contention scheduling system, Norman and Shallice (2000) suggest that the errors that Reason (1977, 1979) characterized might occur for a number of reasons. For example, intrusion errors (inappropriate actions being incorporated into an action sequence) could be due to environmental triggers that are not inhibited by supervisory attentional control. Also, errors of omission (neglecting to include a necessary action into the sequence) could be the result of the insufficient activation of the routine action plan (schema). Regardless however, all of these types of errors are suggested to result from a failure in switching from the ‘auto-pilot’ contention scheduling system to the more consciously controlled supervisory attention system at those critical decision points.

Distractions and the Simon Effect.

Neglecting to switch between these two systems to accomplish action goals has been shown by both Botvinick and Bylsma (2005) and Humphreys, Forde and Francis (2000) to happen more often and with more detrimental results when the routine action is interrupted and/or accompanied by a second task. Others have shown that slips of action are more prevalent when external distracters are present in the environment (Buxbaum, Schwartz & Montgomery, 1998) while others still have suggested that internal factors like worry and boredom lead to increased action and attention slips (Manly, Lewis, Robertson, Watson & Datta, 2002). Both of these situations however are perhaps also explained by the supervisory attention system’s resources being deployed to attend to either the external or internal distracters instead of the central routine task.

Whereas few researchers have attempted to study action errors under conditions of internal preoccupation, a number of studies have been conducted to examine the impact of environmental distracters on action production. Much of the literature in this area is focused on patient populations (primarily victims of traumatic brain injury, patients with apraxia and Alzheimer's as well as patients with dysexecutive syndrome) and there is overwhelming agreement that within these populations, errors on everyday tasks are significantly correlated with an increased presence of distracters (Buxbaum et.al., 1998, Robertson et.al. 1997, Schwartz, 1995). Importantly too, while not investigated in the same patient populations, the physical location of environmental distracters may play a role in predicting the likelihood of upcoming action errors. Simon and Berbaum (1990) describes an effect named in his honour (the Simon effect) that shows that participants are quicker and more accurate when responding to targets whose physical location is compatible with the physical location of the stimulus/cue that instructed the movement. This effect is referred to by others as stimulus-response compatibility (Weigand & Wascher, 2005) and insofar as it is considered to be facilitating when the stimulus and response are compatible, the opposite can be said when they are not. Therefore, if one considers the potential impact of the Simon Effect on everyday action routines, one should be expected to prepare for tooth-brushing quickly and accurately if one's toothbrush and toothpaste are within close physical proximity to each other. However, errors might more likely result if one's toothpaste is, for some reason, located across the room or even worse still, if it is located near his/her comb.

Sustaining Attention

As has been discussed previously, Norman and Shallice's model of action (2000) describes two complementary systems which work together to control both routine and novel actions. Involved in these two systems is an understanding that the contention scheduling system will operate independently and outside conscious control whenever possible to conserve attentional resources. However, at critical checkpoints, the supervisory attention system may be required to intervene into the usual routine to accomplish slightly different action goals. The effectiveness of this intervention appears however to be reliant on having sufficient sustained attention to the task at those critical decision points. Such a seemingly simple process though is complicated by the fact that while new and exciting objects are able to draw our attention away involuntarily, focusing on routine tasks, even for a short period of time, demands a fair amount of effort and will. Now, while focusing one's attention on a tedious task is difficult for the most concerted individual, Robertson and colleagues (1997) have confirmed that patients with traumatic brain injury find the task even more challenging. They asked normal control participants as well as TBI patients to participate in a go-no-go paradigm, the Sustained Attention to Response Task (SART), which involves responding with a key press to a series of digits (one through nine) except when that digit is three. A task of this sort is extremely tedious but still requires each participant to actively attend to each digit as it appears or risk making an error. As such, Robertson and his colleagues (1997) and later Manly, Robertson, Galloway and Hawkins (1999) have asserted that one's ability to avoid making errors on the SART reflects one's ability to maintain consciously controlled, or sustained, attention. On this task the researchers measured accuracy and response time and found that those with traumatic brain injuries made more errors than the healthy controls. In addition, SART performance measures were also strongly correlated with responses on the Cognitive Failures Questionnaire (CFQ) which is a self-report questionnaire about the occurrence of cognitive failures in daily life (Broadbent, Cooper, FitzGerald & Parkes, 1982).

The CFQ which has been shown to be correlated with objective measures of selective attention (Tipper & Baylis, 1987), and was once considered an excellent measure of individual proneness to errors (Martin & Jones, 1984), has more recently received harsh criticisms concerning its ability to predict attention-related cognitive errors (Cheyne, Carriere & Smilek, 2006). In fact, some have suggested that the CFQ measures a number of underlying factors related to everyday errors, only one of which is related to attention. In response to this, Cheyne and colleagues (2006) sought to develop another self-report questionnaire that more specifically looks at errors in everyday life that are attributable to disengaged or insufficient attention. The Attention Related Cognitive Errors Scale (ARCES) was developed from relevant items on the CFQ, questions from Reason's self-report questionnaires as well as the personal experiences of the creators. The result was a very short, twelve item, questionnaire that asks respondents to rate how often certain slips of attention happen to them in daily life on a scale from one (never) to five (very often). Quite happily for the authors of this questionnaire, they found that it was highly correlated with frequency of errors on the SART, which is thought to reflect lapses in attention (Cheyne et.al., 2006). As such, they have suggested that the ARCES reflects one's propensity for making attention-related errors in everyday functioning.

It is important to recognize though that while there are definite advantages to asking for subjective reports of behavior, like with the ARCES or with Reason's diary studies, these reports are at risk of being incomplete and sometimes even inaccurate. However, like mentioned earlier, devising objective means of investigating action slips is difficult at best since they rarely occur in a natural environment, let alone a contrived one. Despite this though, developing effective methods of investigating action slips is important, primarily because of what they can tell us about the underlying mechanisms and processes that govern attention but also because of how they can help us to increase our awareness of what triggers slips, thereby helping us in minimizing their sometimes devastating consequences (Robertson, 2003).

Inducing Slips of Action.

In an effort to develop one such method of examining action slips, members of my lab have designed a paradigm that fundamentally differs from those used by other researchers in the area because our procedure actually induces action slips rather than having to rely on rare, and potentially flawed, recollections of events. The slip induction paradigm was devised to induce slips of action by requiring participants to deviate from a well-learned movement sequence. As such, participants were first taught a sequence of seven hand movements to a series of four targets around a central home location. Subsequent to having learned the sequence, endogenous (directional arrows) and/or exogenous (spatial location of arrows) cues are introduced that instruct participants to move to an unexpected location.

During the learning phase of the procedure, each of the movements within the sequence is indicated by an arrow cue that points to the location of the desired target button (up, down, left or right). In addition, in this phase, the arrow cue is always spatially compatible with the target location. As such, the cues contain both exogenous (the physical location of the arrow) and endogenous (the pointed direction of the arrow) information about the desired target. Using this cue information, participants learn the sequence of movements by practicing it for between 120 and 720 trials after which participants should be able to expect and anticipate each move within the sequence. As such, by the end of the learning phase of the study participants should represent the sequence not on a movement-by-movement basis but instead by the sequence of movements as a whole.

Subsequent to the learning phase, to induce slips of action, participants once again execute the movement sequence a number of times but in this manipulation phase some of the sequences are altered. These alterations can take one of three forms. Firstly, the goal of the movement may change by altering the pointed direction of the arrow cue. For the second type of alteration while the goal of the movement remains the same, the physical location of the arrow cue is changed. As such, when a participant expects an arrow cue to appear on the right that is pointed to the right target, the arrow now, while still pointed to the right target, is located either above, below or to the left of the central button. Finally, the third type of alteration is a combination of the two previously discussed. Consequently, this type of alteration includes both a change in the physical location of the arrow cue and a change in the actual movement goal.

Parakh (2003) has shown that the exogenous and endogenous manipulations to the cue information used in this paradigm are sufficient to induce action slips, particularly when the spatial location of an arrow actually acts as an attractor toward the expected yet inappropriate target. As such, they found that the first type of manipulation, where the goal of the movement is changed, but not the spatial location of the arrow, was most likely to create errors, or slips of action. Importantly though, while this alteration type was most detrimental to accuracy, performance on trials containing either of the other two manipulation types also successfully resulted in significantly more errors than trials that were not altered in any way. Therefore, Parakh (2003) surmised that the slip induction paradigm was an efficient way of inducing a number of action slips in a short period of time.

This particular study was designed to extend the work of Parakh (2003) by not only examining accuracy in more detail but also by looking at the micro-structure of sequence performance. To accomplish this, a procedure very similar to Parakh's (2003) was adopted but timing measures like reaction time, movement time and time to return to the home location were also incorporated. Therefore, while a significant amount of this study is devoted to replicating the results found by Parakh (2003) we are also very interested in the specific timing dimensions that may or may not predict the occurrence of action slips.

Like in Parakh's (2003) paradigm, participants in this study were assigned to one of three practice groups. Those participants pseudo-randomly assigned to group one received one block of practice and therefore practiced the sequence of seven movements only 120 times. In contrast, groups two and three practiced the sequence of movements for three and six blocks of 120 trials, respectively. The amount of practice received in the learning phase was manipulated to assess the potential differences that training may have incurred in the participants' propensity to commit action slips. If one is to believe what Norman and Shallice (2000) purport about the importance of switching from a contention scheduling system to a supervisory attention system to ensure accurate action production, then the degree to which the movement sequence was executed 'automatically' by participants should relate to his/her accuracy. Therefore, we expected that those participants who had received the most training had also learned the sequence more thoroughly and as such should execute the sequence more 'automatically'. As a result of this increase in automaticity we reasoned that participants who received more practice trials would only minimally involve the supervisory attention system to monitor performance and as such, they would be more vulnerable to errors when alterations were encountered. Conversely, participants with less training trials may not have had enough time to adequately learn the sequence and therefore, may not have been as prone to error following the alterations.

Considering the fact that Parakh (2003) observed the most errors after alterations that manipulated a movement goal, we expected that this type of alteration would also be most detrimental to performance. On the surface, making errors after this type of alteration is not surprising as participants are required to completely change their action plan. A change of this sort inherently implicates the supervisory attention system as it is required to inhibit the expected or routine action plan and also create a new plan of action. As we have seen though, the intervention of supervisory attention is not always dependable, especially when conscious attention is lacking (Norman &

Shallice, 2000). Consequently, one might expect that a number of errors would be made in this alteration condition. Importantly though, for this type of alteration, participants in the slip induction task might be even more prone to errors than was previously expected. Since the arrow cue, whose exogenous and endogenous cue information was compatible, is now spatially located in a place that draws attention to the incorrect target, participants might be liable to not notice the actual directional information in the arrow cue. As a result, participants' moves to a target might be simply based on the spatial, exogenous information and therefore they will make an error.

Another type of alteration that was used in this study involved incorporating external distracters into the environment. This was done by keeping the goal of the movement the same as expected but by positioning the arrow cue in an unexpected spatial location. This type of manipulation made it important for the participant to rely on the endogenous information in the arrow cue and to ignore the exogenous, distracting spatial information. It was expected that this type of alteration would induce action slips because of the Simon Effect (Simon & Berbaum, 1990). Simon reported that his participants took longer to respond to a target that was incompatible with its cue because they had to override their natural inclination to move toward the spatial location of the cue. As such, even though it might not be part of the sequence routine, we expected that participants might naturally move to a target that is near the cue, even if the cue does not point to that target. For example, even though the expectancy might be to move to the left target, if the arrow cue appears to the right, one might be likely to move to the right target button even though the arrow was pointed to the expected, left target.

Considering Norman and Shallice's model of action (2000), slips occur if the supervisory attention system does not intervene with the contention scheduling system when the routine behaviors must change in some way. As such, just before an alteration occurs, participants may be more prone to an error if they have gotten into a routine deeply enough that it is difficult to override. Consequently, for behaviors that are overly routine it is efficient for the action production system to prepare for movements in advance, perhaps in the time when a participant is returning to the home location after the previous movement. Therefore, to investigate this hypothesis, the time taken to return to the home button before an altered movement will be analyzed with reference to whether an error was made or not. If in fact it is the case that movements are planned in advance of the actual arrow cue, during the time taken to return to the home button, this return to home time should be significantly longer than for all other trials.

Regardless of the type of alteration that is present in this study, it is expected that participants will exhibit faster reaction times and movement times for trials in which an error is made than when participants are able to adjust their expected movement plan and therefore avoid an error. Robertson and colleagues (1997) have demonstrated this finding with their SART task and since our slip induction task also requires an inhibition of expected response, it is safe to assume that those results should carry over into our paradigm. What is still up for debate however, is the cause of this particular difference in reaction and movement times. While it is entirely possible that shorter reaction and movement times make a participant more prone to committing an error, it is also possible that when a correction to the expected plan is required, those times must be elongated to allow for changes to the plan. If this is the case, in order to allow time for the supervisory attention system to intervene and for the preparation of a new action plan, sequences that are altered but are correctly executed should have a longer movement time, and possible a longer reaction time, depending on when in the timing structure the alteration occurs.

In real life action production, the only way to guard against slips of action is to learn from them when they do occur. In keeping with this, and the fact that participants were motivated to avoid errors, it is expected that participants would attempt to avoid action slips whenever possible. Following an alteration that resulted in an error, it was hypothesized that participants would attempt to prevent subsequent errors by slowing down their movements and focusing more closely on the action sequence. Conversely however, following an altered yet correctly executed trial, it was expected that participants' timing measures would not differ from those associated with non-altered trials.

In addition to looking at the specific timing breakdown of sequences that were altered another focus of this study is to investigate the relationship between our slip induction paradigm and other well-established measures of action and attention slips. Specifically, the SART and ARCES questionnaire were selected and it was expected that number of errors on the slip induction task would predict the number of errors on the SART and also the amount of self-reported action and attention slips in everyday life as measured on the ARCES.

Method

<u>Participants.</u>

Thirty University of Waterloo undergraduate students (12 male) volunteered to participate in this study. All participants were right handed, had normal or corrected-to-normal vision and gave their informed consent before taking part. Twenty-five of the participants were recruited from Kinesiology 330, a University of Waterloo course on research design. These students chose to sign up to participate in this experiment from a list of available studies that were being conducted on campus. As part of their participation, after completing the requirements of the experiment they were given a synopsis of their individual data which they analyzed within a small group and used to write a brief research report for course credit.

<u>Experimental design.</u>

The purpose of this study was to induce a number of slips of action in a short period of time by requiring participants to deviate from a well-learned movement sequence. To accomplish this, participants were asked to learn a series of seven right hand movements to four target buttons located around a central home button. For each movement in the sequence, an arrow appeared either above, below, to the right or to the left of the central button that pointed to the target. As such, for each movement, participants received both exogenous, the physical location of the arrow on the screen, and endogenous, the pointed direction of the arrowhead, information about the target location.

Depending on their subject number, participants practiced the sequence of movements for either one, three or six blocks of 120 training trials. Subsequent to this learning phase, participants were asked to again execute the sequence of movements; however, in this experimental phase, twenty-eight percent of the trials were altered by either changing the spatial location of the arrow cue, by changing the direction of the arrowhead and therefore the actual movement goal or by changing both of these components. As such, three types of alterations to the learned movement sequence were introduced (see figure 1 in appecdix for pictorial representation of alteration types).

Type 1: Positional Alterations

For this type of alteration the goal of the movement remained as expected however, the spatial position of the arrow cue was changed. As such, for positional alterations, when participants expected to see an arrow located to the right indicating a movement to the right target, they actually saw an arrow pointed to the right but located either above, below or to the left of the central home button. Therefore, the only cue information that was unexpected was the exogenous information which is communicated by the spatial positioning of the arrow cue.

Type II: Directional Alterations

Directional alterations were exactly the opposite of positional ones. As such, while the spatial position of the arrow cue was as expected, the direction of the arrow head was changed and therefore the actual goal of the movement was changed from what was expected in practice. This meant that when participants expected to see an arrow located to the right indicating a movement to the right target, they actually saw an arrow pointed up, down or to the left, but yet still located to the right of the central home button. Consequently, this type of alteration only manipulated the endogenous information that the participants received.

Type III: Combined Alterations

In this final type of alteration, both the endogenous and exogenous information that the participant received was changed. Thus, both the spatial location of the arrow cue and the pointed direction of the arrowhead were randomly changed from what the participants expected. As such, for combined alterations, even though participants may have expected an arrow to appear to the right and point to the right target, the arrow cue could have actually appeared in any one of the other spatial locations and was pointed to any one of the other targets.

<u>Stimuli and apparatus</u>

The sequence of arrow stimuli used in this experiment was created using Micro Experiments Laboratory (MEL 2.0). Each of the arrow cues that were displayed using this program measured 20 mm in length, with 10 mm arrowhead fins, and they were displayed 125 mm from the center of the screen in one of the four directions. The sequence of arrow stimuli were shown on a 15 inch flat-screen monitor that was inverted to allow the stimuli to be projected onto a mirror which occluded the participants' hands (see Figure 1 for depiction of this setup). Situated under the mirror was a 16 inch by 16 inch button board equipped with five 2 inch diameter buttons, one located centrally with the others located to the north, south, east and west of the central home button. Participants made movements to and from these buttons and accordingly participants' reaction times and movement times were recorded when these buttons were released and depressed. Each participant was seated directly in front of the apparatus at a distance where the tip of their fingers barely brushed the back of the apparatus and all five buttons were within easy reach.

<u>Procedure.</u>

Upon arriving at the laboratory participants were informed of the general procedures of the study including the risks and benefits that they might incur. This discussion included making sure they knew that the study would require them to come into the lab on two separate days and that the study would take in total up to 3 hours of their time. After giving their informed consent, even-numbered participants completed the Sustained Attention to Response Task (SART) and the Attention Related Cognitive Errors Scale (ARCES) while odd-numbered participants skipped this step and proceeded directly to the training phase of the slip induction task (they would complete the SART and ARCES at the completion of the study).

For the training phase, participants were assigned to one of three practice groups depending on their subject number. Regardless of the amount of training however, the participants received the same instructions and the learning phase always began by quickly getting acquainted with the button locations on the response board. Once familiar with the response board participants were informed that a series of arrows were going to appear and their task was to move as quickly and as accurately as possible to the buttons on the response board that corresponded with those arrows. During this training session, the position and direction of the arrow cues were never manipulated and the participants were informed that this was the case.

Five to eight days after the training phase, each participant returned to the laboratory for the experimental session. This second phase of the study took approximately the same amount of time for each participant as each were required to complete 5 blocks of experimental trials with 120 trials per block. Out of these 600 sequences 28% were altered. Seventy of the trials had the directional alteration, forty-two were altered positionally and another twenty-eight involved the combined alteration. At the beginning of this experimental phase of the study participants were informed that a portion of the sequences would be changed in some way and that their task would be to follow the arrow's instructions. As such, if an arrow appeared that pointed to a new target, they were to move to that new target as quickly and as accurately as possible. Before actually commencing the section of the study in which trials were altered, the participants were first given an opportunity to become reacquainted with the movement sequence on a series of 60 reminder trials which were not altered in any way before beginning the experimental blocks. For the first few minutes of both the training and experimental phases an experimenter stayed with the participant to ensure that they had understood the instructions and they had an opportunity to ask questions while actually experiencing the protocol.

Turning now to the sequence of events for each trial, a fixation cross appeared in the center of the screen at the beginning of each sequence of seven movements. This fixation cross remained for between 500 ms and 1500 ms to ensure participants were not able to predict when the sequence of arrow cues was going to begin. Once the fixation cross disappeared, the participant pressed the central home button, which automatically triggered the onset of the first arrow cue. Upon seeing this arrow cue, the participants released the home button and quickly moved to the target that it pointed to. At the release of the home button the participant's reaction time (RT) was recorded and this also triggered the beginning of the movement time measure. Once reaching the target, participants quickly pressed the button which signaled the end of movement time (MT). They subsequently released the button and immediately returned to the central home button. The time that elapsed between the release of the target button and the next depression of the home button was also recorded and this measure will herein be referred to as 'return to home' (RtH) time. In addition to RT, MT and RtH time, the overall time that it took to complete each sequence (sequence time, ST) was also recorded as was the participants' overall accuracy.

Analyses.

Establishing practice effects.

To investigate the potential effects of the three amounts of practice, a number of analyses were conducted looking at the overall time to complete the movement sequence as well as the number of errors made. Considering first the time to complete the sequence (ST), a one-way ANOVA, accompanied by the Tukey's HSD post-hoc test, was computed using the STs from the participants' last 120 trials of practice. This test was followed by similar within-subjects ANOVA and Tukey's HSD tests which examined the hypothesis that the amount of errors made would decrease with more blocks of training.

Reminder Trials. To determine whether any practice effects in the training phase were maintained into the second day of testing, STs and accuracy were also examined for the 60 reminder trials. For these tests, using accuracy and ST as the dependent measures, the three groups of participants were contrasted using one-way ANOVAs.

Exploring accuracy.

Number of errors made by each subject were tallied and grouped according to whether the error was on a directionally altered trial, a positionally altered trial, a trial with a combined alteration, or a trial that was not altered in any way. These four error frequencies were then converted into percentage accuracy scores by dividing the number of errors made on each type of trial by the total number of possible errors that could have been made. The resulting accuracy scores were tested against each other by using a one-way ANOVA and a follow-up Tukey's HSD post-hoc test. This was used to determine whether the alterations were successful at inducing slips and also to explore whether any certain alteration was more effective than the others.

Examining the timing structure.

In light of the hypotheses about movements immediately proceeding and following alterations that were discussed earlier, a number of statistical tests were conducted. Four groups of statistical tests resulted. One group of tests looked at participants' RT, MT and RtH for the trial immediately before an error was made. A second group investigated these measures right before a correct, yet altered, trial was executed. The third group of analyses explored the same measures for the trial immediately following an error and the fourth group looked at the measures after a correct, yet altered, trial was completed. All of these ANOVAs used amount of practice as a between-subjects factor and the dependent measure of interest (RT, MT and RtH time) as a within-groups factor.

Considering congruence with other measures.

To determine the extent to which other accepted measures of inattention would predict performance on this slip induction task, a number of regression analyses were computed. Firstly, the SART task features measures of hits (situations where a button press is withheld accurately), misses (situations where a button press is not withheld when it should have been) and false alarms (situations where a button press is withheld inappropriately). In addition, participants' subjective reports of experiencing attention failures in daily life were scored on the ARCES questionnaire. Using a step-wise regression to eliminate the potential effect of practice group, each of these components, SART misses, SART false alarms and ARCES score were regressed upon the total number of errors, the number of errors made on altered trials and the number of errors made on unaltered trials. Finally, to ensure that these results were interpretable in light of the literature available on these two measures, correlations were also conducted to investigate the degree to which ARCES scores predicted performance on the SART.

Results

While the means were in place to discard any individual sequences in which participants made more than three errors, those methods were not used as no one executed more than two errors in any one sequence in either the training or the experimental sessions.

Impact of Training Group

Before looking specifically at the ability of our paradigm to induce slips of action it is important to examine participant performance during the training session. As such, the results discussed in this section are grouped according to the amount of training that each participant received. Shown below (in Table 1) are the average times to complete the entire sequence of seven movements for the final block of practice that participants received.

Table 1: Average sequence time (ST) in milliseconds for each training group in their final block of training.

Training Group	Average ST (ms)
1 training block	4250.34
3 training blocks	2934.44
6 training blocks	2916.79

This data clearly suggest a trend that participants with increased training are able to execute the sequence more quickly than those with less training. This trend was confirmed statistically through a between groups ANOVA which showed a main effect of training group on time to complete the sequence, $F(2, 27) = 14.402$, $p < 0.001$. In addition, a Tukey's HSD post hoc analysis of this main effect revealed that while participants with one block of training were significantly slower than those with three or six blocks of training, those groups with more training (groups 3 and 6) were not significantly different from each other (see Figure 2).

Figure 2: Average sequence time (ST) in milliseconds for each training group and for each block of training.

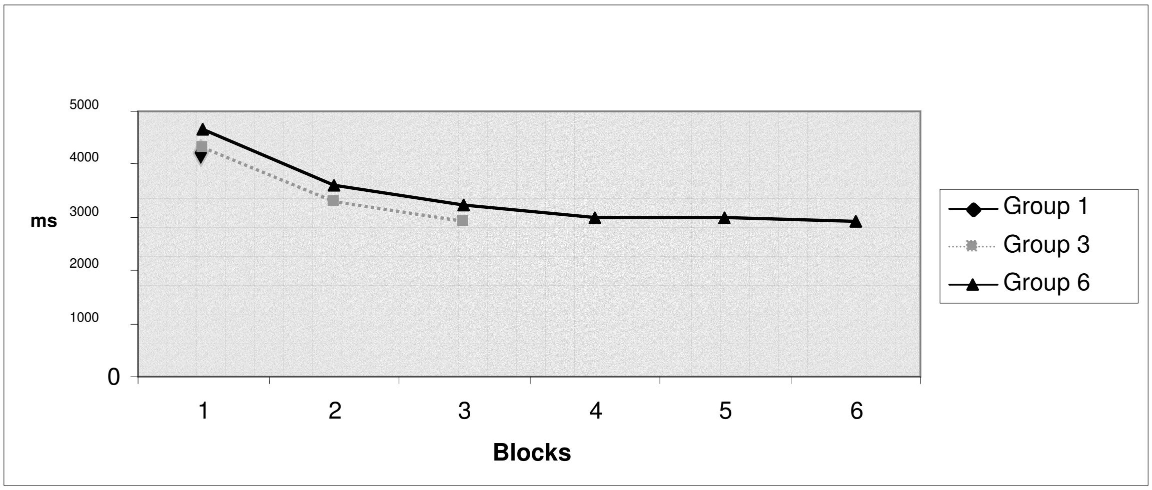

Upon examining whether this effect of training group on ST was carried over into the reminder section of the experimental session a one-way t-test showed even after a delay of five to eight days, those with more training were still performing the sequence significantly faster than participants from group one who received only one block of training, t (28) = 2.085, p = 0.046 (see figure 3).

Figure 3: Average sequence time (ST) in milliseconds for participants with one block of training and participants with three and six block of training during the reminder section of the experimental session.

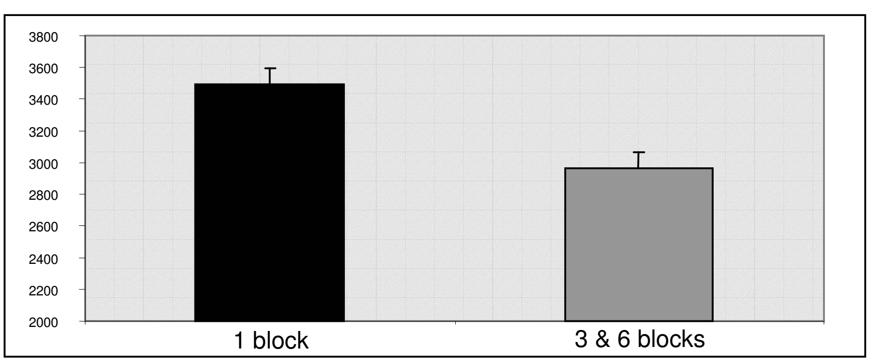

Alterations and Slip Induction

One of the main goals associated with this study was to determine whether the alterations introduced into the well learned sequence would be sufficiently confusing to induce slips of action. Consequently, the results discussed in this section will address this question. Upon examining accuracy in completing the sequence of movement, a main effect of alteration type is observed, $F(1, 29) = 6.643$, $p = 0.015$. As such, our prediction that participant performance would dramatically decline upon the introduction of alterations of any type was confirmed. In addition, post-hoc analyses with Tukey's HSD revealed that each alteration type also significantly differed from each other. Therefore, while accuracy was significantly worse in response to a combined alteration than a positional alteration, accuracy after a directional alteration was also significantly lower than for each of those types previously mentioned (see figure 4). These results support previously reported data collected by Parakh (2003) and further establishes their claims that directional alterations are most detrimental to performance.

Figure 4: Overall performance accuracy across alteration types.

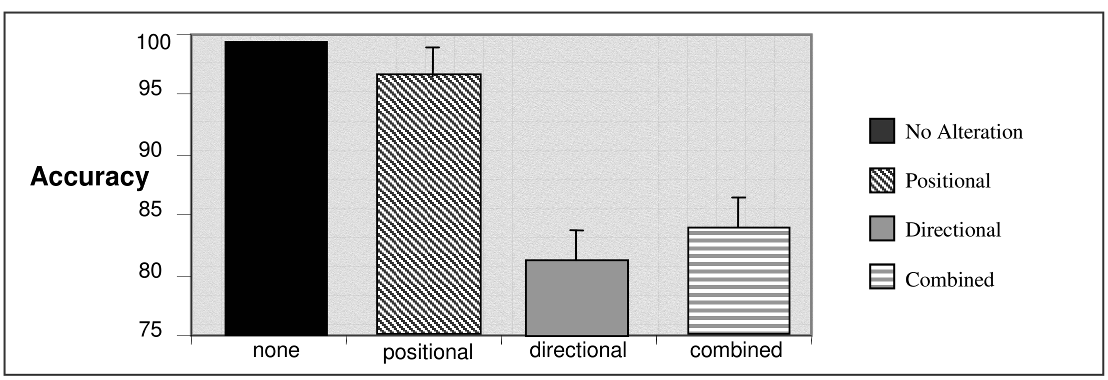

Interestingly however, while directional alterations were most likely to induce action slips, even positional alterations, where only the spatial position of the arrow cue was changed, were able to induce significantly more slips than when no alteration was present, $t(29) = 4.199$, $p < 0.001$. This finding is most likely the result of the Simon Effect which posits that an incompatibility between a stimulus and response can lead to increased movement times as well as increased errors (Simon & Berbaum, 1990).

Accuracy and Training

Considering the effects that were observed during the training and reminder sections of this study, we also predicted that participants with increased training would have learned more information about the sequence, perhaps would have performed the sequence more automatically than the other participants who received less training and therefore would commit more action slips. This hypothesis however, was not supported by the data. In fact, no significant differences were observed between the amount of training received and participants' accuracy performing sequences with either alteration type (see Table 2). Most likely, the strength of this relationship was diminished thanks to relatively few participants in each training group and these participants' insufficient propensity to make errors within a limited set of 720 trials. As a result of these factors, it is impossible to comment on the degree to which one's training on a task interacts with one's accuracy when performing it under altered conditions.

Table 2: Percentage accuracy for each training group when executing sequences that contained either positional alterations, directional alterations or combined alterations.

Training Group	Positional Alteration	Directional Alteration	Combined Alteration
1 training block	82.9	24.3	26.4
3 training blocks	87.6	26.4	38.9
6 training blocks	90.7	25.7	42.1

The Timing Micro-Structure of Trials

The results discussed up until now have focused on either the time to complete an entire sequence or the overall accuracy in completing the slip induction task. In this next section, the ST will be broken down into its micro-structure and as such the individual measures of reaction time, movement time and return to home time will be examined. In addition, each of these timing measures will be discussed for not only the actual trial that was altered but also the trials that preceded and followed the alteration (summary below in Table 3). Finally, for each situation where an alteration was present, the results are grouped according to whether an error was made or whether the sequence was executed correctly.

Table 3: Average ST, RT, MT, and RtH times (in ms) for unaltered trials, altered trials and trials that preceded and followed an alteration.

		Altered Trials		Preceding an Alteration		Following an Alteration	
	No Alteration	Error	Correct	Error	Correct	Error	Correct
ST	3179.30	3620.85	3937.94	3584.38	3925.73	3711.22	3960.16
RT	118.79	112.28	126.32	116.75	119.95	140.55	129.11
MT	184.83	193.59	465.17	171.54	191.3	316.08	314.05
RtH	175.65	163.50	184.85	173.79	185.82	488.62	320.27

Altered Trials.

As is evident in the above table, when considering the trials in which an alteration took place, there is a dramatic increase in the MTs associated with trials that were executed correctly as compared to trials that resulted in errors (see figure 5). While a trend toward increased RT and RtH times is seen for correctly executed trials, those differences were not statistically significant.

Figure 5: Average MTs (in ms) for altered trials that were executed correctly or resulted in errors.

Trials Preceding an Alteration.

Now considering the trials that preceded an alteration it is evident that while a trend might be emerging that RT, MT and RtH times are shorter for trials immediately before an error (see figure 6), no statistically significant results were established. Therefore, while it is possible that reacting and moving more quickly reduces one's likelihood of being able to correct one's expected movement plan, this is not substantiated by the results of this study.

Figure 6: Average RTs, MTs and RtH times for trials immediately preceding an alteration which induced either an error or a correction to the expected movement plan.

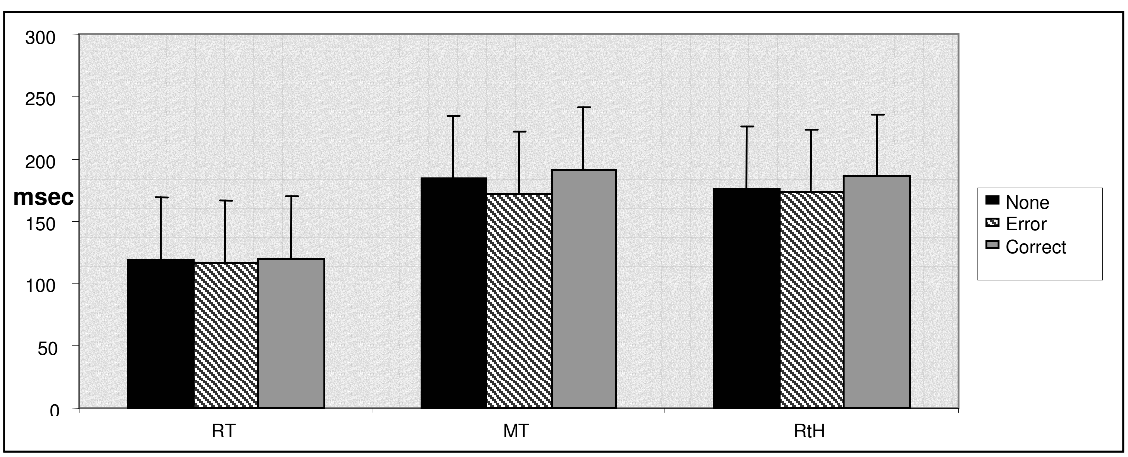

Trials Following an Alteration.

With respect to the breakdown of the timing measures, looking at the trials immediately following an alteration can lend insight into whether participants noticed that they had committed an error and how they reacted to making those errors. As is evident in Table 3 and below in Figure 7, participants' average RTs and MTs appear considerably longer following an error than for trials that were not altered. While this trend is not statistically significant for the reaction time data, it is strongly significant for the movement time data, $t(29) = -12.883$, $p < 0.001$.

Figure 7: Average RTs and MTs (in ms) for trials that were not altered, trials that were altered, and trials that immediately followed an alteration.

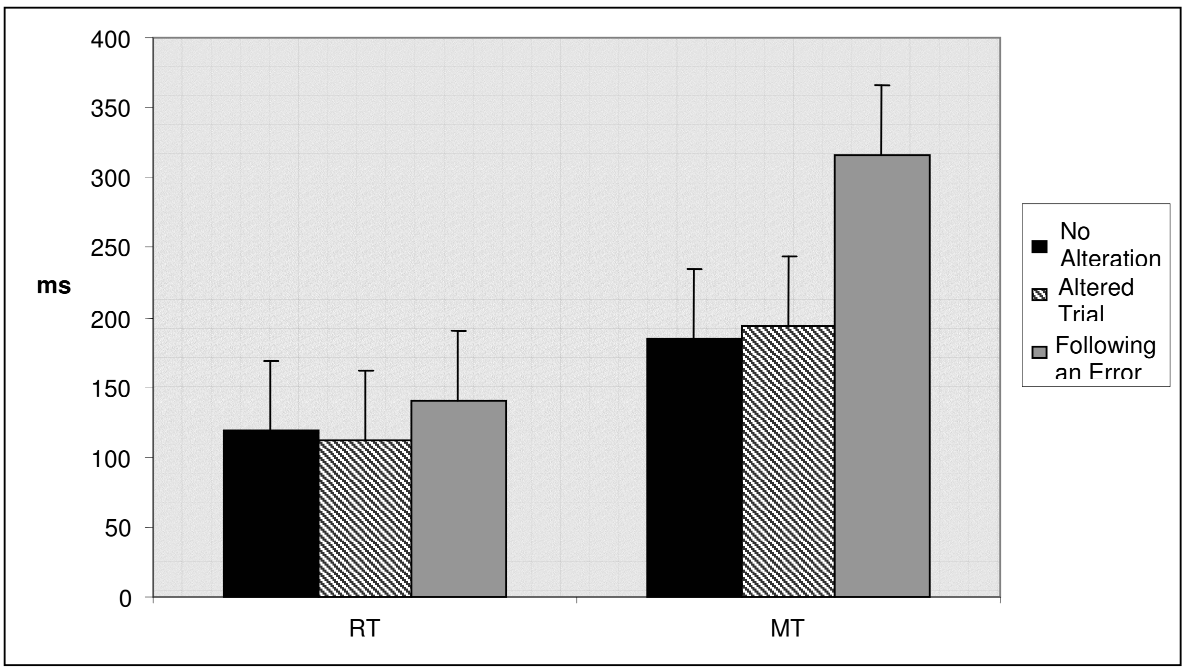

In addition to the effects discussed above with RTs and MTs, the RtH times for trials immediately following an error are also significantly longer than for those trials that were not altered, $t(29) = -13.478$, $p < 0.001$ (see Figure 8). This might suggest multiple things. One potential explanation is that perhaps participants did not even recognize that an error had been made until they were returning to the home button and as such, took extra time processing that slip. Alternatively, after recognizing that an error had been made, participants took extra time returning to the home button to 'reset' the automatized sequence in an effort to prevent subsequent errors.

Figure 8: Average RtH times (in ms) for trials in which no alteration was made, for trials that followed an altered yet correctly executed trial and for trials that were altered and resulted in an error.

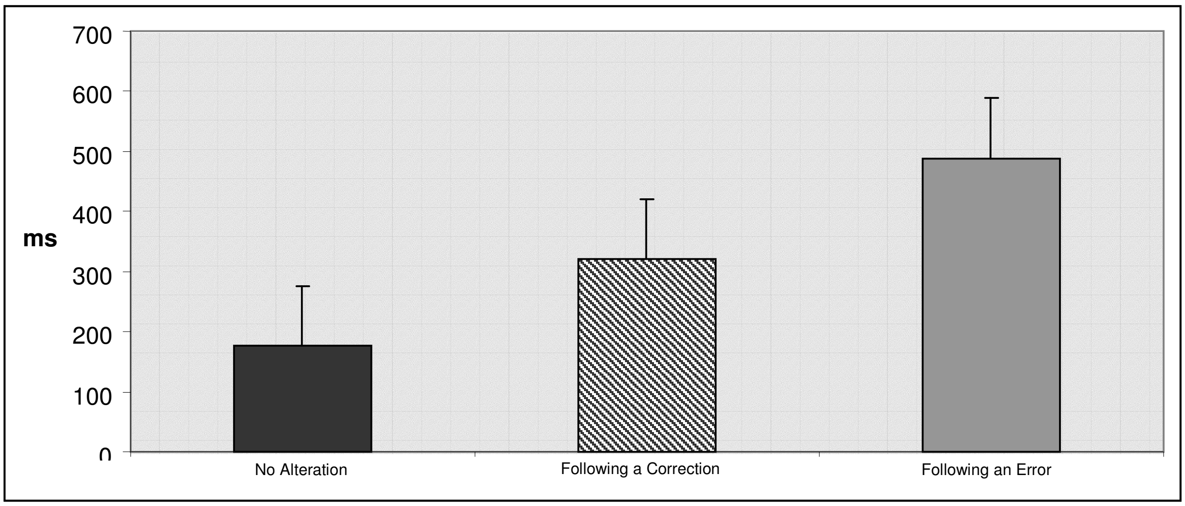

Relationship between the Slip Task and Other Measures of Attention Failures.

Another goal of this study was to examine how our slip induction paradigm is related to other widely used measures of attention failures, mainly, the Sustained Attention to Response Task (SART) and the Attention Related Cognitive Errors Scale (ARCES). In previous studies, scores on the ARCES have been shown to positively correlate with the number of errors made on the SART. As such, before examining the extent to which these measures are related to our paradigm, it made sense to first verify that a similar correlation was with the current population and this was indeed found but to a slightly weaker extent, $r = 0.186$, $p = 0.0325$.

Upon confirming that the SART and ARCES were operating within this population as expected, a number of regressions were performed to examine the relationship of each of these measures with our slip induction task. Regressions were done in a step-wise manner and the amount of training each participant received was included in the model. When looking closely at the regression model in each of the following cases however, the step which included the training group did not significantly explain any more of the variance. Therefore, one can conclude that there was no effect of training group on the ability of our slip induction task to predict SART errors or ARCES scores.

Sustained Attention to Response Task.

The main measures associated with the SART task are total number of errors (misses), total number of false alarms and response time (in ms). As such, each of these measures was examined in relation to the number of slips made on our task when either the trial was altered or was not altered in any way (no expectancies were violated). As expected, the number of errors that participants made on altered trials did significantly predict the number of errors that they would make on the SART task, $r = 0.549$, $p = 0.002$. This finding however was not replicated when looking at the participants' tendency to commit false alarms on the SART, $r = 0.210$, $p = 0.266$.

Attention Related Cognitive Errors Scale

The ARCES questionnaire is a self-report scale of how often failures of attention happen in a participant's daily life. As such, a high score on this questionnaire reflects a participant's feeling that attention failures occur relatively frequently whereas a low score indicates that attention failures are quite rare in their daily life. Upon examining the relationship between this measure and the number of slips made on our task (on either altered or unaltered trials), regression analyses revealed that it is only the number of errors made when a trial is not altered that predicts a participant's ARCES score, $r = 488$, $p = 0.006$. Interestingly, while this effect is quite strong, when considering the number of slips participants made on altered trials, the ability to predict ARCES scores plummets, $r = 0.162$, $p = 0.394$.

Discussion

This study was intended to shed light on the precursors to and results of attentional failures by experimentally inducing slips of action. This was achieved through manipulating a well learned action routine by requiring on-line adjustments to the expectations that participants had about upcoming movements. Previous studies on this topic have been limited by a number of factors which includes but is not limited to a reliance on rare and perhaps inaccurately recalled events. This particular study was designed to supplement previous work in the field by creating an experimental paradigm in which several slips could be induced in a brief period of time. In addition, this study allowed for the collection of detailed behavioral and physical data which helps to enrich the pre-existing database of published studies.

As was discussed earlier in the introduction section of this paper, virtually every person has had some experience where they intend to do one thing yet find themselves doing something else. Many have suggested that these slips of action are most likely in familiar environments (Manly et.al., 1999), while others have stressed the importance of distractions and drifts from conscious control (Broadbent et.al., 1982) in the generation of these errors. Perhaps most notably, Norman and Shallice (2000) have proposed a model of action which highlights the role of the supervisory attention system as a moderator or overseer of the system which is most active when performing highly routine tasks. This contention scheduling system, while extremely efficient, requires the intervention of the supervisory attention system in situations where the routine action at hand must be adjusted in some way because of a new goal or demand. As such, if the supervisory attention system fails to intervene at appropriate times, adequate adjustments to the routine action may be missed and a slip of action may occur.

For this present study, we were motivated to discover the conditions under which the supervisory attention system would fail to intervene. Consequently, we devised a paradigm in which participants were trained on a simple movement sequence and we manipulated this sequence in one of three ways. Firstly, positional alterations were designed to investigate the role of stimulus-response compatibility in generating action slips. Directional alterations were also interested in the role of stimulus-response compatibility to a degree however, this type of manipulation differed in that the actual goal of the movement had to be changed in order to execute the sequence correctly. In other words, an on-line adjustment, which requires active, sustained attention, had to be made to the expected action routine to avoid a slip.

Effectiveness of the Slip Induction Task

One of the initial goals of this study was to replicate the results found by Parakh (2003). In her Master's thesis, Parakh reports that the altered trials used in our slip induction paradigm were successful in generating significantly more action slips than trials that were not altered. This follow-up study found evidence which further supports this claim. In fact, the results of this study suggests that while directional alterations were most detrimental to performance, combined and even positional alterations also resulted in significantly more slips than trials which were not altered in any way.

The ability of this experimental paradigm to actually induce action slips, especially in response to directional and combined alterations, is interesting as it shows that at least when the goal of the action was changed participants were unable to adjust their actions in response to the unexpected changes in the routine. Consequently, it appears that the supervisory attention system was not effective at interrupting the contention scheduling system at the critical decision points in the action routine. For both directional and combined alterations, the goal of the movement had to change in order for the sequence to be executed correctly. It is expected that the reduced accuracy associated with the directional alterations is, at least partially, the result of the fact that the arrow cue was spatially compatible with the expected, yet incorrect target button. In addition, the fact that combined alterations also induced slips, though less frequently than directional alterations, suggests that something about the new physical location of the arrow cue reduced the participants' likelihood of making an error. One potential explanation for how this occurred is that completely violating the participants' expectancies, by changing both the direction and location of the arrow cue, made the changes more salient and therefore more obvious to the supervisory attention system.

In Norman and Shallice's model (2000) and even earlier, in Reason's writings (1977, 1979, 1984) the role of critical decision points is stressed. These points in the execution of any goal-related multi-step action are times when the performer is required to consider the overall goal of the action as well as the steps that have been completed to date to achieve that goal. Should these considerations not be made at the critical times, errors in the routine are likely and they might take the form of repetitions, omissions or outright errors in the course of action. It is probable that the slips of action that resulted from directional and combined alterations were, at least in part, the result of failing to check that the goal of the current action was the same as the goal of the action from the previous sequence. Now, while it is impossible, with this data, to ascertain whether the supervisory attention system failed to activate, or whether the supervisory attention system activated yet failed to interrupt the contention scheduling system, it is fairly safe to conclude that in one way or another, the participants' sustained attention was not always adequate.

Turning now to a consideration of the rarer, but still relatively pervasive, action slips associated with positional alterations, it was expected that participants' performance could be degraded even though the actual goal of the movement remained the same as what was expected. Based on research on the Simon Effect (Simon & Berbaum, 1990) movements made to spatially congruent targets were faster and more accurate than to spatially incongruent targets. Consequently, it was expected that by simply altering the spatial location of the arrow cue in this task, participants' attention would be directed to the cue and they would be drawn to respond congruently yet incorrectly. The results of this study indicate that while this type of error was made following only 13% of the positional alterations, this drop in accuracy is statistically significant. As such, our data suggests that while slips of action are primarily generated by changes in the goal of an action, they can also be generated by unexpected stimuli, or in other words the distracters that are present in the environment. This seems relatively intuitive since a number of action slips in daily life are reported to occur after a distraction (ringing phone in another room) or as a result of an inconveniently placed distracter object (sugar bowl located in the place where the cream normally is). What might not be intuitive however is the fact that positional alterations, those with an obvious environmental distracter, resulted in far fewer errors than the other two types of alterations. One potential explanation for this is that the type of environmental distraction matters. More specifically, for positional alterations the location of the arrow was different than what was expected, yet the location of the arrow did not coincide with the expected target. For direction alterations, which were most detrimental to performance, a few factors were working against participants and this appears to explain why more errors were made. With some thought, this explanation makes sense because it is quite simple to imagine a real-life predicament which would support this explanation. For example, when preparing a cup of black coffee one can imagine being distracted by a phone call and inadvertently adding sugar to the cup, but the chances of this happening are less than if the phone was located in close proximity to the sugar bowl.

The Micro-Structure of Sequence Timing.

The most significant addition that this paper makes to the study completed by Parakh (2003) is that each trial within each sequence was broken down into its individual behavioral measures. As a result, we are able to consider not only STs but also the individual RTs, MTs and RtH times for each trial in the study. This produced an abundance of data which has been categorized according to whether the trial was altered, came before an altered trial, followed an altered trial, or was not closely associated with an alteration. Within this section those trials that were altered and trials that preceded and followed an alteration will be discussed.

Altered Trials. For trials that were altered, RTs, MTs and RtH times were examined and some interesting observations were made. Firstly, as was expected, participants appeared to either move in almost exactly the same fashion as when a trial was not altered, which eventually resulted in an error (no change in MT was observed between trials that were not altered and trials where an error was made), or they made an on-line correction to their expected movement plan, which allowed for a correct trial. Evidence for this on-line correction is seen when examining the MTs for trials which were executed correctly, which are significantly longer than trials where an error was made. Further support for this is found when considering the RTs and RtH times as these measures did not differ for correctly versus incorrectly executed altered trials. Consequently, it appears that for trials where participants were able to adjust their expected movement plan in response to a change in the sequence, the successful intervention of the supervisory attention system occurred during the period of time when they were moving from the home button to the target.

Trials Preceding an Alteration. Earlier in this document the hypothesis was made that movements preceding an error would be highly routinized. Therefore, participants would come to expect movements in advance and would prepare for and program those movements in advance as well. In doing this, participants would have almost entirely programmed an upcoming movement even before the actual arrow cue appeared. This overlap creates a very small if not non-existent critical decision point which would be more rarely interfered with by the supervisory attention system. This hypothesis was not fully supported by these data. The lack of support may be the result of incomplete measures being in place to ensure that participants planned and programmed their movements during periods of time that we were actually recording. For example, in this study, no methods were in place to ensure that participants were unable to linger for a time on the target button. If participants did press and hold this target button while preparing for the upcoming move, no increases in the RtH time or subsequent RT would be necessary. Consequently, while it remains possible that participants planned and programmed their movements in advance of the actual arrow cue, which would have allowed them to react and move more quickly for the upcoming move and increased their likelihood of committing an error, this is not substantiated by the results of this study.

Trials Following an Alteration. Now considering the trials that immediately followed an alteration, our expectations were confirmed as participants' increased reaction times and movement times suggest that they were eventually aware of their earlier error and their speed-accuracy trade-off and probably attempted to prevent subsequent mistakes by slowing down their progression through the sequence. What remains unclear however is how long this effect continued. Knowing how long the impact of making an error helped prevent subsequent mistakes has large practical implications outside the laboratory. One would think that committing an error when making a cup of coffee would alert someone to their propensity to fail to attend to the goal of the action for the remainder of the task at least, but does this effect persist to the next day's coffee making?

Inducing Slips and Examining Them With Other Measures.

As was discussed earlier in this paper, one of the potential disadvantages of any experimental task is the fact that events as they occur in the laboratory are never perfect replications of behavior in real life situations. Because of this, much of the research in the field of attention and action slips has remained within the domain of subjective reporting. Studies of this kind are often extremely rich with descriptive information about the element of study however these descriptions are rarely trusted wholeheartedly. In reaction to this, we attempted to devise an objective method of investigating action slips while also considering previously established methods of looking at attention failures.

Either at the beginning of the training session or at the end of the experimental session, all participants in this study completed both the SART task and the ARCES questionnaire. Before making any claims about how these two tasks relate to the slip induction paradigm we first verified that the ARCES questionnaire was positively correlated with the number of errors on the SART to the degree that was previously reported in the literature. Since these two measures were related to each other as would be predicted we proceeded to examine their correspondence with the slip induction paradigm.

These analyses showed that SART misses, or trials where a response was made when it should have been withheld, were strongly correlated with the number of errors that participants made on our slip induction paradigm. This was not surprising as both of these measures require an ability to withhold and adjust the expected movement plan. It is encouraging however to find that both paradigms were successful in generating slips and both appear to tap into one's ability to inhibit a learned action routine.

In addition, the number of slips made on unaltered trials in the slip induction paradigm was found to significantly correlate with ARCES scores. As such, participants who made more errors on trials that were not altered in any way were also more likely to report attentional failures, like losing the thread of a conversation, in daily life. This relationship is extremely interesting as it suggests that participants who are more prone to errors in everyday routine tasks are also more prone to making errors on our expected, routine slip induction experiment. Considering the overall goal of the ARCES questionnaire in identifying persons who are less mindful and therefore more vulnerable to attention failures, the high correlation between unaltered errors on our task and the ARCES is very promising. Now, with reference to our slip induction paradigm, the high correlation between these measures may suggest that participants who make more errors on unaltered trials are more likely to remember or feel comfortable admitting to making attentional errors in everyday life. Another potential explanation however is that those who make more slips in everyday life are not being accurately identified when we consider only those who make many slips in response to alterations.

Conclusions.

Overall the results of this study on the ability to induce slips of action by manipulating a well learned action routine reveal that action slips can be experimentally induced. In addition, it has also been shown that one's propensity to make these laboratory based slips is highly related to one's likelihood of making attentional and action errors on other experiemental tasks as well as in activities of daily living.

Previous studies of this slip induction paradigm have focused on examining accuracy and overall sequence time in response to perturbations to a practiced action routine. The most significant addition that this paper makes to this literature is that each trial within each sequence was broken down into its individual behavioral measures. As a result, we were able to examine not only the time to complete an entire sequence, but also the individual RTs, MTs and RtH times for each trial in the study. Without this more detailed information, it would have been impossible to test theories about what participants were specifically doing before making an error, how they were able to make adjustments to their expected routine to allow for correctly executed trials and how they adopted more cautious strategies to prevent subsequent action slips.

Considering the effectiveness of recording these timing measures and the experimental alterations to the routine action sequence, this study opens up a number of new research avenues. Firstly, in the study reported here participants were instructed to also make movements according to the pointed direction of the arrow cue. As such, even though one might have expected to move to the left target, if an arrow instructed a new move to the upper target the upper target needed to be pressed in order for the trial to be executed correctly. These instructions resulted in an accuracy pattern where directional alterations were most detrimental to performance while combined and positional alteration also resulted in a fair number of action slips.

A follow-up study will soon be conducted to investigate the accuracy pattern when different instructions are given. In this future study, participants will be instructed to ignore the pointed direction of the arrow cue and instead to move to targets accordingly to the practiced sequence. As such, even if an arrow instructs one to move to a new target, the expected movement goal must be attained to execute the sequence correctly. For this follow-up study, it is reasonable to expect that very few action slips will be made by participants in response to alterations however, if any slips are made, it is sensible to expect that those errors will be in response to positional alterations. By comparing the accuracy pattern described in this paper with that from the projected follow-up study it will be possible to develop a theory about slips of action and whether they are the result of top-down or bottom-up processing. This is possible since in one case participants are required to construct their movements one-by-one in response to the cues while in the other participants are to attempt to override the current cues and move solely in response to their learned action routine.

In addition to this follow-up study, the slip induction task that was described in this paper can also be used to examine the occurrence of attentional failures in healthy older adults. Typically, it is assumed that older adults experience more slips of action and attention in daily life however, this assumption may not be valid as older adults report an increased ability to focus their attention which should actually act as a protective factor against action slips. Considering this, studying the slip induction task as well as the ARCES and SART with this population might help to elucidate a number of the assumptions that are held about action slips, the elderly and overall, a healthy older adult's ability to safely complete activities of daily living.

References

Botvinick, M. & Bylsma, L. M. (2005). Distraction and action slips in an everyday task: Evidence for a dynamic representation of task content. *Psychonomic Bulletin and Review, 12*, 1011-1017.

Broadbent, D.E., Cooper, P.F., FitzGerald, P., & Parkes, K.P. (1982). The Cognitive Failures Questionnaire (CFQ) and its correlates. *British Journal of Clinical Psychology, 21,* 1-16.

Buxbaum, L.J., Schwartz, M.F., & Montgomery, M.W. (1998). Ideational apraxia and naturalistic action. *Cognitive Neuropsychology, 15*, 617-643.

Cheyne, J. A., Carriere, J. & Smilek, D. (2006). Absent-mindedness: Lapses of conscious awareness and everyday cognitive failures. *Consciousness and Cognition, 15,* 578-592.

Humphreys, G.W., Forde, E.M.E., & Francis, D. (2000). The organization of sequential actions. In Monsell & Driver (Eds.)*, Control of Cognitive Processes: Attention and Performance XVIII.* Massachusetts: The MIT Press.

James, W. (1890). *The Principles of Psychology Vol. I.* New York: Henry Holt & Co.

Manly, T., Lewis, G.H., Robertson, I.H., Watson, P.C., & Datta, A.K. (2002). Coffee in the cornflakes: time-of-day as a modulator of executive response control. *Neuropsychologia, 40,* 1-6.

Manly, T., Robertson, I.H., Galloway, M. & Hawkins, K. (1999). The absent mind: Further investigations of sustained attention to response. *Neuropsychologia, 37,* 661-670.

Martin, M. & Jones, G. V. (1984). Cognitive failures in everyday life. In J. E. Harris & P. E. Morris (Eds.), *Everyday Memory, Actions and Absent-Mindedness*. London: Academic Press.

Norman, D. A., & Shallice, T. (2000). Attention to action: Willed and automatic control of behavior. In Gazzaniga (Ed.), *Cognitive Neuroscience: A Reader*. Malden: Blackwell Publishers Inc.

Parakh, R. (2003). *Inducing slips of action through the unintentional activation of a well-learned response.* Unpublished master's thesis. University of Waterloo: Waterloo, ON.

Reason, J. (1977). Skill and error in everyday life. In M. Howe (Ed.), *Adult Learning*. London: Wiley.

Reason, J. (1979). Actions not as planned: The price of automization. In G. Underwood & R. Stevens (Eds.), *Aspects of Consciousness.* London: Academic Press.

Reason, J. (1984). Lapses of Attention in Everyday Life. In R. Parasuraman & D.R. Davies (Eds.)., *Varieties of Attention.* Orlando: Academic Press.

Reason, J. & Mycielska, K. (1982). *Absent-minded? The Psychology of Mental Lapses and Everyday Error.* Englewood Cliffs: Prentice-Hall, Inc.

Robertson, I.H., Manly, T., Andrade, J., Baddeley, B.T., & Yiend, J. (1997). 'Oops!': Performance correlates of everyday attentional failures in traumatic brain injured and normal subjects. *Neuropsychologia, 35,* 747-758.

Robertson, I. H. (2003). The absent mind: Attention and error. *The Psychologist, 16,* 476-479.

Schwartz, M. F. (1995). Re-examining the role of executive functions in routine action production. *Annals of the New York Academy of Science,* 769, 321-335.

Shallice, T., & Burgess, P. (1993). Supervisory control of action and thought selection. In Baddeley and Weiskrantz (Eds.), *Attention: Selection, Awareness and Control.* Oxford: Oxford University Press.

Simon, J. & Berbaum, K. (1990). Effect of conflicting cues on information processing: The 'Stroop Effect' vs. The 'Simon Effect', *Acta Psychologica, 73,* 159-170.

Tipper, S. T. & Baylis, G. C. (1987). Individual differences in selective attention: The relation of priming and interference to cognitive failure. *Personality and Individual Differences, 8,* 667-675.

Weigand, K. & Wascher, E. (2005). Dynamic aspects of stimulus-response correspondence: Evidence for two mechanisms involved in the Simon Effect, *Journal of Experimental Psychology: Human Perception and Performance, 31,* 453-464.

Appendix

Figure 1: Pictorial Representation of apparatus (Parakh, 2003), button board and alteration types.

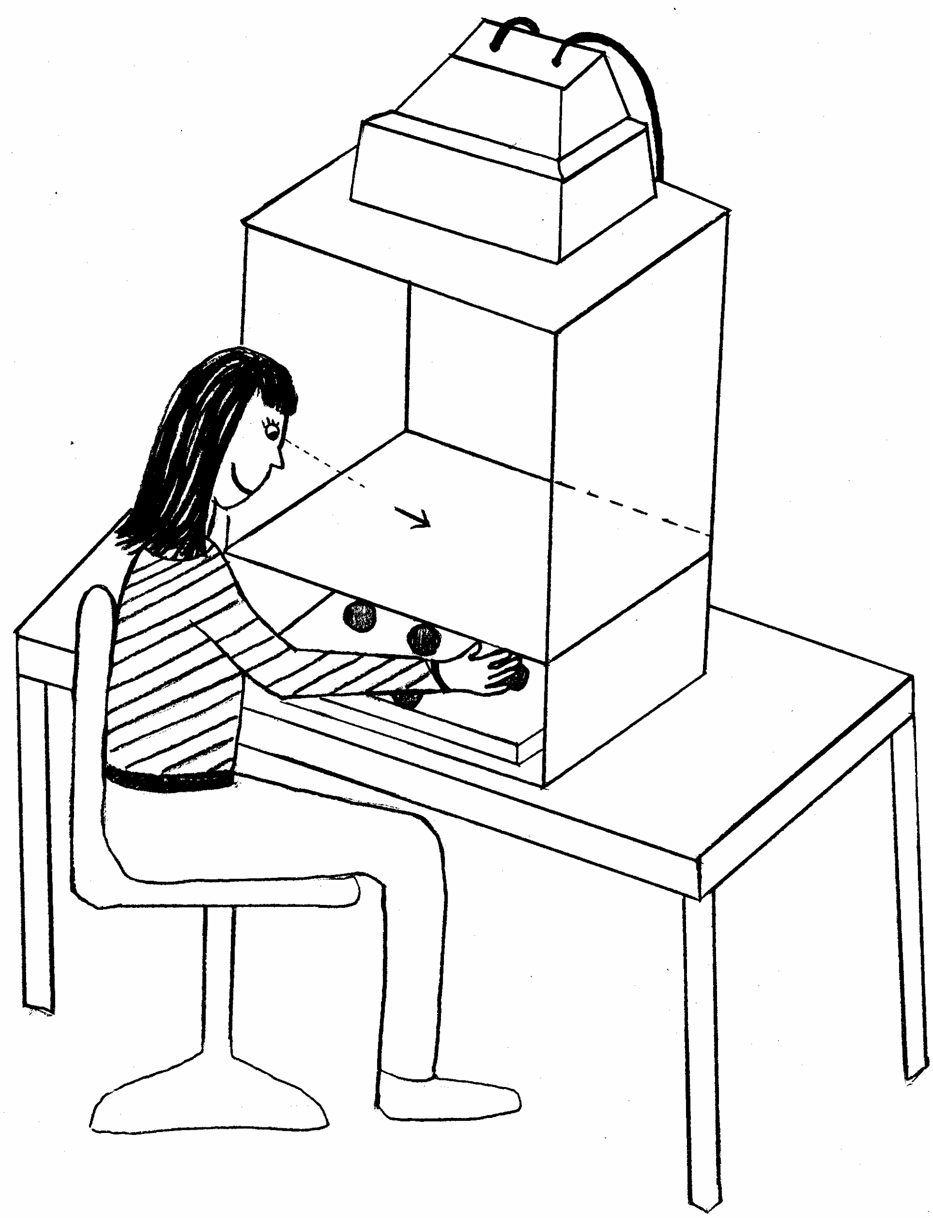

Positional Alteration Condition – note that the expected target (black with white dots) remains the same in this condition and that only the physical location of the arrow cue is changed.

Expected /
Learned Movement

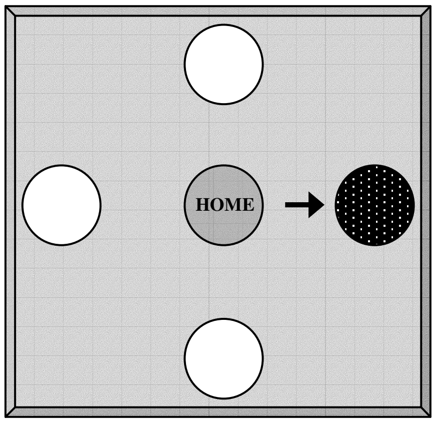

Alteration /
New Movement

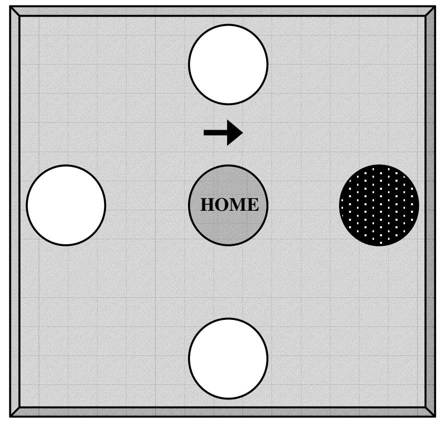

Directional Alteration Condition – note that the arrow cue is located in the expected physical location but that it does not point to the expected target (black with white dots) instead, a new button (black with white horizontal lines) must be pressed for a correct response.

Expected /
Learned Movement

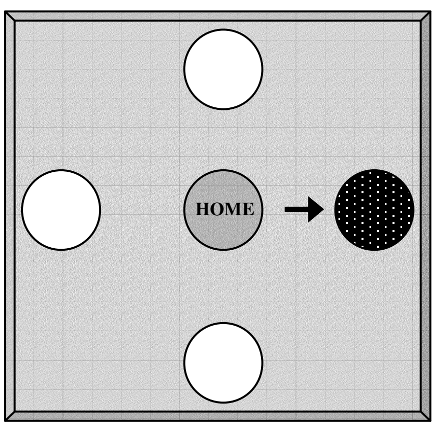

Alteration /
New Movement

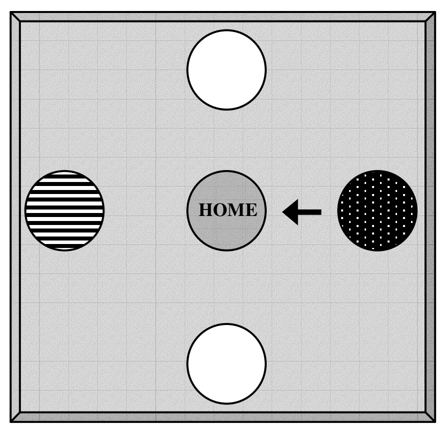

Combined Alteration Condition – note both the correct target (black with white horizontal lines) and the physical location of the arrow cue are not as expected.

Expected /
Learned Movement

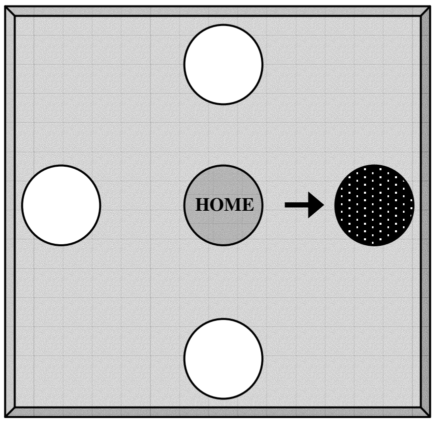

Alteration /
New Movement