Summer 2010

A funny thing happened on the way to the maze: incidental learning of irrelevant information in humans

Rose Halterman Danek
University of Iowa

Recommended Citation

Danek, Rose Halterman. "A funny thing happened on the way to the maze: incidental learning of irrelevant information in humans." PhD (Doctor of Philosophy) thesis, University of Iowa, 2010.
http://ir.uiowa.edu/etd/660.

A FUNNY THING HAPPENED ON THE WAY TO THE MAZE: INCIDENTAL LEARNING OF IRRELEVANT INFORMATION IN HUMANS

by

Rose Halterman Danek

<u>An Abstract</u>

Of a thesis submitted in partial fulfillment
of the requirements for the Doctor of
Philosophy degree in Psychology
in the Graduate College of
The University of Iowa

July 2010

Thesis Supervisor: Associate Professor J. Toby Mordkoff

ABSTRACT

Contemporary learning theories derive much of their explanatory power from the assumption that all stimuli presented vie for associative strength, the assumption of Shared Weight Space (SWS). Theories based on this assumption have proven successful in explaining many of the observed conditioning phenomena in animals. However, work with humans has proven more complex due to outside knowledge, biases, and heuristics (see, e.g., Chapman, 1991; Msetfi et al., 2005; Perales et al., 2004; Tversky & Kahneman, 1974; Viken et al., 2005; Waldmann, 2000 & 2001). The present series of experiments sought to test the assumption of SWS in a task that is less susceptible to the influence of "top-down" factors. An information processing task (i.e., the correlated flankers task) was used so that human participants were completing a central task (i.e., responding to the target) and were unaware as to the importance of the contingencies in the designs, yet were still exposed to them via the irrelevant information (i.e., flankers). Four compound conditioning phenomena were studied in order to test the assumption of SWS. Evidence for the simple predictions coming from SWS theories was mixed. However, a slightly more complex version of these theories can explain the entire pattern of data quite elegantly.

Abstract Approved: ______________________________
Thesis Supervisor

____
Title and Department

____
Date

A FUNNY THING HAPPENED ON THE WAY TO THE MAZE: INCIDENTAL LEARNING OF IRRELEVANT INFORMATION IN HUMANS

by

Rose Halterman Danek

A thesis submitted in partial fulfillment
of the requirements for the Doctor of
Philosophy degree in Psychology
in the Graduate College of
The University of Iowa

July 2010

Thesis Supervisor: Associate Professor J. Toby Mordkoff

Graduate College
The University of Iowa
Iowa City, Iowa

CERTIFICATE OF APPROVAL

PH.D. THESIS

This is to certify that the Ph.D. thesis of

Rose Halterman Danek

has been approved by the Examining Committee
for the thesis requirement for the Doctor of Philosophy
degree in Psychology at the July 2010 graduation.

Thesis Committee:

J. Toby Mordkoff, Thesis Supervisor

Eliot Hazeltine

Shaun Vecera

Edward Wasserman

Thomas Schnell

To Zach, Mom, Dad and Andy

ACKNOWLEDGMENTS

The author wishes to acknowledge the invaluable help given by her thesis advisor, J. Toby Mordkoff, her dissertation committee members, Eliot Hazeltine, Shaun Vecera, Edward Wasserman, Thomas Schnell, and various readers of the many drafts of this work. She would also very much like to acknowledge the help of the undergraduate research assistants who aided in the collection of this data: Jordan Bierman, Lauren Del Boccio, Victoria Lin, Kaitlyn McCoy, Chris Oehler, Lauren Peche, and Tana Truelove.

TABLE OF CONTENTS

LIST OF TABLES

Table

LIST OF FIGURES

Figure

CHAPTER 1

BACKGROUND AND SIGNIFICANCE

The world is full of information. This information can be relevant to a particular goal (i.e., task) or it can be irrelevant. For example, information relevant to a conversation would be the words emanating from your conversational partner’s mouth, while irrelevant information might include the sound of the traffic outside. In order to keep up your end of the conversation you must be processing the relevant speech sounds coming from your partner, while it is not necessary (and may even prove problematic) to process the traffic noise. However, both types of information- relevant and irrelevant- can have relationships with other events or stimuli in the environment.

It has been proposed, and indeed there is a bevy of support for, the idea that contingency sensitive mechanisms exist and allow organisms to recognize associations between events (e.g., stimuli; an event and an outcome) that are either relevant or irrelevant to the task at hand. Organisms such as the *Aplysia* (or marine snail) exhibit the ability to learn associations between stimuli, for example, a tactile stimulus and a shock or the association between a biting response and esophageal stimulation (e.g., Brembs, Baxter & Byrne, 2004; Brembs, Lorenzetti, Reyes, Baxter & Byrne, 2002). These contingency mechanisms have long been studied in non-human animals and humans alike (see Schwartz, Wasserman & Robbins, 2002, pp. 10-22).

What is the nature of learning? One could argue that the process of learning is recognizing associations between events or between events and their outcomes in the environment. The ability of an organism to predict an outcome is a powerful tool for survival. The ability to do this type of learning of associations is critical to an organism’s skill set. For a relationship between events in the environment to exist, it requires that there is a contingency between these events. For example, when a particular bird song is heard, it signals that it is time to mate (Eriksson & Wallin, 1986). Being able to detect the contingency between the sound and the ability to mate increases the chances that a bird that has learned this contingency will pass on its genetic material. On the other hand, the bird's song might also signal a predator that it's a perfect time to eat. Regardless of the particulars of the contingencies present in the environment, organisms have the ability to pick up on contingencies in the environment and use this to their advantage; even in the case of irrelevant information which may not be related to the particular task at hand, but which occurs in the environment and is a reliable predictor of another event. The type of unintentional learning that is a result of other activities is often referred to as "incidental learning"- the learning of irrelevant (or relevant) information that is not intentionally trained. In research this might occur when a participant is unaware of contingencies in an experimental design, but behavior on the task changes as a result of the contingencies present.

Humans exhibit the ability not only for incidental learning, but for more complex forms of learning that involve scenario prediction, the ability to problem-solve using novel instructions, and the use of insight. The current research examines this first type of learning, learning supported by contingency sensitive mechanisms, rather than other types of learning. The incidental learning of irrelevant information provides a unique opportunity to study in humans principles from animal learning paradigms. Understanding this type of learning is a tractable problem that is well-modeled by animal analogues, and has many applications.

Real World Applicability

Having established that the contingency mechanisms which underlie the ability to detect associations between events in the environment are critical to an organism's success, one can turn to the "real-world" applications of these contingency mechanisms (or learning). Two such examples are the ability to categorize objects in the environment, and language learning. It behooves an organism to develop a category for the objects it encounters because when a novel object is encountered and categorized all the knowledge about objects in that category can be applied to the novel object (Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976). Language learning is largely based on trial-and-error or recognizing relationships between sounds that are uttered and their outcomes. In fact, one of the most powerful applications of this type of learning (i.e., delta rule learning which will be discussed in further detail in later sections) was used to model how humans learn language (Rumelhart & McClelland, 1986). In the cases of both category learning and language learning the person makes an initial "guess" as to what the item to be categorized is or what the utterance being heard stands for, and then receives feedback which either strengthens this association or weakens it. In this way, learning about the environment takes place.

The use of contingency information in the environment, and therefore the study of it, has many applications for how we design human-computer interfaces. Computers have been designed to recognize human faces, emotions, and intent; and to aid in efficient human-computer interaction in a variety of areas including the medical, military, and consumer fields (for review see, Sebe, Cohen, Cozman, Geveres & Huang, 2005). Using learning algorithms which incrementally change the *a priori* identification of a face, emotion, or even skin consistency to more closely match the training input is the cornerstone of this type of research. This type of learning is what is being modeled by learning theorists and tested in these real world applications (Sebe et al., 2005). The ability for the computer software to use contingency information from the inputs in the environment is what allows the building of better human-computer interfaces which can allow the automization of many tasks by a computer (Sebe et al., 2005).

Psychologists should be concerned with contingency mechanisms as well. Methodologically speaking all the information in an experiment even information deemed irrelevant by the researcher can be processed by a learning mechanism and this can change responding in a given "higher-order" experiment (i.e., an experiment purporting to study purely cognitive mechanisms). The participant will detect on any contingencies in the experimental design and will be influenced by these such that the phenomenon that is being studied might be altered by the presence of experimental contingencies. It is important to keep issues of contingency learning in mind when designing experiments because contingencies can change how tasks are approached by participants and how results are interpreted by researchers (e.g., Mordkoff & Yantis, 1991; Mordkoff, 1996; Mordkoff, 1998; Mordkoff & Barth, 2001).

Furthermore, research into a learning mechanism can change the questions being asked in other, “cognitive” experiments because depending on the nature of the question different mechanisms might be used for processing. For example, research questions that appeal to explicit cognitive mechanisms which concern information based on arbitrary, instruction-based relationships, or questions that engage incidental learning mechanisms which are based on what contingencies exist in the experiment. The research question and the dependent variable are important as these can influence the interpretation of a results. There are different circumstances that lead to the proposal of a cognitive mechanism versus a contingency sensitive mechanism to explain a pattern of data.

Additionally, contingency mechanisms (as studied, for example, in classical conditioning) can inform us about how organisms represent and organize their worlds. This type of conditioning often reveals what information organisms find to be biologically relevant, how this information is organized into natural categories, and the level at which information is processed. This type of information has ramifications beyond just methodological concerns because it can guide researchers working at the level of perception and action to further develop theories for how organisms organize their worlds.

Finally, an understanding of how contingency mechanisms work has many clinical applications. Contingency mechanisms have been shown to be relevant not only to species-typical behavior, but to abnormal behavior as well. The clinical applications of contingency mechanisms have proven to be both problematic, in that they can produce atypical behaviors, but also therapeutic, in that they have been successfully used in the treatment of said behaviors. One demonstration of the effect of these behaviors is that the pairing of two stimuli can produce life-long phobias in people (e.g., Watson & Raynor's (1920) study of "Little Albert"). In the case of Little Albert, experimenters paired a white rabbit with a loud, startling noise and conditioned Albert to fear the rabbit. Though there has been some debate as to the ability of all phobias to be explained in this manner, demonstrations like Little Albert's are powerful illustrations of how contingency mechanisms might drive the acquisition of phobic behavior (see Thomas & O'Callaghan, 1981 for discussion).

Another demonstration of how contingency mechanisms affect atypical behaviors has been studied in the realm of drug tolerance and addiction. Siegel and Ramos (2002) have proposed that often drug overdoses are the result of a change in the typical environment of the drug user. They propose that environmental cues become conditioned for the drug user (through Pavlovian conditioning which will be discussed in more detail elsewhere). If the environmental cues are consistent they create a typical tolerance level for drugs used in that environment, as well as induce an actual physiological change prior to the drug use. When a drug user overdoses it is often in a new environment and one proposed explanation for why this occurs is because the user's body does not have the conditioned environmental cues that change the body's physiology prior to the drug use so that even a typical dose of drugs in these new surroundings can cause the user to overdose (Siegel & Ramos, 2002). The finding of environmental cues being linked to the presence of a drug has also been found in laboratory experiments with rats. As a natural response to alcohol a rat's body temperature will drop. If a rat has become tolerant to alcohol the natural reaction of the body is to increase its temperature in order to compensate for the coming drop in temperature due to the alcohol. Rats who have repeatedly received alcohol in the same environment will come to show an automatic increase in body temperature in response to an inert substance administered in that environment. Presumably the increase in temperature is because the rat has been conditioned to expect alcohol based on the environmental cues (Siegel, 2005).

The fact that even simple pairings of stimuli can have profound effects on behavior which lead to serious clinical disorders or physiological changes in response to environmental cues supports the idea that some contingency mechanism(s) must be underlying the development of these atypical and often destructive behaviors. On the other side of the coin, clinicians have taken advantage of these contingency mechanisms in order to treat a variety of psychological problems, most notably fear and anxiety disorders which seem to respond well to this type of treatment. Classical conditioning (via a contingency mechanism) in this context can be therapeutic and result in the correction of distressing behavior. For example, a technique known as systematic desensitization can be applied to treat phobias. This technique consists of pairing a previously aversive stimulus (e.g., snakes) with a positive stimulus (e.g., food) thus decreasing the fear associated with the previously aversive stimulus (e.g., Wolpe, 1958; Cover Jones, 1924a, 1924b as cited in Thomas & O'Callaghan, 1981). Similarly, behavioral therapy has a number of instantiations all of which appeal to the pairing of stimuli in order to evoke a change in an undesired behavior. Another example of this technique is the pairing of deep breathing and other relaxation techniques with feelings of anxiety in order to eventually "automatically" trigger a relaxation response when anxiety-provoking thoughts are present (Benson, Beary & Carol, 1974). There has been additional research to suggest that breaking the environmental cues associated with drug or alcohol use can help to reduce withdrawal symptoms (physiological changes) and can help with reducing or eliminating addiction (e.g., Siegel & Ramos, 2002; Siegel, 2005).

Learning Phenomena

There has been a great deal of study of contingency sensitive mechanisms. Work with conditioning in animals has produced a wide variety of learning phenomena that have been the catalyst for the creation of learning theories which seek to provide rules governing when learning will take place, and ultimately seek to characterize the mechanism of learning. Learning theorists use observations of the basic learning phenomena as a means on which to base the structure of a learning mechanism. A good theory of learning will therefore be able to explain the learning phenomena, predict the conditions under which they should and shouldn't occur, and predict new patterns of data. The goal of the learning theorist is to make explicit the underlying psychological principles that bring about the observed phenomena of learning. The goal of the present work was to apply some of these learning theories to determine if the mechanism underlying these observed learning phenomena is in fact the same one that is underlying the observed incidental learning elicited by the present information processing task.

The following are the phenomena or conditions of learning that will be used to test if the mechanism underlying the current task which seeks to examine incidental learning is the same as the generalized learning mechanism studied in the animal literature[1]. In order to provide a context for the more complicated

---

[1] An important point to consider when looking at the current research is that although it was originally motivated by the similarities between the incidental learning exhibited by participants in the current research paradigm and classical conditioning, the current research (as so much research in this area) is seeking to understand associative learning and does not have to be tied to a particular demonstration of either classical or operant conditioning. Classical conditioning involves the association between stimuli (or events), while operant conditioning involves the relationship between an organism's behavior and an outcome. The two types of conditioning can be better characterized as the physical conditions necessary to elicit the phenomena, not the theoretical accounting. This task could be modeled using a variety of learning theories, just as those theories can account for both classical and operant conditioning. In general, classical conditioning paradigms are used to study associative learning,

phenomena that will be discussed, the basic foundation of conditioning is presented. Acquisition is perhaps the simplest of all the learning phenomena and can be defined as the learning of an association between stimuli or events as in classical conditioning or the association between a behavior (event) and an outcome as in operant conditioning (Schwartz et al., 2002). If the unconditioned stimulus is not paired with the CS, as when the CS is then presented in the absence of the US, the association will begin to weaken and the CR will become less frequent until it ceases to occur in response to the CS. This phenomenon is called extinction (Schwartz et al., 2002).

Acquisition and extinction involve the conditioning of single CSs, but it is also possible to obtain conditioning with compound CSs. Each of these more complicated phenomena will be discussed in greater detail in their respective chapters, but for the purposes of setting the stage for the methodology being used are discussed briefly here. Overshadowing is an example of compound conditioning in which there is a reduction or elimination of the strength of the association of a weak stimulus when it is presented simultaneously with a strong stimulus (compared to the strength of the individual association when trained alone). The terms weak and strong are often subjective for different species of animals and the determination of the amount of association they can support is

---

but originally associative learning theories were established for operant conditioning (Thorndike, 1898). Classical conditioning became more popular because it was more tractable, but both can be explained with associative theories (Pearce & Bouton, 2001).

often defined in terms of psychological salience which is a somewhat circular definition, but for the point of demonstrating overshadowing, one of the stimuli should show evidence of little or no conditioning, while the other stimulus should retain most of its associative strength (Schwartz et al., 2002). Blocking can be thought of as a special case of overshadowing in which prior experience plays a key role. Blocking is when a stimulus gains more associative strength by virtue of being trained alone before being trained in a compound with another stimulus. In this case, the stimulus that was given training alone prior to being paired in the compound exhibits evidence of conditioning, while the second stimulus does not (as compared to a condition in which both are trained separately) (Schwartz et al., 2002).

Further compound conditioning is elicited by two additional paradigms. Overexpectation is said to occur when two stimuli are each trained (on separate trials) until they reach a full strength association. Then the two stimuli are tested as a compound pair and it has been found that their association is weaker as the pair compared to when they occurred separately, even though this association was always reinforced (Schwartz et al., 2002). A final phenomenon that will be examined is that of conditioned inhibition which occurs when one CS is paired with a US when presented as a single stimulus, but then not paired with the US when presented in a compound with a second stimulus. In this case, the second stimulus acquires a negative association or the opposite association from the first CS (Schwartz et al., 2002).

Models of Learning

One proposal of this work is that the contingency mechanisms underlying an information processing task in which irrelevant information affects behavior can be explained using the same mechanisms which are proposed to explain models of animal learning. Many models are based on the delta learning rule or error-based learning. This means that the system starts with an initial prediction about the nature of the relationship between a cue and an outcome (or stimuli), then observation of the relationship takes place and the prediction is altered in order to become more indicative of the true observed relationship. The delta (or change) portion of this rule comes from the continual adjustment of the associative weights between events, until the system is able to predict correctly the relationship between events (or objects). An equation that captures this rule is $net_i = \sum_j a_j w_{ij}$ in which this relationship is described. Essentially, the net input to the output $unit_i$ is a summation of the products of the activation of the input from $unit_j$ and the weight between $unit_j$ and $unit_i$. This allows changes in the relationship of $unit_j$ and $unit_i$ until threshold is reached and learning asymptotes (Rumelhart & McClelland, 1986).

We now turn to some of the learning theories which have been proposed to explain the mechanism underlying learning. Though all of the theories that have been proposed will not be presented here, the major highlights will be presented in chronological order to give evidence of the scope of this research, and to set the stage for how the present research might be integrated into the large body of research concerning a generalized learning mechanism.

Historically, the first theory for a learning mechanism came from Pavlov. Pavlov's proposal for the explanation of a learning mechanism involved the simple principle of temporal contiguity. In this view the only necessary conditions for the eliciting of a CR was that the CS and the US appeared closely together in time (Pavlov, 1927). Contiguity theories like Pavlov's were proposed to explain how acquisition occurs and indeed were successful in doing so. However, contiguity theories lack the ability to explain other learning phenomena, such as those that involve compound cues. For example, it is hard to see how a theory of simple temporal contiguity could explain the observed effect that if a stimulus (CS_A) is presented in the presence of another more salient stimulus (CS_B), the first stimulus (CS_A) will show virtually no conditioned responding (i.e., overshadowing). Additionally, contiguity theories have difficulty explaining blocking in which previous experience with a stimulus (CS_A) will interfere with subsequent training of another stimulus (CS_B). According to contiguity theory, the mere fact that the CS-US pairings have occurred should be enough to elicit the conditioned response.

The inability of contiguity theories to explain many of the observed learning phenomena led many to abandon it in favor of a theory that included the notion of contingencies. In this conceptualization the learning occurs not just based on what events are occurring, but takes into account what events are not occurring. Essentially, the circumstances necessary for conditioning to occur involve the probability of a US given the absence or presence of a CS. A formulation of this is given by: $\Delta P = P\,(US \mid CS) - P\,(US \mid No\ CS)$, which states that the change in probability (and thus the contingency) is the difference between the probability that the US will occur given the CS minus the probability that the US will occur given that the CS has not occurred. To the extent that there is a greater probability of the US occurring given the occurrence of the CS there will be excitatory conditioning. If the probability of the US occurring in the absence of the CS is higher there will be inhibitory conditioning. If the probability of the US occurring is equally likely given the absence or the presence of the CS, then no conditioning will result (Rescorla, 1968). Like contiguity theories, contingency theories lack the ability to explain interactions between compound stimuli.

Associative theories of learning were developed primarily to explain the cue competition effects that were observed (e.g., blocking, overshadowing, overexpectation, conditioned inhibition, etc.). Previous theories were unable to explain the various phenomena observed when multiple stimuli were presented. Associative theories derive much of their explanatory power from the assumption that all stimuli presented vie for associative strength, an assumption I am calling, the assumption of Shared Weight Space (SWS; use of this term is discussed on pp. 33 - 34).

There are two classes of associative learning theories which have been proposed, elemental theories (e.g., Mackintosh, 1975; Rescorla & Wagner, 1972; Wagner, 1981) and configural theories (e.g., Pearce, 1987; 1994). Broadly speaking, these two classes differ in the way in which compound stimuli are processed (Mehta & Russell, 2009). Elemental theories propose that compound stimuli are treated as a sum of their parts. For example, the associated strength of compound AB is treated as the sum of the associative strengths of its components (i.e., A + B) (see Rescorla & Wagner, 1972). Whereas configural theories propose that the compound is an entity onto itself, different from its component parts and that the only relation between the compound and the parts is their resemblance to each other (see Pearce, 1987; 1994 for review). That is, AB is an entirely different entity from A or B and the relation between A and B to AB is through any resemblance these component parts may have to the compound as a whole.

While these theories differ on the surface in the way in which compound stimuli are processed, a deeper examination reveals that there is a fundamental similarity between the classes of theories. Both depend on the assumption of Shared Weight Space (hereafter: SWS). Theories that rely on SWS have in common the idea that there is a finite amount of associative strength available for a given unconditioned stimulus (US) and the idea that all conditioned stimuli (CS) that are presented in a conditioning paradigm vie for this associative strength. This assumption that all CSs vie for associative strength (or the ability to evoke a conditioned response) allows both elemental and configural theories to explain many learning phenomena that previous theories, which did not include SWS, could not explain. Most notably, associative theories are able to explain conditioning to compound stimuli composed of two or more CSs.

One of the most influential associative theories, and one that has spawned the greatest amount of research, was proposed by Robert Rescorla and Allan Wagner in 1972 (Siegel & Allan, 1996). Rescorla-Wagner theory is captured by the following equations: $\Delta V_A = \alpha_A \beta_1 (\lambda - V_{AX})$ and $\Delta V_X = \alpha_X \beta_1 (\lambda - V_{AX})$ where a summation of both formulas is necessary to capture a compound stimulus (ΣV). Taking as an example the first formula (and extending this include all stimuli), the formula states that the change in the strength of the association of the CS (ΔV_A) is the product of the salience of the CS (α_A) multiplied by the associability of the US or the learning rate parameter of the US (β_1) multiplied by the difference in the asymptote of learning and the learning that has occurred thus far ($\lambda - V_{AX}$). As can be noted from the addition of the formula for another element in a compound and the notion of the V_{AX}, there is the assumption that the strengths of the associations for the both elements are related to each other. V_{AX} is said to be composed of the strengths of the components in a simple additive fashion $V_{AX} = V_A + V_X$. According to the original version of the Rescorla-Wagner model, all values are bounded by 0 and 1 (Rescorla & Wagner, 1972).

Learning theorists have studied the assumption of SWS in animals in order to gain an understanding of the underlying contingency mechanisms. There is evidence that the same mechanisms studied in animals are present in humans as well (e.g., Arcediano, Matute & Miller, 1997; Chapman, 1991; Chapman & Robbins, 1990; Dickenson, Shanks & Evenden, 1984; Pineño, Denniston, Beckers, Matute & Miller, 2005; Price & Yates, 1993; Shanks, 1985;

Williams, Sagness & McPhee, 1994). Indeed, the parallels between animal conditioning and human learning have been amply demonstrated (see Wasserman, Kao, Van Hamme, Katagiri, & Young, 1996 for review). However, there is some difficulty adapting the conditions for learning used in animals to human paradigms because humans have the ability to use other information in the world on which to base responding.

The Study of Contingency Mechanisms in Human Participants

Assuming Shared Weight Space (SWS) allows for the explanation of many phenomena that have been observed in animal conditioning paradigms. However, the observation of SWS in humans has proven to be difficult given that humans have other processes that are occurring that provide the incredible flexibility that allows humans to follow arbitrary instructions and perform a task correctly on the first try. Previous work has shown that humans do change responses based on contingencies that are present in experiments much like their animal analogues. This finding has been characterized in two different ways.

One example of the way in which contingency information has been studied is as an experimental confound. There are many instances in which the unintentional contingencies in an experimental design have been used by a participant to produce a faster or more accurate response (e.g., Mordkoff & Yantis, 1991; Mordkoff, 1996; Mordkoff, 1998). It has been shown that performance can be affected by unintentional contingencies in experimental designs. Specifically, it was found that the way in which redundant information extracted from two separate sources can come to activate a specific response was not through a summation of the activation of the information presented, as was previously thought, but through simple contingencies activating associations. One theory to explain the way in which divided attention operates in order to extract information from multiple sources was that of coactivation. According to this theory information from multiple objects pooled its activation to produce a faster response for redundant targets (i.e., the two objects presented were both targets) than single targets. However, when the ability for the system to form associations was decreased (i.e., no contingencies were present- all pairings of the redundant information were equally likely) then the theory of how divided attention operates was changed (see Mordkoff & Yantis, 1991).

It has also been shown that work with "pre-pulse inhibition" which led many to propose a theory about sensory gating and attentional capture could also be explained by unchecked contingencies in the experimental design (Mordkoff & Barth, 2001). The "pre-pulse inhibition" is the finding that the involuntary eye blink that occurs when a participant receives a tap on the glabella is reliably decreased in intensity when a tone (or less intense stimulus) is presented just prior to the tap. In the literature this finding was related to the idea of a form of attentional capture – which proposed that the tone captured attention, making less resources available for switching attention when the tap occurred, such that "sensory gating" of the tap information occurred and reduced the involuntary response when preceded by a warning. The temptation then was to use to the magnitude of the pre-pulse inhibition as a measure of attentional capture (Mordkoff & Barth, 2001). However, Mordkoff and Barth (2001) found that in all the studies reporting pre-pulse inhibition there was a perfect contingency between the tone and the up-coming pulse which meant that the contingency information was confounded with the phenomenon of interest and that an alternative explanation for the reduction in response could be conditioning rather than a model of sensory gating. When additional tones were added that served to weaken the contingency between tones and the tap to the glabella it reduced the observed inhibition. The weakened correlation didn't completely remove the effect of the tones on the startle response, but the automatic component (i.e., that which was proposed to be due to attentional gating not due to conditioning) was much smaller (Mordkoff & Barth, 2001).

Another example of contingencies influencing cognitive tasks is found in the manipulation of interference effects. The term "interference effect" refers to the finding that information that is irrelevant to making a response can come to affect responding to relevant information. One example is the Stroop task. The Stroop task requires reporting a single dimension of a multidimensional stimulus. In order to be successful at this task the irrelevant stimulus dimension(s) must be ignored. The typical Stroop effect is the finding that if the irrelevant stimulus dimension conflicts with the reported dimension, then response time (RT) is slower than when the irrelevant stimulus dimension does not conflict or is compatible with the reported dimension. For example, in a task in which participants must respond by naming the color of a colored stimulus, if the colored stimulus is a color word that is incompatible with the to-be-named color then RT is slower than if the color of the word and the form of the word are compatible. That is, participants would be faster to respond to the word "blue" printed in blue ink than if it was printed in red ink (Stroop, 1935/1992). This is also true for spatial words and locations as well. If the word "above" is presented above an imaginary horizon it is faster to be responded to than if it appears below the horizon (Palef & Olson, 1975). The accepted explanation for this finding is that this task is performed by attention being directed to the relevant dimension. The slowed response time arising from the conflict between the dimensions is seen as a failure of selection attention or as being due to the automatic processing of the irrelevant dimension (see Logan & Zbrodoff, 1979 for discussion).

There is, however, an alternative (or at least additional) explanation for the Stroop effect. Evidence for the alternative explanation comes from experiments in which the number of incompatible trials was increased. The magnitude of the Stroop effect was found to be altered by the inclusion of more incompatible trials relative to compatible trials. The compatible stimuli were processed faster when incompatible trials were rare, but when incompatible trials were frequent incompatible stimuli were processed more quickly (Lindsey & Jacob, 1994; Logan & Zbrodoff, 1979). One possible explanation, and one that is put forth by some researchers, is that there are two components to the processing of the irrelevant dimension of the Stroop stimuli. One component is automatic, while the other is an attentional component that is strategically varied by the participant (Logan & Zbrodoff, 1979). This finding suggests that there is at least one mechanism that is not cognitive in nature, but is instead conditioned. The alternative explanation for the change in the magnitude of the Stroop effect is that the participant learns through an associative mechanism to ignore information coming from the irrelevant dimension. This is similar to work done with priming. Priming refers to the fact that people are faster to respond to a word given that a related word has been previously presented rather than an unrelated word. For example, a faster RT for responding that "nurse" was a word would be observed given that "doctor" was previously presented. Whereas no advantage would be observed for "nurse" given that "bird" was previously presented. However, it has been found that participants can be primed for a target word by an unrelated category if the unrelated category and the target category have been correlated in the experiment (Neely, 1977). These findings suggest that there is more to the interference effects observed than just an appeal to cognitive resources. There appears to be an associative process that underlies these effects as well.

While, contingency information has been studied as a confound in previous research with humans, contingency mechanisms have been studied in animals for decades (e.g., Pavlov, 1927). The assumption underlying this research is that the same mechanisms operating in animals are operating in humans as well. These parallels between the learning observed in animals and humans can provide evidence of a generalized learning mechanism that is demonstrated to be operating in animals as well as humans. However, a generalized learning mechanism that operates on associations between stimuli might not be the only mechanism responsible for the effects observed in human contingency learning. It has been argued that contingency learning, as evidenced by human participants, might depend on "both rational reasoning and simple associative mechanisms" (De Houwer & Beckers, 2002, p. 306). The authors argue that people are able to make rational inferences about contingencies and causality. However, they argue that rational influences will only be observed if participants have motivation and opportunity to use those (De Houwer & Beckers, 2002).

Given that humans can (and often do) use reasoning or outside knowledge to make contingency judgments it is often difficult to study contingency learning mechanisms that have been studied in animals using human participants. The difficulty arises because humans have a variety of "top-down" factors which can interfere with the study of contingency mechanisms (e.g., Alloy & Tabachnik, 1984; Brewer, 1974; Chapman & Chapman, 1967; Smedslund, 1963). Human participants bring a large amount of prior knowledge, biases, and expectations to bear on situations in which contingency information is also present. There have been many demonstrations of humans responding to contingency tasks in ways that are not commensurate with the true contingencies present in the task. This is due in large part to the operation of explicit mechanisms which allow humans to make arbitrary responses to stimuli on the basis of instructions alone. It has been pointed out that there are a large number of ways in which the explicit nature of a task fails to produce the same types of data that we would except, and would need evidence for, in order to say that an incidental contingency mechanism is operating in humans (see Brewer, 1974 for discussion). There is also a large body of work that shows how human participants' knowledge about the structure of the world, personal biases, and heuristics based on experience can change how contingency information is used (e.g., Chapman, 1991; Msetfi, Murphy, Simpson & Kornbrot, 2005; Perales, Catena & Maldonado, 2004; Tversky & Kahneman, 1974; Viken, Treat, Bloom & McFall, 2005; Waldmann 2000; Waldmann, 2001). In fact, this "cognitive" information can influence responding to the true contingencies inherent in the tasks.

Evidence that participants' knowledge about structured relations in the world plays a role in responding in human contingency learning tasks comes from studies in which differences in the perceived relationships of the experimental stimuli change the perception of the contingencies present in the task. Waldmann (2000; 2001) presented participants with identical stimuli (hereafter: cues), outcomes and feedback information, and only varied the instructions given to the participants in order to manipulate whether the cues were predictive or descriptive of the outcomes. For example, participants in one experiment were presented with information regarding the presence or absence of a fictional substance in the bloodstream of some animals and given information about a fictitious disease that was related to these substances. In the *predictive condition* participants were told that the substance(s) were causes of the disease. In the *diagnostic condition* participants were told that the substances were the effects of the disease. In this task, in which all the factors were identical except for the characterization of the cues as either causes or effects, different results were obtained. In a paradigm which would typically produce conditioning effects in animals (e.g., overshadowing and blocking), it was found that these effects only occurred when the cues were predictive, but not when they were diagnostic (Waldmann, 2000; 2001). This is in contrast to the predictions of associative models of animal contingency mechanisms which would predict that these effects would occur in both conditions given that the cues, outcomes and correlations between the cues and outcomes were identical between the conditions (but, cf. Van Hamme, Kao & Wasserman, 1993). These experiments demonstrate the effect of beliefs about causality on conditioning in human participants.

Another demonstration of the power of causal knowledge to affect contingency ratings comes from Perales and colleagues (2004). They conducted a study in which a cue (i.e., substance in the blood) and an outcome (i.e., fictitious disease) were trained in a single training phase. In the next training phase, a different cue (i.e., different substance in the blood) and the same outcome (i.e., same disease) were trained. In each phase, participants provided ratings of the causality between the cue and outcome. Then at the end of the experiment participants were given the unexpected task of rating the contingency between the cues. As in the studies by Waldmann (2000; 2001), there were two conditions used in the task, one was a *predictive condition* and the other was a *diagnostic condition*. Critically, participants reported a correlation between the cues (which had never been paired in the experiment) if the cues were characterized as effects of a common cause, the *diagnostic condition*, but did not report the correlation between the cues if they were characterized as causes of a particular effect, the *predictive condition* (Perales et al., 2004). Interestingly, this finding of an illusory correlation between cues did not occur for all participants. One-third of the participants did not give an illusory correlation rating for the two cues. The authors argue that this subset of participants exhibited the type of behavior predicted by associative models. That is, they were sensitive to the true contingencies present in the experiment (Perales et al., 2004).

Evidence that human participants are differently biased in their causal judgments based on “outside” knowledge comes from a variety of sources. In Chapman’s (1991) experiment in which participants were presented with fictitious symptoms (cues) and then asked to rate the ability of these cues to predict a disease (outcome) she observed that participants rated two symptoms significantly differently even before these cues had been presented in a training phase. That is, the cues had different predictive strengths based on some factor other than an experimental manipulation. She attributed this to a pre-experimental bias that existed (Chapman, 1991). Although she controlled for this bias in the subsequent analyses, the fact that it existed speaks to the notion that in human causal judgment tasks, different amounts of “outside” knowledge might be utilized by the participants. In another experiment, it was found that participants spent more time considering information that was against their initial biases and less time considering information that confirmed an existing bias. These biases were directly assessed in the experiment as well and were found to exist (Levin, Wasserman & Kao, 1993). While these biases can be controlled for when analyzing data from these types of experiments, they are not always assessed in every experiment and their existence is problematic for making claims about human contingency judgments.

Though one might argue that other species have the same types of biases to respond to given stimuli (i.e., biological constraints on learning), these patterns of responding are generally species-wide and consistent for each organism. For example, Garcia (1966) found that it was easy for rats to pair a given taste with illness, but that pairing a taste with shock was difficult to learn. Conversely, rats show rapid conditioning to a pairing of a light and shock, but do not learn the light and illness association, a finding which has been replicated many times (Garcia & Koelling, 1966). However, this finding is not particular to a given rat, but is instead a general "principle".

In work with humans, participants have demonstrated pre-experimental biases toward stimuli, reminiscent of the "Garcia effect" in animals, which might lead one to conclude that this type of bias in humans is the same as the bias sometimes demonstrated with animals. However, the types of biases occurring in these human participants are not unconditioned biological responses which have a greater likelihood of being shared across individuals of a species. Instead, these biases are acquired through experience and knowledge about the world which might vary between participants. Additionally, the extent to which participants are using this knowledge or being influenced by it in a given experiment is likely to vary between participants as well.

Additional evidence for the idea that biases can play a role in casual judgment tasks comes from work by Alloy & Abramson (1979). In this study, participants were asked to rate the degree of contingency between their response and an environmental outcome. It was found that depressed participants accurately judged the contingencies in all of the experiments. Whereas depressed participants were accurate in their assessments of contingencies, non-depressed participants overestimated the contingency between their response and an outcome in situations in which the relationship was non-contingent, but the outcome was frequent or desired. Non-depressed participants also underestimated the contingency between their response and an outcome when the outcome was undesired (Alloy & Abramson, 1979).

Interestingly, these findings have been disputed in more recent years. It has been shown that the differences between depressed and non-depressed participants are not that depressed participants are better at detecting contingencies, but rather they tend to discount information that generally leads to overestimation of a contingency. Specifically, it has been proposed that participants might be processing the information coming from intertrial intervals (ITI) in experiments. This can be likened to a "no-response/no-outcome" trial which serves to shift the balance of a probability estimation to be positive (see Msetfi et al., 2005 for discussion). One idea about why depressed participants produce less skewed estimations of contingencies in experiments with ITIs is that they are prone to rumination and often display a lack of attention which would allow them to more easily discount the information coming from the ITIs and thus be less influenced by it (Msetfi et al., 2005). Another theory for why depressed participants have been shown to exhibit a more accurate prediction of an experimental correlation in some cases, and not in others, is that depressed participants have a higher threshold for providing a yes response. That is, depressed participants have a change in the threshold at which they will make a response compared to those who are non-depressed, which in some cases can look like depressed participants have a more accurate representation of the contingencies of an experiment when in fact they just need more information that non-depressed participants in order to make a response which leads to less overestimation in most situations (see Allan, Siegel & Hannah, 2007 for discussion). However, the argument about whether depressed or non-depressed participants are more accurate at using information is not of central importance to the point being made herein. Instead the value of this research is to show that there are biases in the information that people use in these causal judgment tasks.

Further evidence of biased responding in human contingency judgments comes from work by Tversky & Kahneman (1974) which suggests that people are biased to respond based on heuristics and stereotypes, and often ignore the true contingencies present. One such example comes from an experiment in which participants were asked to judge the personality profiles of 100 professionals. In one condition, the fictitious group of professionals consisted of 70 engineers and 30 lawyers. In the other condition, the group consisted of 30 engineers and 70 lawyers. Therefore the odds that a particular description belonged to an engineer would be higher in the first condition than in the second. However, participants in both conditions produced essentially identical probability judgments. The authors conclude that these ratings were based on stereotypes about the personality of engineers which were stereotypically described in the passages rather than the experimentally defined probabilities (see Tversky & Kahneman, 1974 for an extensive review of the use of biases and heuristics in judgments).

Recent work into the clinical applications of cognitive measures also suggests that people can be insensitive to the true contingency information in a design. Viken and colleagues (2005) have found that when participants are asked to make contingency ratings for a group of materials (e.g., judging the correlation between affect and body size) they tend to overestimate the contingency information present in the design based on cultural biases. In a set of materials in which the true correlation between affect and body size was zero, participants consistently described a negative correlation between the variables consistent with the cultural bias in the American population to ascribe positive affect to thin people (Viken et al., 2005).

These findings demonstrate how "outside" cognitive factors can influence contingency ratings in human participants. The existence of these top-down influences makes it difficult to study the contingency mechanisms that we believe exist in humans as they do in animals. What is needed is a way to study the contingency mechanisms in humans using a task that is less susceptible to influence from explicit cognitive factors. Fortunately there is a way in which we can engage these explicit mechanisms and study contingency mechanisms, and that is to use the correlated flankers task.

The use of the correlated flankers task in this line of research is an attempt to serve two purposes. First, using this task can provide an in-depth examination of the mechanism underlying the processing of irrelevant task information that has a contingent relationship with the response. Secondly, the use of the correlated flankers task can serve as a bridge between bodies of knowledge to be able to gain greater understanding of the learning mechanisms which have been proposed to exist in animals and humans alike using a task that humans do not recruit top-down resources in order to complete.

A New Approach: The Correlated Flankers Task

Miller (1987) introduced a task called the correlated flankers task. In this original version of the correlated flankers task, visual stimuli (i.e., white letters) were mapped onto manual response buttons (i.e., keyboard buttons). The correlated flankers task is by definition a selective attention task because there is one item to which attention is to be directed and other items which are to-be-ignored. A selective attention task is one that involves a target-defining attribute that indicates which item or items to attend to, and a reported attribute on which the response is to be based. The target-defining attribute in the correlated flankers task is location (i.e., respond only to the center item), while the reported attribute is identity (i.e., the shape of the letter). In Miller's (1987) task, a trial began with the presentation of a large, white circle in the center of the computer screen. After a delay, the white circle would disappear and the trial display would appear. A display consisted of three stimuli: a target and two flankers. The flankers (i.e., other white letters) were always identical to each other, but different from the target. Each target had a corresponding response assigned to it. Three targets were assigned to the right button press and three targets were assigned to the left button press. This is commonly known as a 6:2 mapping because there are six targets mapped onto two buttons (three targets on each button). The participant's task was to respond to the target (the center item) with the correct button press (depending on the instructions), while ignoring the two flanking stimuli. The participant is told in the instructions to only respond to the target item and to ignore the "irrelevant" flanking stimuli. Response time (RT) is the measure of interest in this experiment.

The targets in this paradigm have an instruction-based relationship with the response. That is, the targets are arbitrarily assigned to the responses via the instructions. In order for the perception of the target to be turned into an overt behavior (i.e., the button press), the stimulus (i.e., target) must be translated from its perceptual code to the motor code necessary for response. This process is called Stimulus-Response translation (or S-R translation).

The flankers however, do not have an instruction-based relationship with a given response. Instead, the flankers are other letters that are contingent by virtue of the frequency of a particular response given the occurrence of a given flanker. For example, imagine that the letters **A**, **B**, and **C** are assigned to the *left* response. That is, whenever the target is an **A**, **B** or **C**, the correct response is to press the left button. And the letters **X**, **Y**, and **Z** are assigned to the *right* response. The correlated flankers would be other letters, for example, **D** and **W** that are not assigned by the instructions to any response. By manipulating the frequency with which **D** occurs when the correct response is a left button press a contingency can be set up between the *left* response and **D**. That is, when the response indicated by the target is a left button press there will more often be a D in the display, then *left* response and **D** will accrue a positively contingent relationship. The **D** in the display is often a valid predictor that the correct response will be a left button press. The same will occur for the *right* response and **W**. The finding from experiments of this type is that when the flanker is positively correlated with the response (e.g., **D** is the flanker and the target is an **A**, **B**, or **C**, and the correct response to the target is a left button press), RT is faster than when the flanker is negatively correlated with the response (e.g., **D** is the flanker and the target is a **X**, **Y,** or **Z** and the correct response to the target is a right button press) (Miller, 1987). This finding is somewhat surprising given that participants are told to ignore the flankers, and the only relationship the flankers have with the response is a contingent one. This experimental design is evoking incidental learning about the flankers and their relationships to the responses.

Preliminary evidence for the ability to study contingency mechanisms (like those demonstrated to underlie animal conditioning) using a task which engages and distracts the explicit mechanisms has been obtained using the correlated flankers task. The evidence that the correlated flankers task is being processed by incidental contingency mechanisms and not the typical explicit information pathway that arbitrary, instruction-based information is processed with comes from two sources. First, participants in a correlated flankers task are unable to report the contingency relationships for the flankers and responses although they show a behavioral effect (Miller, 1987). At the end of his experiments, Miller (1987) probed participants for verbal reports about the flankers. He found that while RT was affected by the contingency relationships between the flankers and responses, participants were often unable to report the identity of the flankers, the frequency of their occurrence, or their relationship to the response (Miller, 1987).

Additional evidence that the correlated flankers task involves a different mechanism than the one that has been studied traditionally comes in the form of a double dissociation between the standard flankers (an arbitrary, instruction-based task) and the correlated flankers task. Miller's (1987) correlated flankers task is an adaptation of Eriksen and Eriksen's (1974) "standard" flankers task. The standard flankers task is similar to the correlated flankers task. Similarly to the correlated flankers task, the standard flankers task involves the presentation of a target and two flankers on every trial. The target-defining attribute is location (again, the center item is to-be-responded-to), while the reported attribute is identity (i.e., shape of the letter). Unlike in the correlated flankers task in which the flankers never share identity with the target, standard flankers share their identities (reported attribute) with the targets (i.e., targets and flankers are the same letters). The only thing that differentiates the target from the standard flankers is the target-defining attribute of location. For example, if the targets are **A**, **B**, **C**, (mapped onto the *left* response) and **X**, **Y**, **Z** (mapped onto the *right* response) then the flankers would also be **A**, **B**, **C**, **X**, **Y,** or **Z**. The participant is told to respond to the target which is always the center item, and to ignore the items located to the left and right, the flankers. The general finding is that when the target and the flankers are mapped onto the same response (e.g., **ABA** or **BBB**), participants are faster to respond than when they are mapped onto opposite responses (e.g., **ZCZ** or **AXA**) – the so-called "standard flanker effect" (Eriksen & Eriksen, 1974).

Two modifications to the experimental procedures described for the standard and correlated flankers produce different results than those that are obtained with the typical designs for those tasks. In previous versions of the standard and correlated flankers tasks, the reported attribute was only a single dimension (i.e., identity; shape dimension). However, the two tasks can be modified such that the reported attribute is a conjunction of two dimensions (e.g., shape and color are necessary to determine the appropriate response). Recall that in both tasks, using a single dimension reported attribute produced a faster RT on trials in which the target and flankers were **congruent** (compared to the **incongruent** trials) (see Kornblum, 1992 for a review on classifications of this nature). That is, for the standard flankers task in a **congruent** trial the target and flankers were mapped onto the same response button. For the correlated flankers task, a **congruent** trial would mean that the target was mapped onto a given response and the flankers presented had a positively contingent relationship with that same response. **Incongruent** trials on the other hand are those on which the target and flankers are mapped onto different responses (standard flankers task) or the target is mapped on to one response and the flankers have a positively contingent relationship with the opposite response (correlated flankers task). The finding that RT is faster on congruent trials than incongruent trials is the typical flanker effect.

However, if the reported attribute needs a conjunction of two dimensions to define the response (e.g., shape *and* color), then standard flankers do not show any evidence of being able to affect RT. In this case, the RT for the congruent trials is roughly equal to (and not statistically different from) the RT for the incongruent trials (Mordkoff & Halterman, 2008). If the paradigm is changed again so that the targets are all the uppercase forms of letters (e.g., **A**, **E**, **R**, and **G**) while the lowercase forms of the letters are the flankers (e.g., **a**, **e**, **r**, and **g**) then the standard flanker effect is still observed. That is, the standard flanker effect crosses case. Even though the letters are perceptually very different (A vs.

a) they are conceptually the same -both **A** and **a** share the same identity (Mordkoff & Danek, in prep).

The opposite pattern holds for the correlated flankers task. Correlated flankers do show effects on RT when they are conjunctions of features (Mordkoff & Halterman, 2008). If the targets are red or green diamonds or squares mapped onto different responses then if yellow or blue pound signs or tilted pound signs are made to have a contingent relationship with the response, RT on the congruent trials is faster than the incongruent trials. However, the correlated flanker effect does not transfer between upper- and lowercase versions of the same letters. When the training of the contingencies is conducted with uppercase letters and then the flankers switch to being lowercase versions, the correlated flanker effect disappears (Mordkoff & Danek, in prep).

The most parsimonious explanation for the double dissociation between standard and correlated flankers is that these tasks are being processed by two different information pathways. One system uses an explicit mechanism and processes arbitrary, instruction-based information, while the other system is conditioned and is sensitive to contingency information. The first (explicit) mechanism has been studied for years, but this second system which is sensitive to contingency information has not been studied in great detail previously.

Overview of the Present Research

The present research is the first to examine the effect of compound flankers on behavior. The main focus of the work is to understand if the mechanism underlying the incidental learning of irrelevant information (as exhibited in the correlated flankers task) is the same as that which is operating in animal (in conditioning paradigms). This will be accomplished using the correlated flankers task to demonstrate analogous phenomena of compound conditioning from animal learning with humans. Each time that demonstrations of the same phenomena are evoked in humans using analogous paradigms from animals, it strengthens the argument that the same mechanism that is operating in animals is operating in humans as well (Williams et al., 1994).

Many contemporary theories can successfully explain conditioning phenomena due to the assumption of SWS. This research uses a task that is less susceptible to top-down influence in order to investigate the proposition that the mechanism underlying incidental learning in humans relies on SWS. Specifically, tests of compound conditioning phenomena will be incorporated into the correlated flankers task in order to supply the evidence needed to understand how issues of SWS are active in humans.

Now that the general conditions for learning have been discussed, the application of how these will be carried out in the information processing task at hand – the correlated flankers task- is in order. In the next chapter, a discussion of the general method that will be employed across all experiments is presented. It is detailed there for orientation purposes. This general method should provide the reader with the building blocks for understanding how each of the more complicated applications of this method is accomplished. One final technical note will be addressed in this introduction that of terminology.

The Use of Terminology

Various terminology may have specific meanings to those within a particular field of study and those outside the particular field. As this paper is an attempt to bridge between two different fields, some discussion of the terminology chosen is in order. There are two cases which deserve particular attention.

First, the term "incidental learning" is used throughout the paper rather than the term "implicit learning". The choice to use incidental learning rather than implicit learning was made because of the underlying assumptions that each term evokes. Implicit learning has been defined as being characterized by two critical features: "(a) It is an unconscious process and (b) it yields abstract knowledge", (Reber, 1989, p. 219). Incidental learning has been defined as the "unintentional or unplanned learning that results from other activities", (Kerka, 2000, p. 3). Some of those who study implicit learning are quite specific about the underlying processes that allow for implicit learning to take place, it is an unconscious process that yields abstract knowledge. This characterization of implicit learning includes a description of the form of knowledge representation as well (for a review, see Reber, 1989). Incidental learning, on the other hand, is a term that describes the parameters necessary for a particular type of learning to take place, learning that is informal and unintentional. The term incidental learning is used throughout this paper because it is less specific about the underlying mechanisms of the learning that is taking place. It is more a description of the type of circumstances that evoke this particular type of learning. Instead of being tied to the implications that come from using the term "implicit", the descriptive term, "incidental" is more suited to the purposes of this paper. The conditions necessary to evoke incidental learning were met in the experimental designs, but a greater commitment to the exact representation of this knowledge and the processes used to acquire it is not assumed.

Although the choice of incidental versus implicit learning was made in an attempt to disengage from the potential assumptions that each term invokes, the second place in which different terminology is used takes the opposite approach. The term "shared weight space (SWS)" is used throughout this paper. It could be argued that this term is used to describe what has previously been labeled elsewhere as “associative (or cue) competition”. SWS is used in this paper because it is a term that carries with it an assumption about the mechanism that allow for multiple stimuli to be processed. Theories that assume, what I am calling, SWS, have at the core this assumption that there is a very specific way in which multiple stimuli interact. Specifically, there is a finite amount of associative strength available and when one stimulus acquires that associative strength the other stimulus must necessarily, and in a *systematic* fashion, lose associative strength (or not acquire it). While the term “associative competition” is descriptive of the events taking place when multiple stimuli are present (i.e., the stimuli compete for associative strength), it does not convey the underlying assumptions about the specific mechanisms occurring. Associative competition describes *what* occurs whereas the assumption of shared weight space is an attempt to describe in more detail the *how*. In this case, SWS is used because it is a term that carries with it assumptions about underlying mechanisms which are being tested in this paper.

CHAPTER 2

GENERAL METHOD

The series of experiments reported use the same general method to study these phenomena. Essentially, a modified correlated flankers task was used, and various conditioning paradigms were incorporated. Evidence of conditioning is measured by changes in response time. The basic method for the series of experiments that will be described herein will be similar across the phenomena under investigation. The details of this general method will be employed throughout, noting any changes in procedures for each individual experiment.

In all the experiments there were three types of blocks: practice, training and test blocks. There were 2 blocks of 48 trials each that were practice blocks. The practice blocks allowed participants to get used to the task and to learn the mapping of the targets-to-responses. Each participant then served in 6 blocks of 96 trials each with a forced mid-block break of 9 seconds that were training blocks and two blocks of 96 trials each that were test blocks. Breaks between blocks were self-paced.

Procedure

For each participant, six letters (**A**, **U**, **X**, **M**, **O**, and **T**) were randomly assigned to be one of six targets. Three of the targets were mapped onto the "go" response – the response which required the participant to press the response button, and three of the targets were mapped onto the "no-go" response to which the participant did not make a button press response.

Practice Blocks. The participants were told that they would always be presented with a target (one of six letters) appearing in the center of the screen, and that they were to respond to this central letter. They were told that there would be other items appearing around the target, but that they were to ignore these items and only respond to the middle letter. They were told to respond as quickly as possible while making few errors. Participants were shown the mapping of the targets-to-responses for 7 seconds. In the practice blocks, only the target was presented.

Each trial began with the presentation of a warning stimulus which alerted the participant that the trial was starting. The warning stimulus was presented for 350 ms and was followed by a blank screen for 150 ms, and finally the presentation of the final display. This display remained on the screen until the participant responded or the 1500 ms deadline had passed (for "no-go" trials), at which time feedback was given only on error trials for 1 second and then followed by the mapping of targets-to-responses for an additional 3 seconds for the first block of practice trials. In the second block of practice trials, only the error feedback was presented for 1 second, but the mapping was not reshown. Feedback was only given on error trials. The next stimulus display appeared about 1.5 seconds after the offset of the feedback for the previous trial. An example of the events of a trial in a practice block is presented below.

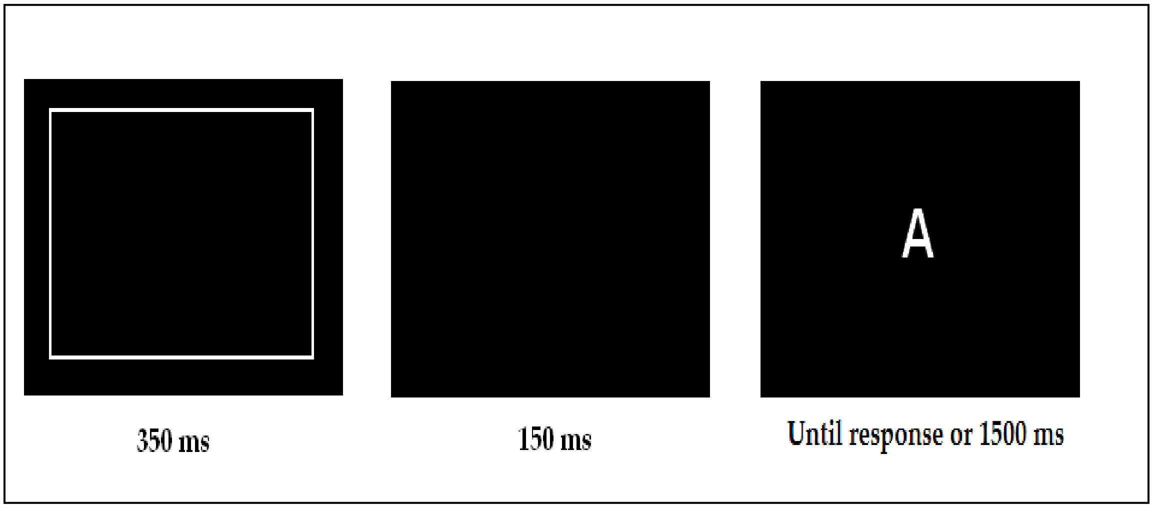

Figure 1. Trial sequence for practice blocks.

Training Blocks. As with the practice blocks, each trial began with the presentation of a warning stimulus which alerted the participant that the trial was starting. The warning stimulus was presented for 350 ms. In the training blocks, the flankers onset immediately after the offset of the warning stimulus, and 150 ms prior to the onset of the target. The final display consisted of the flankers and the target and this final display remained on the screen until the participant responded or the 750 ms deadline had passed (for "no-go" trials), at which time feedback was given only on error trials for 1 second. The next stimulus display appeared about 1.5 seconds after the offset of the feedback for the previous trial. An example of the events of a trial in the training blocks is presented below.

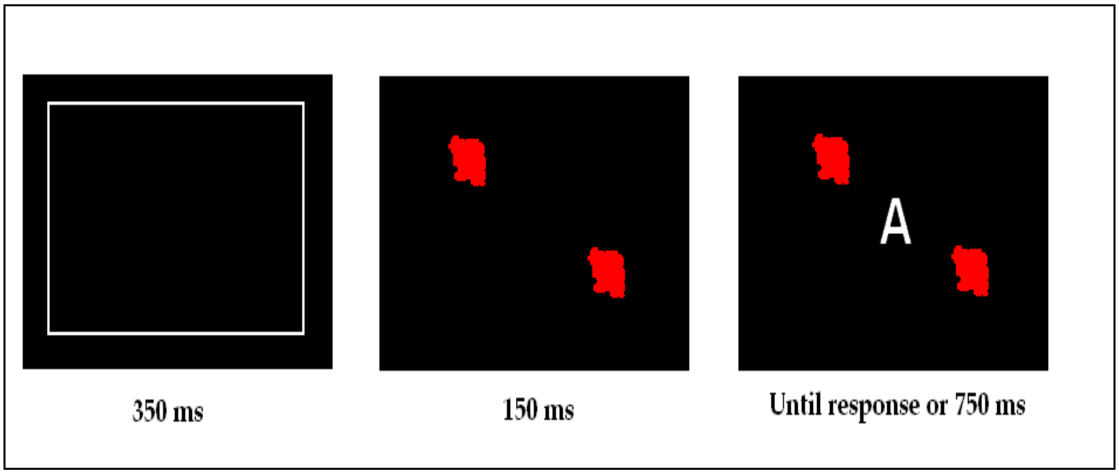

Figure 2. Trial sequence for training blocks.

Test Blocks. As with the training blocks, each trial began with the presentation of a warning stimulus which alerted the participant that the trial was starting. The warning stimulus was presented for 350 ms, immediately following the offset of the warning stimulus, the flankers onset for 150 ms prior to the onset of the target. Both the target and flankers then remained on the screen until the participant responded or the 750 ms deadline had passed (for "no-go" trials), at which time feedback was given only on error trials for 1 second. The next stimulus display appeared about 1.5 seconds after the offset of the feedback for the previous trial. Note that the figures illustrating the training and testing blocks are identical as they have been depicted here. The following example is simply meant to illustrate one example of a trial in a test block, in this case, a trial that presents single color flankers. However, trials which include a set of shape flankers are also presented (but not depicted here). An example of the events of a trial in the testing blocks is below.

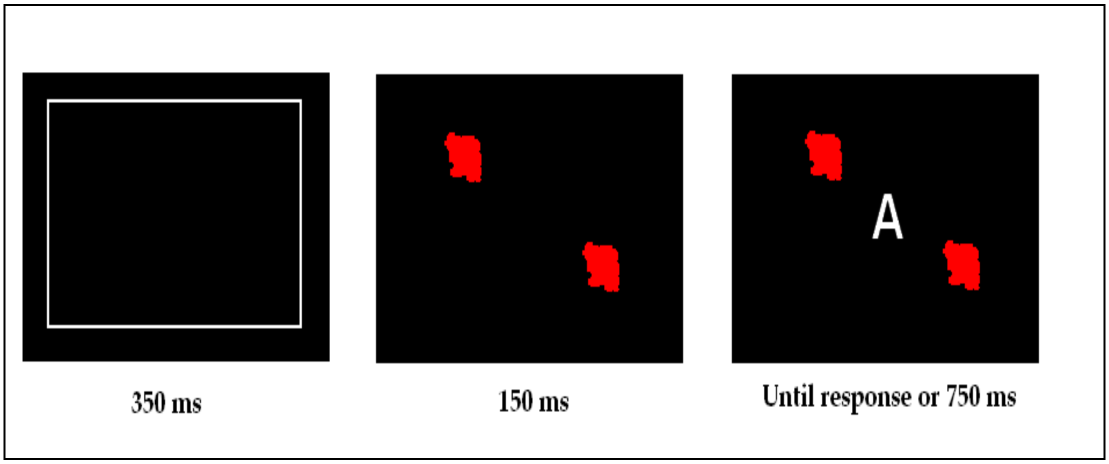

Figure 3. Trial sequence for testing blocks.

Stimuli and Apparatus

Stimuli were presented and responses and response times were recorded by a Dell Optiplex 745/755 series computer using E prime software for stimulus display and data recording, and the E prime S-R box to register responses. Displays were viewed from a distance of about 74 cm. The warning stimulus at the beginning of the trial was a large, white rectangle subtending 19.56° square. The experimental stimuli consisted of three sets of stimuli: targets, color flankers, and shape flankers all presented on a black background. The target stimuli were the capital letters: **A**, **U**, **X**, **M**, **O** and **T** in a sans serif font presented in white. All letters subtended visual angles of approximately 0.98° wide and 1.48° high. The colored flankers consisted of red, blue, green and yellow (MS paint, bitmap color palate) "blobs". These irregular, non-descript shapes were approximately 1.48° for both height and width. Finally, the shape flankers consisted of white outlines (3 point thickness per line) of a pound sign, a titled pound sign (45° rotation of the pound sign), a square and a diamond (45° rotation of the square). The shape flankers were approximately 1.72° for both height and width.

The target was always presented in the center of the screen. All flankers appeared on an imaginary grid with a center-to-center distance of 3.95°. Each flanker could appear at one corner of the imaginary grid with opposite pairs in use at any time (e.g., identical flankers appearing in the upper left and lower right hand corners or the upper right and lower left hand corners). Participants made their responses by either pressing the "white" key (this key was identified as the only key on the box with white paper under it to give the appearance of a white surface) on the S-R box with the index finger of their favored hand or withholding a response.

Issues Pertaining to RT Tasks

There were several changes from the original version of the correlated flankers task as proposed by Miller (1987) to the current experiment. Miller's

(1987) correlated flankers task was a forced-choice RT task. A forced-choice task involves producing a response on every trial, generally a left or right button press. This task is not completely analogous to the type of responding required of animals in conditioning paradigms. It would be the most convincing to investigate the contingency mechanism in humans, while demonstrating that these same mechanisms are operating in animals, by successfully importing theories and knowledge gained from animal learning into this information processing task. Therefore, using a task that more closely resembles that which is used in the animal learning paradigms would be helpful in demonstrating that these same mechanisms are occurring in humans and animals alike. A task that more closely resembles the tasks most often performed by animals in a conditioning paradigm is the go/no-go task in which a response is either made or withheld on every trial. This is more similar to the type of responding which occurs in traditional conditioning experiments –either the animal exhibits the conditioned response or it does not.

A second change from Miller's (1987) paradigm was to present the flanker in advance of the target. This was done to increase the efficacy of the correlated flankers. It has been shown that providing irrelevant information slightly prior to the relevant information increases the effect of the irrelevant information (e.g., Eriksen & Shultz, 1979; Halterman, 2006). This has been shown both for irrelevant information that has an instruction-based relationship with the response (i.e., standard flankers; Eriksen & Shultz, 1979) and for irrelevant information that has a contingent relationship with the response (Halterman, 2006). One plausible explanation of the increased efficacy of the irrelevant information when presented slightly prior to the critical information is that increases in the viewing time of the irrelevant stimulus allow it to overcome deficits in visual acuity. These processing deficits are a product of being presented to either side of the fovea. Response time increases as the distance of a stimulus from foveal fixation increases (Eriksen & Schultz, 1977). In the case of the flankers that have a contingent relationship with the response, presenting the flankers ahead of the target is more akin to typical conditioning procedures in which the presentation of the CS prior to the US is more likely to produce conditioning than if they are presented simultaneously (Schwartz et al., 2002). In the original correlated flankers task, the target and flankers onset simultaneously, but in the current experiment the flankers preceded the target by 150 ms.

Another issue that had to be addressed in the current research was the possibility of confounds. For example, it has been shown in previous research that increasing the probability of an expected display speeds response time to the display (i.e., display frequency bias; Miller & Pachella, 1973). Because of the display frequency bias, Miller's (1987) design used two types of targets. One set of targets were "inducing" targets in which the specific targets appeared with specific flankers 30 out of 32 times of the time. For example, Flanker 1 might appear with Targets 1 and 2 on 30 out of 32 trials in a block whereas Flanker 2 would appear with Targets 4 and 5 on 30 out of 32 trials. These targets are inducing targets because they were not used for data analysis, but were only used to induce the contingency. The other two targets appeared equally often with both Flankers (16 trials with each flanker). These targets were the "test" targets because they did not have any display frequency bias. A difference in RT between trials on which the flanker that was correlated with a particular response was present in the display and the RT from trials on which the flanker correlated with the opposite response was present in the display yielded the correlated flanker effect (CFE; Miller, 1987).

Pilot testing of the conditions for the current series of experiments revealed that there was a great deal of variability in the RT data. In an attempt to decrease the variability, it was necessary to optimize the conditions necessary for finding a difference (CFE). One could either run a large number of people in a given experiment or collect more observations from each participant. In the interest of efficiency more observations were collected from each participant. Thus the design was changed to a design in which all targets were essentially "test" targets in that each had a biased display frequency- a particular flanker was paired with a particular target set. However, in the critical testing phase- each display (combination of flankers and targets) occurred equally often. In this way any problems with display frequency effects were avoided (see Miller & Pachella, 1973 for discussion) in the data being analyzed. This means that any effect observed during the testing blocks was due to the conditioning that was occurring, and not due to a bias to respond more quickly to a display that was presented more often.

A final issue which had to be addressed before using this RT task was that of "lag-one repetition" effects. It has long been observed that a repetition of a stimulus from a single trial to the one immediately following it will produce a faster RT for the repetition trial than for a nonrepetition trial in which the stimulus is different from trial N-1 to the current trial, N (see Pashler & Baylis, 1991 for a review). These repetition effects are hypothesized to reflect a "shortcut" from the stimulus to the response, skipping the response selection phase (Pashler & Baylis, 1991). In order to avoid the influence of lag-one repetition effects, that is, the effect of performance on the N-1 trial on the current trial's performance, the flankers never repeated from trial N-1 to trial N. This allows the flanker effects due to lag-one repetition to be avoided.

This was accomplished by forcing the flanker on the current trial to be different from the trial immediately preceding it (N-1). There were multiple values for a given flanker dimension (e.g., different hues for the color flankers and different shapes for the shape flankers) which alternated from trial-to-trial. This meant that on every trial, the flanker was never identical to the one preceding it, although the relevant dimension of target identity was random. In order to not repeat dimensions between trials, four levels of each dimension (e.g., color or shape) were chosen. Two of the levels were assigned to serve essentially as a single flanker in that two of the flankers had a contingent relationship with the "go" response and two of the flankers had a contingent relationship with the "no-go" response. The flankers were arranged in a way that made one flanker from each pair of flankers assigned to one of the two contingency values (i.e., positive or negative flankers). While it is true that the identity of the flanker on the current trial reduced uncertainty about which flanker would appear (limited the possibilities for the flankers on the following trial), it did not alter the uncertainty about whether a positive or negative flanker would appear on the current trial.

Design

As the major thrust of this work is to understand the effects of compound stimuli – multiple stimulus dimensions must be used in order to determine if the assumption of SWS in humans is valid. Therefore, participants were randomly assigned to be either "color" or "shape" participants (exceptions are noted in individual experiments). For each participant, two of the color flankers and two of the shape flankers were randomly assigned to have a contingent relationship with the "go" response or a contingent relationship with the "no-go" response depending on whether the participant was in the color or shape group. For example, a color participant could be assigned to have “go” responses be positively contingent on the appearance of red or green in the display, and “no-go” responses contingent on either blue or yellow appearing in the display. A shape participant might have “go” responses have a contingent relationship with square and pound sign, and “no-go” responses have a contingent relationship with diamond and titled pound sign. For the participants in the color group, the colors were always presented in a contingent relationship with the responses in the training phases, while the shapes only appeared in the test phase and appeared equally often with both responses therefore were “untrained” in that they did not have a contingent relationship with any response. For the shape participants, the shapes were trained in the training phases, and colors appeared equally often in the test phases (and were untrained). This allowed for testing of each flanker to determine if any produced an innate response[2].

For all experiments, there were always two flankers that had a contingent relationship with the “go” response and two flankers that had a contingent relationship with the “no-go” response. These flankers had the same contingencies and were identical in the design except for the fact that they varied on some level of the dimension (i.e., hue – red, blue, green or yellow; or shape – pound sign, tilted pound sign, square or diamond). In the design tables presented, these two flankers will be shown as “odd or even” flankers. This odd or even designation simply means that there were different flankers presented every other trial to ensure a flanker repetition was not possible. The Flanker 1

---

[2] Each stimulus dimension (color: red, green, blue, yellow, and shape: square, diamond, pound sign, tilted pound sign) was presented equally often with each of the targets in what is being called the “untrained” condition. The means for each of these dimensions were not different from each other, F(1,7) = 0.15, p =.994 and the mean RT to each was 380 ms. A table of the means and proportion of errors for each is included in Appendix A.

(from the odd trials) and the Flanker 3 (from the even trials) have the same relationship to the response. Flanker 2 (from the odd trials) and Flanker 4 (from the even trials) share the same contingency relationship to the response.

In order to label these flankers for later discussion and to be consistent in how they are categorized, the flankers will be discussed as they relate to the primary ("go") response for which a measurement of RT can be obtained. "Positive" flankers are designated as such because they have a positive contingent relationship with making the "go" response by virtue of the fact that more often when a "positive" flanker is in the display the "go" response is being made rather than the "no-go" response. Conversely, "negative" flankers have a negatively valenced contingent relationship with making the "go" response. That is, more often when a "negative" flanker in the display, the correct action is to *not* press the response button. Flankers that were untrained (i.e., occurred equally often with all targets in the training phases) were arbitrarily assigned the role of either "positive or negative" flanker in order to determine if any effect of untrained stimuli was observed.

The following tables represent the basic design of the experiment. Each of the tables represents half the trials (48 trials) of the block because two of the flankers were presented on odd trials (Flankers 1 and 2) and two of the flankers were presented on even trials (Flankers 3 and 4). Illustrated in the tables are the flankers without respect dimension (color or shape) because for a given participant this would be assigned at the start of the experiment. In essence, one set of flankers (Flankers 1 and 3) will be correlated with responding (i.e., the "go" response; pressing the response button) while the other set of flankers (Flankers 2 and 4) are correlated with not responding (i.e., the "no-go" response; not making any overt button press response). In this way, Flankers 1 and 3 are the positive flankers and Flankers 2 and 4 are the negative flankers.

Table 1. Odd trials in a training block.

		Number of Trials Per Block	
Target	**Correct Response**	**Flanker 1** ***Positive***	**Flanker 2** ***Negative***
Go 1	*Button Press*	7	1
Go 2	*Button Press*	7	1
Go 3	*Button Press*	7	1
No-go 1	*No Response*	1	7
No-go 2	*No Response*	1	7
No-go 3	*No Response*	1	7

Table 2. Even trials in a training block.

		Number of Trials Per Block	
Target	**Correct Response**	**Flanker 3** ***Positive***	**Flanker 4** ***Negative***
Go 1	*Button Press*	7	1
Go 2	*Button Press*	7	1
Go 3	*Button Press*	7	1
No-go 1	*No Response*	1	7
No-go 2	*No Response*	1	7
No-go 3	*No Response*	1	7

As is demonstrated in the tables, the number of times a given flanker appeared in the display when a particular response was made establishes the contingency between the flanker and the response. To make this more concrete, here is an example of a particular mapping that a participant might receive. Let's assume that a participant is in the Color Group and receives the following mapping: **A**, **M** and **O** are assigned to the "go" response (i.e., if an **A**, **M**, or **O** is present in the display, press the response button). **T**, **U**, and **X** are assigned to the "no-go" response when one of these letters is present in the display the correct response is to not respond (make no button press).

An analysis of the contingency relationship reveals that when a "positive" flanker is present there is a greater than baseline probability that the correct response is a button press, but when a "negative" flanker is present there is a weaker than baseline probability that the correct response is a button press. This relationship is captured by the following equations. A combination of the probabilities of the two types of flankers in order to determine the contingency present for each yields: $\Delta P = P(R|F+) - P(R)$ for the positive flankers and $\Delta P = P(R|F-) - P(R)$ for the negative flankers. These equations state that the change in probability (P) for a response (R) is captured by the probability of the response when the positive flanker is present (F+) minus the base probability of the response. The negative flankers follow a similar equation. Because the baseline probabilities (P(R)) cancel, the combined equation yields $P(R|F+) - P(R|F-)$. Given that half of the trials will be "go" trials and half will be "no-go" trials, there is a 0.5 base probability that a button press is required and a 0.5 base probability that no button press is required. When the flankers are added to the displays, they change this probability.

Following the example mapping started above, the colors green and blue will be assigned to be "positive flankers". That is, when the flankers are green or blue amorphous shapes are in the display then the probability that the response button is to be pressed is 0.875 (ΔP = 0.375). The colors red and yellow will be assigned to be "negative flankers" which means that the probability of a button press when a red or yellow amorphous shape is in the display is only 0.125 (ΔP = - 0.375).

On a given trial for the participant above, there will be a target diagonally flanked by an identical pair of one of the sets of flankers. For example, on trial 1, a participant would receive a display consisting of an **A** at the center location, diagonally flanked by identical red amorphous shapes. The participant's correct response would be to press the response button. This display would be an example of a negative flanker trial because the color red has a negative contingency with responding based on the number of trials on which it will appear when a "go" response is made. On trial 2, the participant receives a display consisting of an **M** diagonally flanked by identical green amorphous shapes. Again, the appropriate response is a button press, but now this is a positive flanker trial because the color green is positively contingent with responding based on the number of trials on which it will appear when a "go" response is made. On trial 3, the participant may receive a display in which a **T** is diagonally flanked by blue amorphous shapes. This would be an example of a "no-go" trial in which the appropriate response is to do nothing, no button press is required. Finally, on trial 4, the participant receives a display in which an **X** is diagonally flanked by yellow amorphous shapes. Of course, all the targets and flankers are randomly assigned and the order of the presentation of the targets and positive vs. negative flanker trials is randomly determined.

In order to determine if learning is taking place, the Correlated Flanker Effect (CFE) will be calculated. The CFE is calculated by subtracting the RT for the positive flanker trials from the negative flanker trials. Only the RT from the "go" trials can be used because on the "no-go" trials, no response time is collected. One piece of data from the "no-go" trials is the false alarm (FA) rate for these trials. That is, the number of times that a participant responded on a trial when no response was to be made.

Another commonality for all experiments will be the testing phase. In the testing phase both dimensions, color and shape, will be tested for all participants regardless of their initial training with one dimension over another (i.e., Color Group or Shape Group). As is shown in the following tables, in each testing phase, both color and shape will be tested on single trials. That is, there will be color flanker trials and shape flanker trials each consisting of half of the trials per block. The testing phase will always be what can be termed an "extinction" testing phase because all the targets and flankers will appear equally often with each other in every possible combination of targets and flankers such that no contingency will exist in the testing phase. Looking at the tables for the testing phase reveals that the base probability for responding to a "go" target is 0.5 and the base probability of responding to a "no-go" target is 0.5. Notice, that unlike in the training phase, the addition of the flankers to the display does not change this probability because each flanker appears equally often with each target. The flankers and responses do not have a contingent relationship in the testing phase. Additionally, because each target and flanker combination occurs equally often, there is no effect of display frequency. This means that any bias to respond faster to the displays based on the sheer number of times the display has occurred is removed. This allows the data collected in the testing phase to be a bias free estimate of the magnitude of the correlated flanker effect which comes from any learning of the association between the flankers and response in the training phases of the experiment.

Table 3. Odd trials in a testing block.

		Number of Trials Per Block			
Target	**Correct Response**	**Flanker 1 Color**	**Flanker 2 Color**	**Flanker 1 Shape**	**Flanker 2 Shape**
Go 1	*Button Press*	2	2	2	2
Go 2	*Button Press*	2	2	2	2
Go 3	*Button Press*	2	2	2	2
No-go 1	*No Response*	2	2	2	2
No-go 2	*No Response*	2	2	2	2
No-go 3	*No Response*	2	2	2	2

Table 4. Even trials in a testing block.

		Number of Trials Per Block			
Target	**Correct Response**	**Flanker 3 Color**	**Flanker 4 Color**	**Flanker 3 Shape**	**Flanker 4 Shape**
Go 1	*Button Press*	2	2	2	2
Go 2	*Button Press*	2	2	2	2
Go 3	*Button Press*	2	2	2	2
No-go 1	*No Response*	2	2	2	2
No-go 2	*No Response*	2	2	2	2
No-go 3	*No Response*	2	2	2	2

A pilot experiment to test the design described above and to determine if the changes to the original correlated flankers task would still produce the predicted findings of the correlated flanker effect (CFE) was conducted. Mean RTs for the Color Group on the color dimension were 386 ms[3] for the negative flanker trials and 372 ms for the positive flanker trials. While the mean RTs for the Color Group on the shape dimension were 383 ms for the negative flanker trials and 377 ms for the positive flanker trials. Mean RTs for the Shape Group on the color dimension were 380 ms for the negative flanker trials and 380 ms for the positive flanker trials. While the mean RTs for the Shape Group on the shape dimension were 387 ms for the negative flanker trials and 369 ms for the positive flanker trials. An ANOVA between the groups revealed that there was no significant difference for dimension (F(1, 95) = 1.19, p = .277), but there was a significant difference for training (F(1, 95) = 10.29, p = .002) (i.e., the trained dimension showed a CFE, while the untrained dimension showed no difference between the negatively contingent flanker trials and the positively contingent flanker trials). There was not an interaction between dimension and training (F(1, 95) = .057, p = .811). In the color group, the trained dimension (color) showed a 13.81 ms, t(23) = 4.41, p < .001 correlated flanker effect (CFE; RT from negative flanker trials minus RT from positive flanker trials), while the untrained dimension (shape) did not show a significant effect mean CFE = 5.90 ms, t(23) = 1.43, p = .165. The overall error rate was less than 1% which made the false alarm rate too infrequent to be analyzed as those accounted for only a portion of the errors. For the shape group, the trained dimension (shape) showed a 17.97 ms, t(23) = 3.89, p = .003 CFE, while the untrained dimension (color) did not show a

---

[3] All mean RTs are reported as rounded to the nearest whole number values. A table of all the mean RTs for every experiment is reported in Appendix C.

significant effect, mean CFE = - 0.35 ms, t(23) = 0.210, p =.836. Again, the overall error rate was less than 1%.

All of the results will be reported in this manner and will be represented in a figure as well. The following figure illustrates the results previously described.

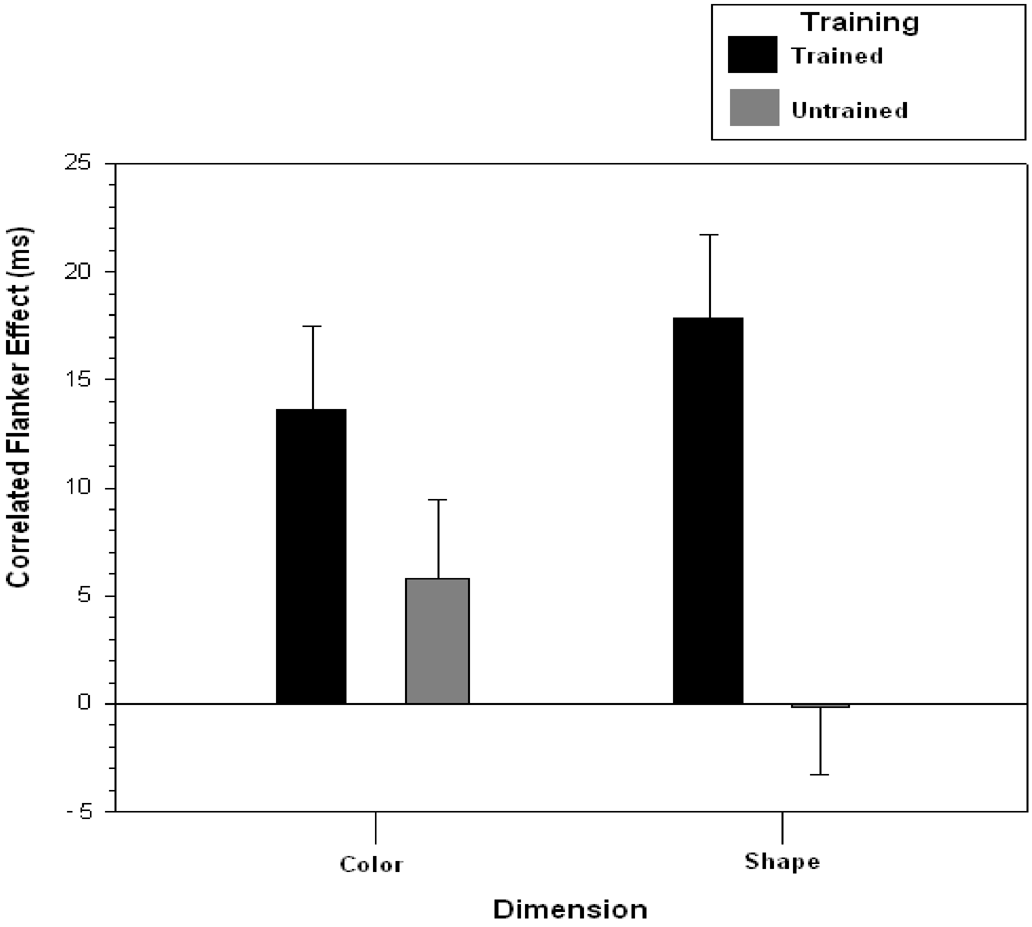

Figure 4. Pilot data. Data come from two groups of participants (color or shape). The data for these groups is split across the plot by the dimension (color or shape) and the training type (trained or untrained). The key pieces of data to focus on are the trained color and shape dimensions (pictured as black bars).

In most of the graphs of the data, flanker dimension (color or shape) will be partitioned into subpanels while the type of training (trained or untrained) will be presented in subplots (exceptions will be noted). The magnitude of the CFE (taken as evidence of associative strength) is plotted on the y-axis. Essentially, this figure shows that there was associative strength for the trained dimensions for both Color and Shape Group participants, but not for the untrained dimensions.

Finally, the following notations will be used to capture the relationships between the flankers and the responses in each experiment. This information will also be provided in Appendix B for reference. These relationships are illustrated in the various design tables for each of the experiments. **C** stands for color, **S** stands for shape and **R** stands for response. **C -> R** will be used to indicate when the color flankers have a contingent relationship with the response. **S -> R** indicates the shape flankers have a contingent relationship with the response. Finally, a **C** or **S** not followed by an **R** indicates that there is not a contingent relationship between the flankers and the responses. This will occur primarily in the testing phase though can occur in the training phase if there is not a relationship between the flankers and the responses. Examples of the notations for the pilot experiment follow.

Table 5. Experimental design for the color group in the pilot experiment.

Phase 1	**Phase 2**	**Test Phase**
C -> R	C -> R	C
		S

Table 6. Experimental design for the shape group in the pilot experiment.

Phase 1	Phase 2	Test Phase
S -> R	S -> R	C
		S

Recall that only the color dimension was trained for the color group, as represented by the **C -> R** designation, but that both color and shape were tested in the final phase. Only shape was trained for the shape group, as represented by the **S -> R** designation, and again, both color and shape were tested in the final phase. These notations will be used for all of the following experiments. There are only two other conditions that may occur. **C,S -> R** will be used to indicate a compound trial in which both shape and color flankers are presented simultaneously and have a contingent relationship with the response. An example of the trial events for a compound trial appears below.

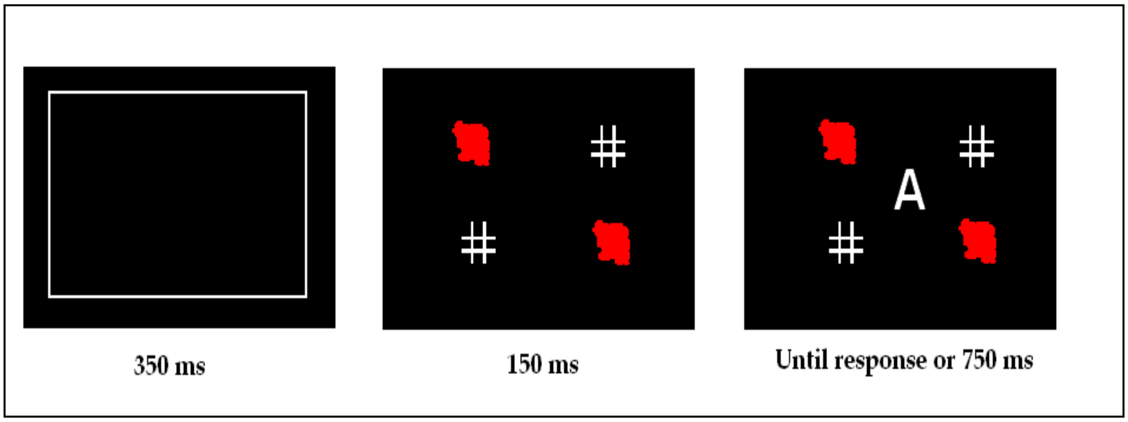

Figure 5. Trial events for compound trials (training or testing blocks).

C->R, S->R is used to indicate that there are single trials which have color flankers and single trials which have shape flankers. Both flankers have contingent relationships with the response though each is presented on its own.

The following chapters will detail the experiments conducted in order to determine if the conditions for learning that have been shown to work for animals will be able to be demonstrated with the correlated flankers task. Each chapter will address some of these basic phenomena and how these experiments were conducted. Finally, the paper will wrap up with the conclusions that can be drawn for the work thus far and future directions for research.

CHAPTER 3

OVERSHADOWING AND BLOCKING

While previous research (e.g., Miller, 1987) has demonstrated acquisition in the correlated flankers task, more complex conditioning phenomena have never been attempted using the correlated flankers task. Yet, compound conditioning shaped many contemporary learning theories. Indeed, demonstrating that humans process compound stimuli in Shared Weight Space (SWS), as has been demonstrated in animals, will strengthen the assumption that the mechanisms underlying animal learning underlie the learning of irrelevant information in humans as well. Two phenomena which involve the presentation of multiple CSs are overshadowing and blocking. The following series of experiments sought to demonstrate overshadowing and blocking in the correlated flankers task.

Overshadowing

Some of the most widely researched phenomena in both animal conditioning and previous conditioning experiments with human participants involve the presentation of multiple CSs in a compound. Theories that assume shared weight space (SWS) predict that when two conditioned stimuli (CSs) are presented simultaneously the available associative strength, that is, the maximum associative strength supported by the unconditioned stimulus (US), will be split between the CSs. The allotment of strength to each of these stimuli will depend on the salience of the stimuli. It has been shown experimentally, and is predicted by theories that assume SWS, that the more salient of the stimuli will obtain more associative strength than the less salient of the stimuli. In general, theories that appeal to SWS for their explanatory power predict that the more salient stimulus will gain more associative strength because stimulus salience is a multiplicative value in many of those equations and those theories assume a finite, and common, associative weight for all CSs presented. This prediction has been observed experimentally and labeled, “overshadowing”.

Overshadowing is a widely discussed phenomenon that refers to the finding that often if two stimuli are presented as a compound with a US then one of these CSs might fail to produce the CR even though the presentation of the compound does produce a CR (Schwartz et al., 2002). For example, a light (CS_A) and a tone (CS_B) are presented together (such that CS_{AB}) with food (US) until salivation occurs (CR). Later when testing occurs and the presentation of just the light (CS_A) or just the tone (CS_B) occurs the less salient of the two stimuli will not produce salivation (CR). The evidence for this reduction or elimination of the CR to the less salient CS is obtained by comparison with a control group. The experimental group is compared to a control group(s) in which the elements of the compound are paired individually with the US, rather than receiving both in compound, and in this case no evidence of a reduction of associative strength for either CS should be exhibited for the control group.

Taking a popular example of a theory that assumes SWS, Rescorla-Wagner theory, the prediction that the more salient stimulus will gain more associative strength comes from the salience variable. According to Rescorla Wagner theory, learning of two CSs (CS_A and CS_B) is captured by the equations: $\Delta V_A = \alpha_A \beta(\lambda - V_{AB})$ and $\Delta V_B = \alpha_B \beta(\lambda - V_{AB})$ where the change in associative strength (ΔV) is the product of the CS salience (α), the learning rate (β) and the maximum available associative strength (λ) minus the current associative strength of all CS on a given trial (V_{AB}). The learning rate is constant for both stimuli as is the maximum associative strength supported by the US (λ). The more salient stimulus gains more associative strength because on each trial, the total associative strength is multiplied by the salience and if this number is greater more associative strength is assigned to the stimulus (see Appendix E for an example of this formula).

Overshadowing can be demonstrated with the correlated flankers task by presenting the two dimensions previously tested in the first experiment (i.e., color and shape) in a compound. The magnitude of the correlated flanker effect from this experiment will be compared to the correlated flanker effect from the pilot experiment in order to determine if any associative strength has been lost for either dimension now that they have appeared in a compound. *A priori* there is no reason to propose that either color or shape will overshadow each other, but this experiment will test for that possibility. One possibility is that color will overshadow shape given that it has been shown through work with event-related potentials (ERPs) that while shape seems to be processed first, color has a stronger effect when it is finally processed (Mordkoff, Miller & Roch, 1996). Shape could also overshadow color because it showed a non-significant trend to being stronger in the pilot experiment and could be more salient given that the dimension of shape is somewhat relevant to the target dimension (i.e., also shape, but nameable letters rather than geometric shapes). If either shape or color overshadows the other dimension this will be evidence for theories that rely on SWS. Alternatively, if both shape and color are of have a similar level of salience then mutual overshadowing, in which both stimuli exhibit a decrease in associative strength, will be observed (see Appendix E for calculations).

Experiment 1

Experiment 1 was designed to look at overshadowing in the correlated flankers task. The magnitude of the CFE in this experiment will be compared back to the pilot experiment as the control group. In the current experiment, the conditions necessary for evoking overshadowing were used. Namely, both flanker dimensions (i.e., color and shape) were presented on every trial. These dimensions were presented in separate objects (i.e., colored flankers and shape flankers) simultaneously on each trial.

Design

The design is illustrated in the following table (for abbreviations refer to General Method, pp. 53-55 or Appendix B, p. 112).

Table 7. Experimental design for Experiment 1.

Phase 1	**Phase 2**	**Test Phase**
C,S-> R	C,S -> R	C
		S

Generally, these training blocks can be thought of as occurring in two phases (training phase 1 and training phase 2), but both the pilot experiment and Experiment 1 have identical training phases 1 and 2. Notice as well that the flankers presented in this case are compounds meaning that both a color flanker and a shape flanker were presented on each trial as shown above. This also necessarily means that there was a perfect contingency between a given color flanker and a given shape flanker (i.e., the pieces of a particular compound always occurred together).

The testing phase remains constant throughout the experiments, always testing both flanker types (color and shape). In the overshadowing experiment this means that both the color and shape flankers were tested as individual trials for each participant. Each of the testing blocks consisted of 96 trials. There were two testing blocks for the experiment. Again, there is no contingency being used in the testing phase, so all targets and flankers occur equally often.

Participants

Twenty-four undergraduates (11 women, 13 men; ranging in age from 18 to 21 years) participated in a single session lasting about 45 minutes. All of the participants provided informed consent but were naïve as to the study's design and purpose. All reported normal or corrected-to-normal, full color vision. The majority of participants were right-handed according to self-report (4 left-handers, 20 right-handers).

Procedure

The procedure for this experiment was similar to that described in the General Method except for the training blocks. Participants in this experiment were not assigned to either a color or shape group as all participants saw both color and shape flankers in each display in the training blocks. The practice and testing Blocks were identical to those described in the General Method.

Training Blocks. As with the practice blocks, each trial began with the presentation of a warning stimulus which alerted the participant that the trial was starting. The warning stimulus was presented for 350 ms. In the training blocks immediately following the offset of the warning stimulus, the flankers onset for 150 ms, and then the target was presented. This final display remained on the screen until the participant responded or the 750 ms deadline had passed (for "no-go" trials), at which time feedback was given only on error trials for 1 second. The next stimulus display appeared about 1.5 seconds after the offset of the feedback for the previous trial. Any error trials were discarded and not rerun in the block. Participants saw color and shape flankers on every trial. The flankers with the same dimension (i.e., color or shape) appeared in random locations on the imaginary grid, but always appeared in opposite corners to each other.

Results and Discussion

A summary of the data is provided in Figure 6. This experiment is compared to the pilot experiment which served as a control. In this experiment, the correlated flanker effect (CFE; negatively contingent flanker trials minus positively contingent flanker trials) for the color flankers was 5.67 ms (ns) and the CFE for the shape flankers was 2.54 ms (ns). Mean RTs for the color dimension were 383 ms for the negative flanker trials and 378 ms for the positive flanker trials. While the mean RTs for the shape dimension were 378 ms for the negative flanker trials and 376 ms for the positive flanker trials. False alarm rates were too infrequent to be analyzed.

The data in this overshadowing experiment were compared to the data from the pilot experiment. The data in the pilot experiment were meant to serve as a control for this experiment. It was found that there was a significant decrease in the magnitude of the CFE for the shape dimension, F(2, 22) = 3.45, p = .003. It was also found that there was a significant decrease in the magnitude of the CFE for the color dimension, F(2, 22) = 2.56, p = .024. This comparison between the CFE for each dimension when it is trained alone (i.e., only shape flankers are present on every trial or only color flankers are present on every trial) and the CFE for each dimension when it is trained as a compound (i.e., both color flanker and shape flankers appeared on every training trial) is used to detect overshadowing. Comparing the two experiments, evidence for mutual overshadowing is obtained.

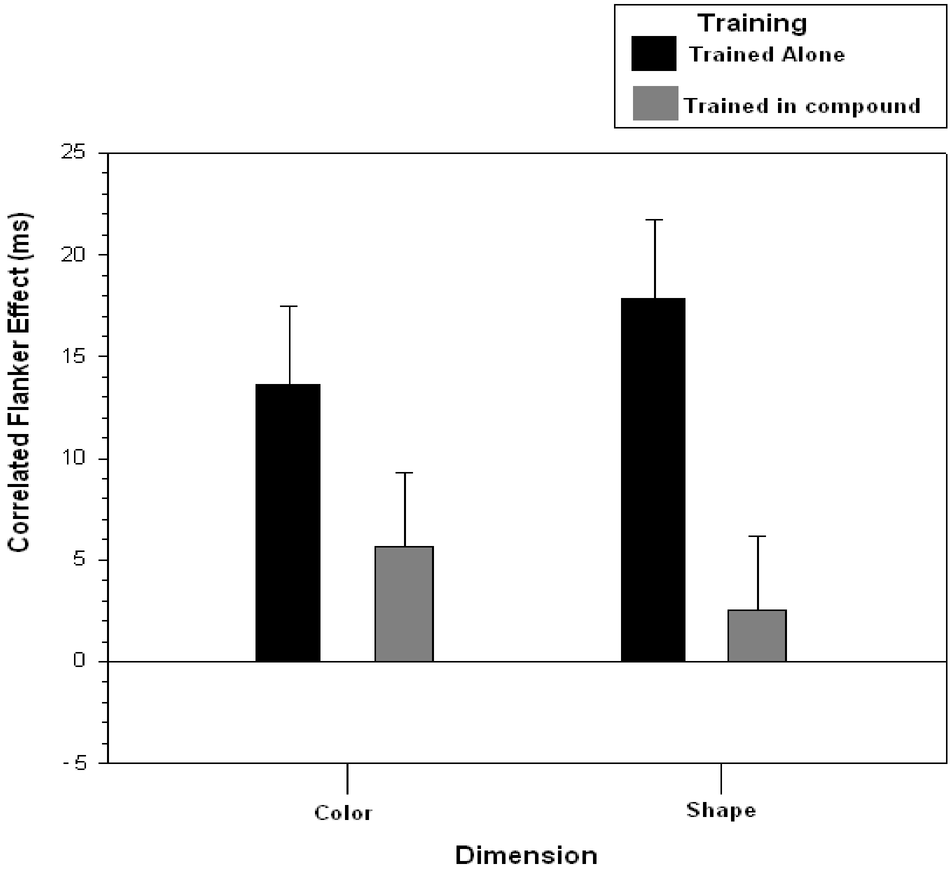

Figure 6. Results for first overshadowing experiment. The data serving as the trained alone data are from the pilot experiment which was designed to serve as a control for Experiment 1 (pictured as black bars). The data from Experiment 1 are plotted as the trained together data meaning that both color and shape flankers were present during all training phases (pictured as gray bars). The key piece of data from this figure is the difference that training alone versus training together makes on the CFE.

The typical overshadowing effect predicted by theories which assume SWS is that the more salient stimulus will demonstrate learning while the less salient dimension will show little-to-no learning. In terms of the current experiment, a demonstration of this typical pattern of overshadowing would be if one of the two dimensions exhibited the same magnitude of CFE as in the control experiment while the other dimension would exhibit a significant decrease in the magnitude of CFE – possibility exhibits no CFE at all. In this experiment however, both dimensions exhibited a significant decrease in the magnitude of the CFE such that neither dimension produced a significant CFE. This is evidence of mutual overshadowing in which both stimuli which were of similar salience (evidence of similar salience comes from the pilot experiment in which there was not a significant difference in the CFEs for the shape and color dimensions).

Although the more typical pattern of overshadowing was not found in the current experiment, a pattern of mutual overshadowing was obtained. The conclusion from this first experiment could be that mutual overshadowing is obtained. However, data from another experiment speak to a different explanation as well which is also consistent with theories that assume SWS. One can assume that in this experiment, participants were treating the displays not as two elements (color *or* shape), but as a compound of color and shape. Given the constraints of the design, the color and shape compounds were always the same (e.g., blue *and* pound sign would always appear together). This means that there was a perfect contingency between the flankers. From an information processing point of view, given that the flankers were colors and shapes which have little chance of interfering with each other, the colors and shapes could have been combined to form a single item (e.g., a blue pound sign). This combination of two features (color and shape) into a single object defined by those two features is an example of a conjunction. Conjunction stimuli are defined here as single objects that have multiple features that comprise the object’s identity (i.e., there is a color *and* a shape that define the object). There is some evidence that participants do process color and shape information together rather than keeping them separate (for review see Cohen & Shoup, 1997; Mordkoff & Danek, in review). In the training phase, participants viewed displays that consisted of color flankers and shape flankers, presumably these flankers were processed as a single color/shape conjunction. Following the previous example, while blue pound signs have a positively contingent relationship with pressing the response button, neither blue nor pound sign has a relationship with a response because neither was ever trained. In the testing phases, when single flanker trials of blue and pound sign were presented no evidence of learning was observed because these stimuli were novel. Because the displays were processed as conjunctions, it is as though the two stimuli (e.g., blue amorphous shape and white pound sign) did not occur prior to the testing phase in which there is no contingent relationship between the flankers and responses as all combinations occur equally often. This assumption does not invalidate SWS theories, it simply raises an issue regarding how stimuli are processed.

An additional experiment was designed to discourage participants from processing the color and shape information as a conjunction, but instead to keep the color and shape information as separate components. This experiment will test the idea that participants are treating the color and shape flankers as a single unit (e.g., blue pound sign rather than blue *and* pound sign). If the flankers are treated as separate entities within a compound then a theory that proposes SWS would predict that the typical pattern of one stimulus overshadowing the other would occur if there was a difference in salience for the two dimensions.

Experiment 1A: Overshadowing/Blocking Control

The following experiment serves two purposes. It was originally designed to be the control group for comparison with the experimental blocking condition (Experiment 2). However, given the interesting results obtained in this experiment as compared to the original overshadowing experiment it is also an experiment that allows for the color and shape flankers to be processed separately. The data from this experiment will also be used as the control for the blocking experiment in due course.

Design

The design is illustrated in the following tables (for abbreviations refer to General Method, pp. 53-55 or Appendix B, p. 112).

Table 8. Experimental design for Experiment 1A.

Phase 1	**Phase 2**	**Test Phase**
N-> R	C,S -> R	C
		S

Contrary to the experiments presented previously, the current experiment did have two different training phases. In training phase 1, "neutral" (**N**) stimuli were used. The stimuli are designated as "neutral" because they are stimuli that will only be used in the training phase and will never be tested again. Additionally, the neutral stimuli have an equivalent relationship with both the color and shape dimensions as the neutral flankers are deliberately composed of both color and shape. This was done to ensure that the neutral stimuli were not biased toward one dimension over the other.

Participants

Thirty-six undergraduates (18 women, 18 men; ranging in age from 18 to 24 years) participated in a single session lasting about 45 minutes. All of the participants provided informed consent but were naïve as to the study's design and purpose. All reported normal or corrected-to-normal, full color vision. The majority of participants were right-handed according to self-report (5 left-handers, 31 right-handers).

Procedure

The procedure for this experiment was the same as that described in the General Method except for the training blocks. For the first three training blocks, participants were presented with a single flanker dimension (e.g., a "neutral" stimulus). For the next three training blocks, participants saw both the usual color and shape flankers in the same display (just as in the previous Overshadowing).

Stimuli and Apparatus

The typical stimuli and apparatus are the same as those described in the General Method. The neutral stimuli were a combination of colors (peach, turquoise, pale purple and light green) in non-descript shapes which subtended visual angles of approximately 1.48° for both height and width. These "shapes" consisted of veridical representations of the "shape" and 45° rotations of those same "shapes". Each neutral stimulus was a combination of all colors, but generally had a "main color" – a color that covered more than 50% of the entire pixel area.

Results and Discussion

The results for this experiment (as compared with the data from the pilot experiment) are presented in Figure 7. Mean RTs for the color dimension were 381 ms for the negative flanker trials and 368 ms for the positive flanker trials. While the mean RTs for the shape dimension were 379 ms for the negative flanker trials and 375 ms for the positive flanker trials. A 2x2 ANOVA revealed an interaction between dimension type (color or shape) and training (trained separately or trained together as a compound), F(1, 22) = 2.36, p = .002. Subsequent t-tests revealed a main effect of dimension. The CFE for the color dimension was 12.58 ms, t(23) = 4.48, p <.001, while the CFE for the shape dimension was 3.87 ms, t(23) = 0.984, p = .335. Thus revealing that in this experiment, the color dimension was able to overshadow the shape dimension. False alarms were too infrequent to be analyzed.

The results from this experiment show the typical overshadowing pattern as would be expected by a theory that assumes SWS if the two stimuli had different salience values. Evidence from this experiment indicates that the color dimension was more salient than the shape dimension and so retained more of its associative strength, while the shape dimension did not exhibit significant learning as evidenced by the magnitude of the CFE for the shape dimension. There is evidence from other experiments that suggests that color and shape are dimensions with different saliency values. This experiment was designed in order to serve as the control for a different experiment, but it also is relevant to the question of how the two stimuli are processed. One explanation for the results in Experiment 1 is that both the color and shape features are conjoined to form a single object which then acquires associative strength. In the testing phase, this single object is never presented again as each trial consists of either a shape flanker or a color flanker.

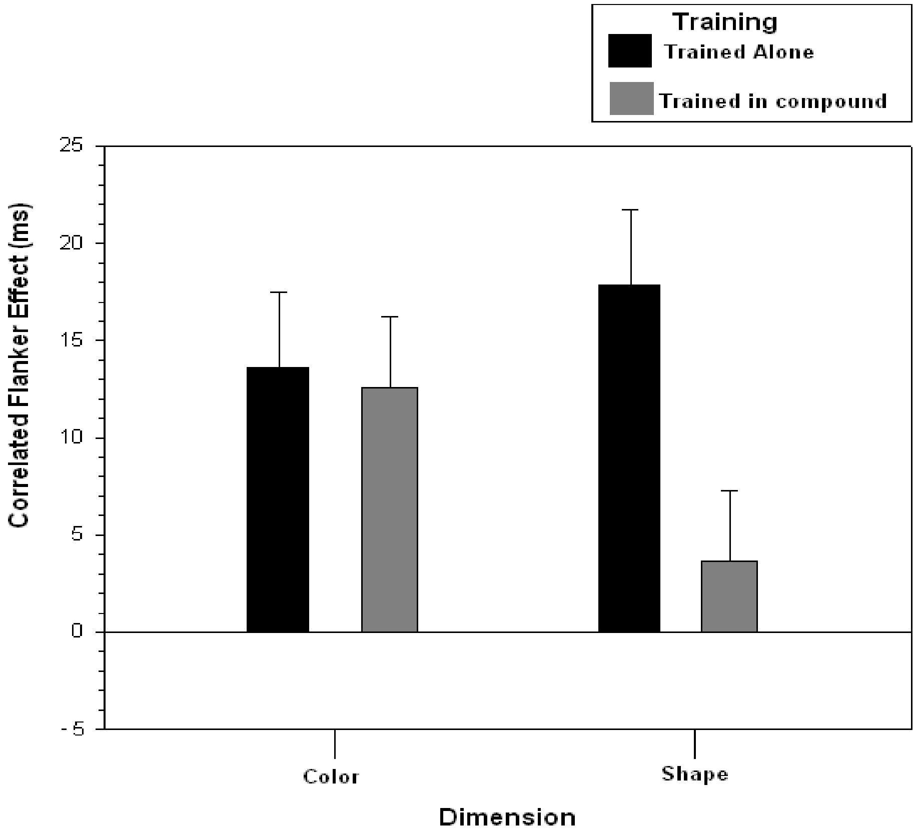

Figure 7. Results for second overshadowing experiment. The data plotted here come from the pilot experiment (trained alone condition; black bars) and from Experiment 1A (trained together as a compound; gray bars). The key piece of data to focus on is the difference between the trained alone conditions versus the trained together as a compound conditions for each dimension.

Why did this change occur? One answer is that the participants in this experiment were processing the displays in an elemental fashion. The only change in this experiment is that there was a training phase (consisting of three blocks) in which a "single stimulus" was presented on every trial (the neutral stimulus described in the method). This neutral stimulus seems to have changed the way in which participants were processing the stimuli from conjunctions of two features to two separate features. One explanation for how this happened is that participants received many exposures to a single set of flankers appearing diagonally from each other on every trial. This presentation of a single set of flankers (presented randomly on either opposing corners of an imaginary square) created for the participant, the notion that a “single set” of flankers occurred diagonally from each other. When presented with the compound trials in which now a set of shape flankers and a set of color flankers appear on every trial, the participants did not combine the two flankers, but kept them separate dimensions, and the stronger of the two dimensions was able to overshadow.

The finding that color was able to overshadow shape is consistent with previous research. In a previous experiment comparing the processing of different visual dimensions, color exhibited stronger processing than did shape (Mordkoff et al., 1996). This is one reason why color may have been more salient than shape because the information coming from color is processed more rapidly than the shape information. Another, less mechanistic, reason why color might be more salient than shape is through experience in the world. From a Gibsonian point of view, color is a more stable cue to an outcome in the environment than shape. For example, a stop light has the same shape for each of the different lights, it is the color that indicates whether to stop, slow, or keep going. It is possible than that through experience, color has become a more reliable predictor of events in the environment and draws its salience from this fact.

Experiment 1B: Overshadowing Control

The preceding experiment was used as a demonstration of overshadowing. However, a fair comparison group for the preceding experiment is one in which there is only one phase of training for the separate dimensions. The pilot experiment which was designed to be the control experiment for the first overshadowing experiment had two phases of training for the separate dimensions which makes it an unfair comparison for Experiment 1A in which participants only had one phase of training with the overshadowing stimuli (i.e., the compound stimuli). The current experiment was designed to be a better comparison for Experiment 1A because it only had one phase of training of the separate dimensions.

Design

The design is illustrated in the following tables (for abbreviations refer to General Method, pp. 53-55 or Appendix B, p. 112).

Table 9. Experimental design for the color group in Experiment 1B.

Phase 1	**Phase 2**	**Test Phase**
N-> R	C -> R	C
		S

Table 10. Experimental design for the shape group in Experiment 1B.

Phase 1	**Phase 2**	**Test Phase**
N-> R	S -> R	C
		S

Participants

Seventy- two undergraduates participated in a single session lasting about 45 minutes. All of the participants provided informed consent but were naïve as to the study's design and purpose. All reported normal or corrected-to-normal, full color vision. Random assignment was used to assign 36 participants to the Color Control Group (21 women, 15 men; ranging in age from 18 to 25 years) and to assign the 36 participants to the Shape Control Group (17 women, 19 men; ranging in age from 18 to 21 years). The majority of participants in both groups were right-handed according to self-report (2 left-handers, 34 right-handers for the color group and 4 left-handers, 32 right-handers for the shape group).

Procedure

The procedure for this experiment was similar to that described in the General Method except for the training blocks. For the first three training blocks, participants were presented with a single flanker dimension (e.g., a "neutral" stimulus). For the next three training blocks, participants saw either the color flankers (Color Group) or the shape flankers (Shape Group). The practice and testing blocks were identical to those described in the General Method.

Results and Discussion

The results for this experiment (as compared with Experiment 1A) are presented in Figure 8. Mean RTs for the color dimension were 376 ms for the negative flanker trials and 365 ms for the positive flanker trials. While the mean RTs for the Color Group on the shape dimension were 374 ms for the negative flanker trials and 361 ms for the positive flanker trials. A 2x2 ANOVA revealed an interaction between dimension type (color or shape) and training (trained separately or trained together), F(3, 104) = 4.30, p = .04. Subsequent t-tests revealed a main effect of dimension. The CFE for the color dimension was 10.48 ms, t(35) = 3.19, p = .003, while the CFE for the shape dimension was 12.75 ms, t(35) = 3.58, p = .001. In this control experiment, color and shape were both able to have an effect. False alarms were too infrequent to be analyzed.

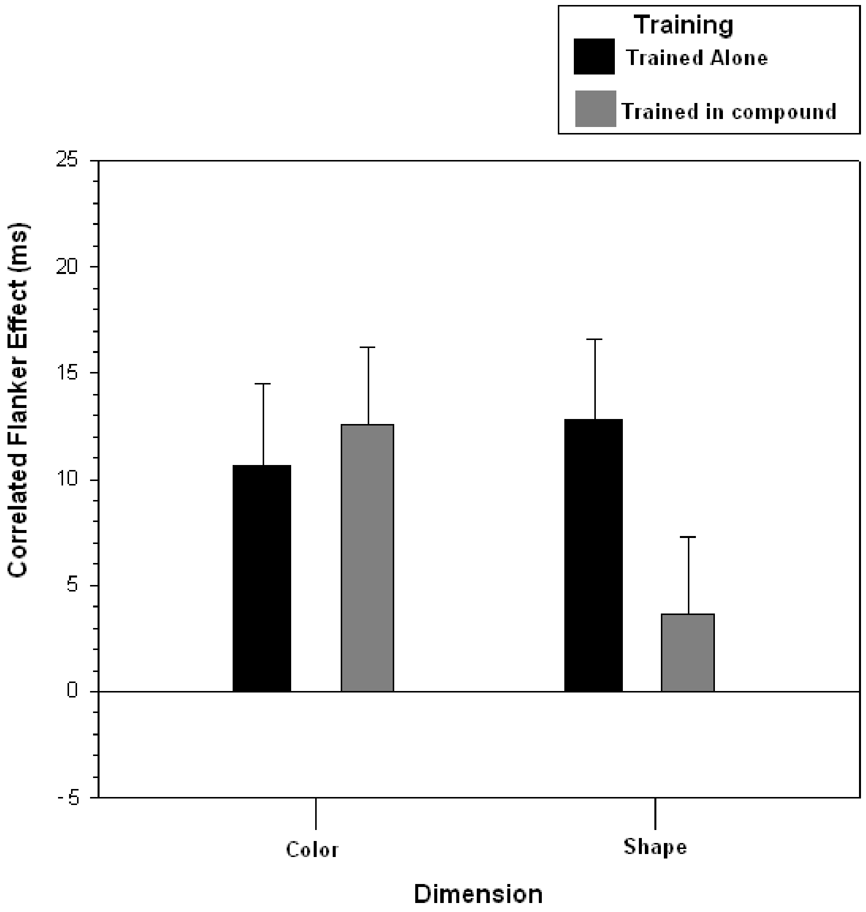

Figure 8. Results for overshadowing experiment with appropriate control. The data plotted here come from this experiment, 1B (trained alone condition; black bars), and from Experiment 1A (trained together as a compound; gray bars). The key piece of data to focus on is the difference between the trained alone conditions versus the trained together as a compound conditions for each dimension.

The design of this experiment was more similar to the design of the second overshadowing experiment and as such was a more appropriate control for Experiment 1A. As was found previously, the comparison between the control group and the experimental overshadowing group still exhibited a pattern of overshadowing. Again, the color dimension for the experimental group was more salient than the shape dimension for the experimental group as evidenced by the reduction of the CFE for the shape dimension in the experimental group (Experiment 1A) compared to the control group (Experiment 1B). This experiment strengthens the conclusion drawn from the series of overshadowing experiments that when participants are biased to process stimuli as single elements rather than conjunctions, they exhibit learning predicted by SWS theories.

Blocking

Another phenomenon looked at in the animal learning literature that involves compound stimuli is that of blocking. First presented by Kamin (1968), blocking refers to the effect of training with a stimulus before presenting the stimulus in a compound with another CS. For example, this procedure involves the presentation of a light (CS_A) and food (US) until salivation (CR) is reliably produced. The next phase involves presenting the light (CS_A) in combination with a tone (now CS_{AB}). This compound is then paired with the food (US). At testing the tone (CS_B) is presented in isolation and the effect is observed. It has been shown that there is little or no conditioning of the tone in this case. The light is said to have "blocked" the tone from an association with the food (Kamin, 1968).

Kamin's (1968) explanation of this finding was that in order for conditioning to occur a stimulus must have a surprisal value. It must be providing new information about the possibility of the US occurrence in order for learning to occur. In the above example the light is completely predictive of the food so the inclusion of the tone in the light-tone compound is providing no additional information about the upcoming food (Kamin, 1968). This evidence for blocking is obtained by comparing the conditioning of CS_B for the experimental group with conditioning of CS_B for a control group. The control group consists of a pre-training phase in which a different stimulus (e.g., CS_Z) is trained before presentation of the compound or in which there is no training of a stimulus in the first phase of training. Then the compound training is conducted and the CR to the target CS (CS_B) in the control experiment should be larger than the CR to the target CS_B for the experimental group.

Theories that assume shared weight space (SWS) predict that experience with a stimulus will produce changes in the associative strength for that stimulus which has been demonstrated experimentally. These theories work by assuming a finite amount of associative strength and the portioning of this associative strength among all stimuli present, predicts that a previously experienced stimulus will gain more associative strength than a second stimulus that is presented in a compound with the previously trained stimulus. If the mechanism underlying the correlated flanker effect is the same as that underlying the learning in general then we would expect a demonstration of blocking with the correlated flankers task given that associative theories which assume SWS and are powerful predictors of learning behavior predict that blocking will occur.

An example of a theory that assumes SWS is Rescorla-Wagner theory. Rescorla-Wagner theory explains blocking in a similar manner to overshadowing (see Appendix E for details of this equation). Whenever a compound stimulus (CS_{AB}) is presented there is competition between the elements of the compound (CS_A and CS_B) for associative strength. If one stimulus receives training prior to the presentation of the compound stimulus, then the associative strength will start to approach the maximum. When a second stimulus is added to make a compound, there is little associative strength leftover before reaching asymptote. Once the maximum level of associative strength has been reached, no more learning takes place. This results in the second stimulus of the compound having a low associative value (as compared to a condition in which it is presented alone), which is the blocking effect (Rescorla & Wagner, 1972).

As color has already been demonstrated to be able to overshadow shape, showing that color can block shape is uninteresting. Shape will be tested for its ability to block using the correlated flankers task. This can be accomplished by presenting single flanker shape trials for the first three blocks of training. In the second phase of training, compound trials of color and shape flankers will be presented. Then the typical testing blocks will test the dimensions of color and shape separately and compare them to Experiment 1A – the control experiment for a blocking condition.

In order to determine if shape is able to block color, the comparison will be between the CFE for the shape dimension when it is presented first (as the blocker) to the CFE for the shape dimension from the control experiment and the CFE of the color dimension (the blockee) for the current experiment to the CFE of the color dimension of the control experiment. If the second dimension suffers the loss of associative strength (as measured by a decrease in the magnitude of the CFE) then evidence for blocking will be obtained. The predicted results follow.

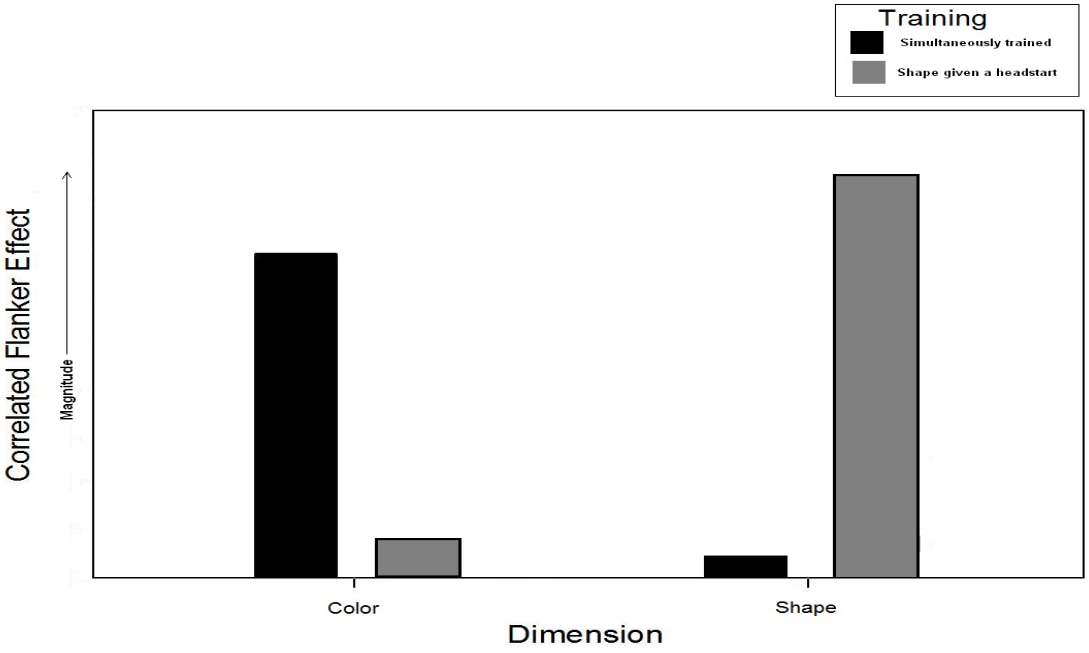

Figure 9. Predicted results for blocking experiment. The presented data are idealized data so there are not specific numbers for the CFE plotted. Instead the focus is on the relative magnitudes of the conditions.

Experiment 2

Experiment 2 was designed to look at the phenomenon of blocking in a correlated flankers task. In the overshadowing experiment (Experiment 1A) it was demonstrated that color flankers will overshadow shape flankers when they are presented simultaneously. It would not be surprising then if color flankers were able to block shape flankers. In this experiment the shape flankers will be presented alone prior to the presentation of the compound in order to determine if shape flankers can block color flankers.

Design

The design of the experiment is illustrated in the following table (for abbreviations refer to General Method, pp. 53-55 or Appendix B, p. 112).

Table 11. Experimental design for Experiment 2.

Phase 1	**Phase 2**	**Test Phase**
S-> R	C,S -> R	C
		S

Participants

Thirty-six undergraduates (19 women, 17 men; ranging in age from 18 to 21 years) participated in a single session lasting about 45 minutes. All of the participants provided informed consent but were naïve as to the study's design and purpose. All reported normal or corrected-to-normal, full color vision. All of the participants were right-handed.

Procedure

The procedure for this experiment was similar to that described in the General Method except for the training blocks. For the first three training blocks, participants were presented with a single flanker dimension (shape). For the next three training blocks, participants saw both the color and shape flankers in the same display (just as in the overshadowing experiment). The practice and testing blocks were identical to those described in the General Method.

Results and Discussion

Mean RTs for the color dimension were 389 ms for the negative flanker trials and 381 ms for the positive flanker trials. While the mean RTs for the shape dimension were 391 ms for the negative flanker trials and 376 ms for the positive flanker trials. The CFE for the shape dimension for the current experiment was 14.80 ms, t(17) = 2.80, p <.012, while the CFE for the color dimension for the current experiment was 8.19 ms, t(17) = 1.709, p = .106. Compared to the control experiment (Experiment 1A) in which shape was 3.87 ms, t(23) = 0.984, p = .335, while the CFE for the color dimension was 12.58 ms, t(23) = 4.48, p <.001. This shows that when shape is presented first, it is able to block learning to the color dimension. False alarms were too infrequent to be analyzed.

The comparison between the shape dimension in the control experiment in which shape and color were always presented as a compound of shape flankers and color flankers and the current experiment in which there was a phase of training in which only shape flankers were presented demonstrates that experience with the shape dimension prior to presentation in a compound allows shape to block color. When shape is given a phase of training prior to presentation in the color/shape compound it is able to produce evidence of a significant CFE. This is the pattern that would be expected if shape was able to block color. The CFE for the color dimension was approximately 8 ms, but it was not significant in this experiment providing further evidence that shape was able to block color in this experiment. A demonstration of color's ability to block shape was not conducted as color was able to overshadow shape in the previous overshadowing experiment.

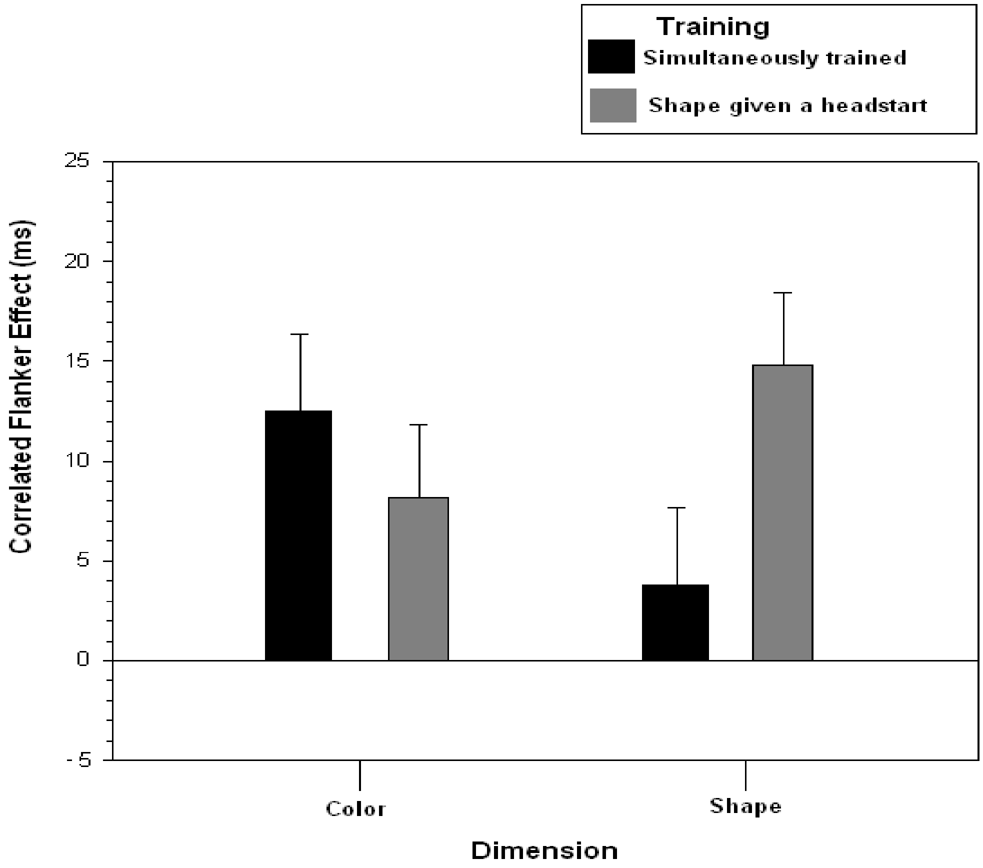

Figure 10. Results for blocking experiment. The data for the simultaneously trained condition are from Experiment 1A in which both shape and color were presented as a compound. The data from the current experiment, Experiment 2, are for each dimension when participants are given prior training with shape before the training of the color/shape compound (pictured in gray). Focus is on the difference between the conditions for each dimension.

This experiment provided evidence of blocking in human participants as would be predicted by theories assuming SWS. Blocking occurs when an organism is given experience with a stimulus prior to that stimulus being included in a compound conditioning stimulus. Recall that theories which assume SWS assume that all CSs that are presented vie for associative strength.

Rescorla-Wagner theory, for example, predicts that blocking should occur because one element in the compound is presented alone prior to the compound and this CS gains associative strength, leaving little-to-no remaining associative strength for the second element in the compound (Rescorla & Wagner, 1972). A demonstration of blocking in the current experiment comes in the form of changes in the magnitude of the CFE, a finding which was obtained.

Data from the overshadowing experiment indicates that the shape dimension is less salient than the color dimension. However, when shape flankers are presented ahead of compound stimuli, then the shape dimension exhibits the associative strength it would have gained if presented alone (cf. pilot experiment data) and the color dimension exhibits a decrement in CFE magnitude. This finding is evidence of blocking in the correlated flankers task.

Thus far, experiments which have involved the presentation of compound stimuli have shown effects that would be predicted by learning theories which assume SWS. Finding evidence of blocking with human participants is further evidence that the same mechanism underlying animal conditioning is underlying the learning occurring in the current task as well.

CHAPTER 4

OVEREXPECTATION AND CONDITIONED INHIBITION

Two more complicated conditions for compound learning involve the presentation of single CSs and compound CSs. The mixture of these two types of trials leads to interesting results. In one case, presenting single CSs which are paired with the US and then presented in a compound, reduces the associative strength of the compound (even though its components were paired with the US). This is a case of overexpectation. Another condition involves having a single CS which is paired with the US when presented alone, but is not paired with the US when presented in compound with a second CS. This case is called conditioned inhibition because the second CS acquires a negative associative strength. Both phenomena are tested here using the correlated flankers task.

Overexpectation

In choosing phenomena to import from the animal learning literature into the correlated flankers task the first consideration is choosing phenomena which show robust conditioning across species and experimental designs. Another consideration was to use phenomena that have been influential in the field for advancing theories about the action of a basic learning mechanism. Using compound conditioning phenomena is therefore an important consideration because theories which assume SWS can explain these phenomena when other theories do not. Each of the conditioning experiments proposed thus far have fit both these criteria. Another compound conditioning phenomenon demonstrated in the seminal Rescorla-Wagner (1972) paper was that of overexpectation.

Overexpectation refers to the finding that when two single stimuli (CS_A and CS_B) are each presented with a US on different trials until asymptotic levels of conditioning are reached then upon presenting these stimuli in a compound

(CS_{AB}) and pairing the compound with the US, responding to the two single CSs is reduced even though they have been continually paired with the US even in the compound (Rescorla & Wagner, 1972). For example, subjects receive trials in which a light (CS_A) is paired with food (US) and trials in which a tone (CS_B) is paired with food (US) until asymptotic responding to both is observed. Then the light and tone are presented together (CS_{AB}) and paired with food (US). Responding to either the light (CS_A) or the tone (CS_B) will now be less than it was prior to the reinforcement of the light and the tone in the compound (CS_{AB}) even though all have been paired with the food (US).

Theories that assume SWS predict that overexpectation will occur because an unconditioned stimulus can only support a finite amount of associative strength which according to the formalization of many models (e.g., Rescorla-Wagner Theory) is bounded by 0 and 1. This means that the maximum amount of associative strength that a conditioned stimulus can acquire to the unconditioned stimulus is 1. This would be an indication of perfect learning for the contingency between the CS-US pairing. Therefore, if each independent CS obtains an associative strength of 1 because they have been conditioned fully in separate trial pairings, then when put in a compound they would have an associative strength of 2. At first, the compound stimulus elicits stronger conditioned responding than either of the single stimuli. However, this cannot be supported by the US so the associative strength of the single CSs and the compound decreases until it reaches a point that can be supported by the US compared to a condition in which just the two CSs are trained independently (see Appendix E for calculations).

Overexpectation can be demonstrated with the correlated flankers task by presenting single color dimension and single shape dimension flanker trials during the training phases of the experiment. This should allow both the color dimension and the shape dimension to be conditioned. In the testing phase, both single dimensions (color and shape) will be tested, and now the compound will be presented and tested as well. Given that the testing phase is not sufficiently long enough to extinguish responding (as evidenced by the previous experiments), responding to the compound should be in its initial phase and should be stronger than either of the single stimuli. The predicted pattern of data is depicted in the following figure.

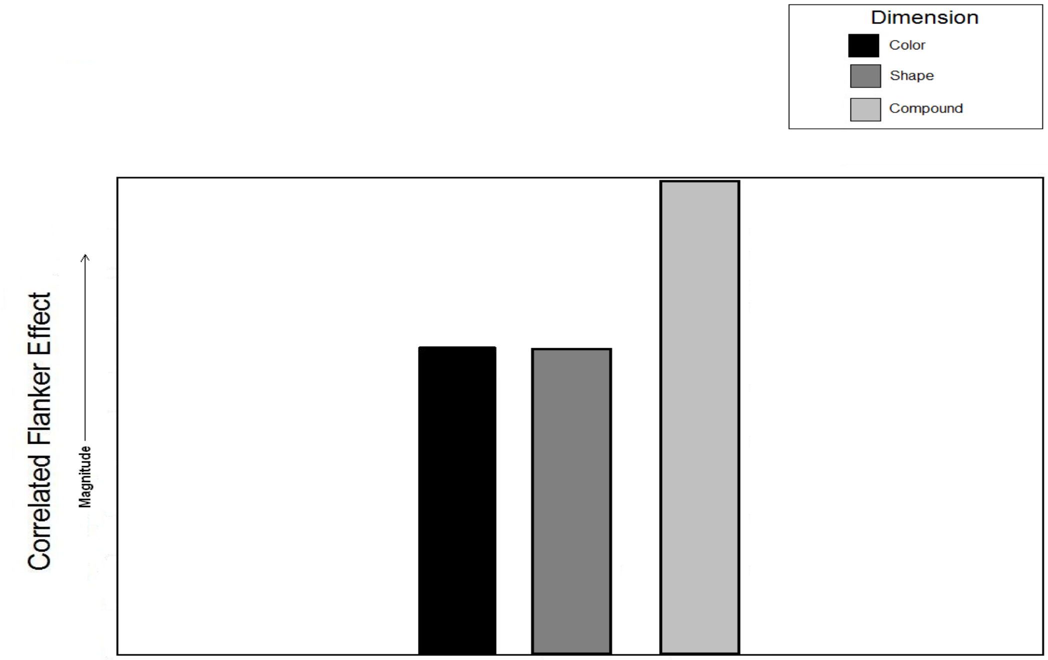

Figure 11. Predicted results for overexpectation experiment. The presented data are idealized data and so there are not specific numbers for the CFE plotted. Instead the focus is on the relative magnitudes of the conditions.

Experiment 3

Experiment 3 was designed to look at overexpectation in the correlated flankers task. In order to investigate overexpectation, there will be three types of trials: single color flanker trials, single shape flanker trials, and color and shape flanker trials.

Design

The design is illustrated in the following table (for abbreviations refer to General Method, pp. 53-55 or Appendix B, p. 112).

Table 12. Experimental design for Experiment 3.

Phase 1	**Phase 2**	**Test Phase**
C-> R	C -> R	C
S -> R	S -> R	S
		C,S

The first and second training phases were identical (six blocks of 96 trials each); the participants were all trained with color and shape flankers on separate trials. In this experiment, the testing phase differed from previous experiments because the single dimensions (color and shape) were tested on one-quarter of the trials each, but the compound of both the trained color and shape was also tested for one-half of the trials.

Participants

Twenty-four undergraduates (13 women, 11 men; ranging in age from 18 to 20 years) participated in a single session lasting about 45 minutes. All of the participants provided informed consent but were naïve as to the study's design and purpose. All reported normal or corrected-to-normal, full color vision. The majority of participants were right-handed according to self-report (3 left-handers, 21 right-handers).

Procedure

The procedure for this experiment was similar to that described in the General Method except for the training blocks and testing blocks. Participants in this experiment were not assigned to either a color or shape group as all participants saw both color or shape flankers in each display in the training blocks. The practice blocks were identical to those described in the General Method.

Training Blocks. As with the practice blocks, each trial began with the presentation of a warning stimulus which alerted the participant that the trial was starting. The warning stimulus was presented for 350 ms. In the training blocks immediately following the offset of the warning stimulus, the flankers onset for 150 ms, and then the target was presented. This final display remained on the screen until the participant responded or the 750 ms deadline had passed (for "no-go" trials), at which time feedback was given only on error trials for 1 second. The next stimulus display appeared about 1.5 seconds after the offset of the feedback for the previous trial. Any error trials were discarded and not rerun in the block. On every trial in the first six blocks of training (both phases of training), participants saw either color or shape flankers on every trial. The flankers could appear anywhere on the imaginary display, but always appeared in opposite corners to each other.

Testing Blocks. As with the testing blocks in the General Method, on every trial the color or shape flankers were presented. However, in these testing blocks one-half of the trials were compound trials (presentations of both color and shape –the ones that were correlated with each other and the same response), one-quarter of the trials were color flanker trials and one-quarter of the trials were shape flanker trials.

Results and Discussion

Mean RTs for the color dimension were 384 ms for the negative flanker trials and 371 ms for the positive flanker trials. While the mean RTs for the shape dimension were 381 ms for the negative flanker trials and 370 ms for the positive flanker trials. Mean RTs for the compound dimension were 381 ms for the negative flanker trials and 370 for the positive flanker trials. The color, shape and compound data were not statistically different from each other, F(3, 21) = 1.09, p = .776. The CFE for the color dimension was 12.83 ms, t(23) = 2.40, p = .025, the CFE for the shape dimension was 10.64 ms, t(23) = 2.15, p = .043, and the CFE for the compound of both dimensions was 12.58 ms t(23) = 3.64, p = .001. All the data presented in the graph are from the current experiment. All comparison conditions in the experiment were obtained from a single group of participants. Each participant provided a CFE for the color dimension, the shape dimension, and the compound of color and shape. In the current experiment, the CFEs for each of the dimensions were approximately equal across all conditions, averaging an 11 ms difference between the negative flanker trials and the positive flanker trials. This is not a demonstration of overexpectation as was predicted by theories that assume SWS.

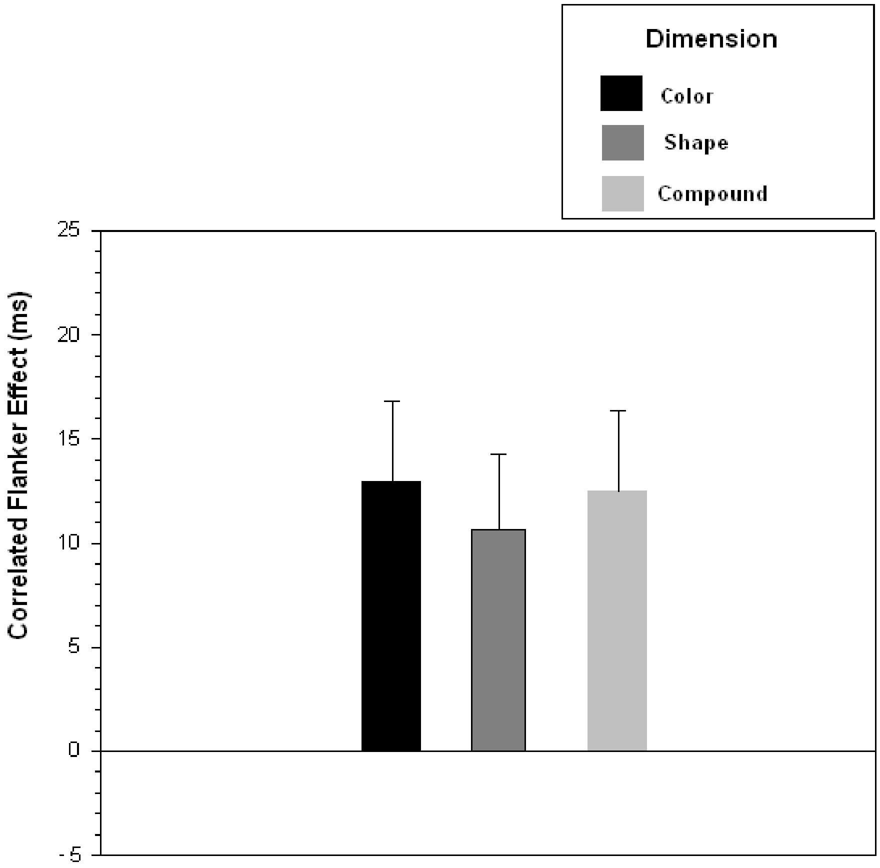

Figure 12. Results for the overexpectation experiment. Each dimension is plotted separately. The key to focus on is that there are not reliable differences in the magnitude of the CFEs for any of the single dimensions (color or shape) or the compound dimension.

Theories that assume SWS predict that the conditions of the current experiment would produce evidence of overexpectation. Specifically, given that the associative strength for the compound should have been in its initial stages, the compound should have exhibited a greater CFE than either of the two single stimuli. However, the compound produced a CFE that was almost identical to that of the color dimension. On the surface, this finding appears to indicate that the mechanism underlying the correlated flankers task is not one which uses shared associative weight space.

Delving deeper into the findings and taking all the data collected thus far into account, however, there is an alternate interpretation. It appears from the current data that participants processed the single flanker trials and both dimensions exhibited learning (i.e., demonstrated a similar magnitude CFE as when they were trained in separate groups). But on the compound trials, participants effectively processed only one stimulus dimension (color *or* shape) on each trial. This assumption comes from the breakdown of the CFEs for all the conditions. If one assumed that on average participants processed shape on half the trials and color on half the trials then the CFE for the compound trials would be the average of the CFE for the shape dimension and the CFE for the color dimension. Breaking down the experiment into its component trials makes clear why this should occur. If on half the compound trials only the color dimension was processed then the CFE for the compound for those trials would be around 13 ms. Assuming that on the other half of the trials only the shape dimension was processed then the CFE for the compound on those trials would be around 11 ms. Averaging these two CFEs produces the total CFE for the compound, approximately 12 ms. Overall, the observed CFE for the compound is approximately the average of the CFE for the shape flanker and the CFE for the color flanker. This interpretation of the data does not rule out the assumption that the color and shape flankers are sharing associative weight space, instead it can be taken as evidence of preferential processing of one dimension over another. Taking the entire body of results so far, one might ask why the compound trials weren't treated as conjunctions of the color and shape dimension and therefore treated as novel stimuli as they appeared to be in the first overshadowing experiment. The answer to this comes from the treatment of the separate dimensions. The training blocks in this experiment were composed of the single color and single flanker trials. Recall that when single flanker trials of neutral stimuli were presented in the overshadowing experiment the evidence suggests that the participants were able to process both flanker dimensions separately. In the present experiment, the processing of both flanker dimensions could have been kept separate allowing for each dimension to have an effect, but only the fastest dimension to be processed on each trial influenced responding.

Experiments from the information processing domain are highly relevant to the current work. Specifically experiments concerning divided attention. An example of a divided attention task is one in which the participant must respond with a button press whenever one or more targets is presented, and to do nothing if no targets appear. It has been shown that responses are faster and more likely when more than one target appears in the display (see Mordkoff & Miller, 1993 for a discussion). As a more concrete example, suppose that the task is to respond with a button press when the color green is in the display or when the shape "X" is in the display (or when both are in the display). Trials in which both the color green and the shape "X" are present are responded to more quickly than those trials in which only one of the targets appears. One explanation for this finding is that the activation present for both targets is pooled across processing channels and allowed to *coactivate* such that a response is triggered more quickly on trials with redundant targets. The alternative explanation for these findings is that information from both targets competes with each other and that the faster of the two processes wins the race to trigger a response. Interestingly, evidence for both explanations has been obtained, although the balance of evidence seems to favor coactivation models (see Mordkoff & Miller, 1993 for discussion).

Findings from the current experiment seem to provide evidence that activation is not pooled across stimuli, but instead competes to trigger the response as the CFE for the compound trials lies somewhere between the two CFEs for each of the single CSs instead of being larger than either of the two single CFEs. Seemingly this occurs because participants are processing only one set of flankers on each trial. Due to stochastic variation, on half the trials the color dimension is processed and on half the trials the shape dimension is processed.

Conditioned Inhibition

Contrary to most of the discussion so far which focused on phenomena that were exhibited as a conditioned response, a final phenomenon tested using the correlated flankers task will be conditioned inhibition. Conditioned inhibition involves the reduction of a CR in response to a CS. In general a stimulus becomes a conditioned inhibitor when it is presented in a compound with another stimulus that has been previously shown to be predictive of the US, but in the presence of the new CS the US now fails to occur (Schwartz et al., 2002). For example, if a light (CS_A) is paired with food (US) for a given number of trials and then a light and a tone are presented together (CS_{AB}) and now no food is present (absence of US), now the tone (CS_B) is a conditioned inhibitor which acts to reduce behavioral responding.

There are three ways in which to test for what can be termed the "absence" of a behavior. First, as advocated by Pavlov (1927), a summation test can be conducted which involves pairing the conditioned inhibitor with a stimulus that has already been conditioned as an excitor (i.e., a CS that has been trained to elicit a CR) and seeing if the conditioned inhibitor will now reduce the

CR. If this test is performed with a novel CS and the conditioned inhibitor, the effect should be similar, in that a CR to the novel CS should be reduced by the presence of the conditioned inhibitor (Schwartz et al., 2002). Another test, proposed by Rescorla (1968), is a retardation test in which the conditioned inhibitor is then paired with a US and the amount of training required to produce the CR to the conditioned inhibitor is measured. If the training in this case takes longer than for a control group, inhibition is assumed to be present (Rescorla, 1968). This test can also be conducted using an approach/withdrawal paradigm in which it is assumed that excitatory conditioning produces an approach response, while inhibitory conditioning produces a withdrawal response (Wasserman, Franklin & Hearst, 1974).

Theories that assume SWS predict that conditioned inhibition will occur because all the stimuli presented contribute to the associative strength observed. Conditioned inhibition occurs because the associative strength of the single element (CS_A) being paired with the US is being driven to an associative strength of 1, while the compound (CS_{AB}) is remaining at 0. In order for this difference to be explained, the second element CS_B must be negative, such that the associative strengths of the two elements will combine to equal zero (see Appendix E for calculations).

In the current research, a summation, retardation or approach/withdrawal test are not necessary in order to test for conditioned inhibition. In animal learning paradigms researchers are looking for the presence or absence of a behavior. Detecting the absence of the behavior necessitates a test to demonstrate the absence of a behavior. However, in the current task, the variable of interest (RT) is collected for every flanker making it possible to determine if the CS that only appears in the compound has the opposite sign from its paired counterpart. The following figure illustrates the predicted results for the conditioned inhibition experiment.

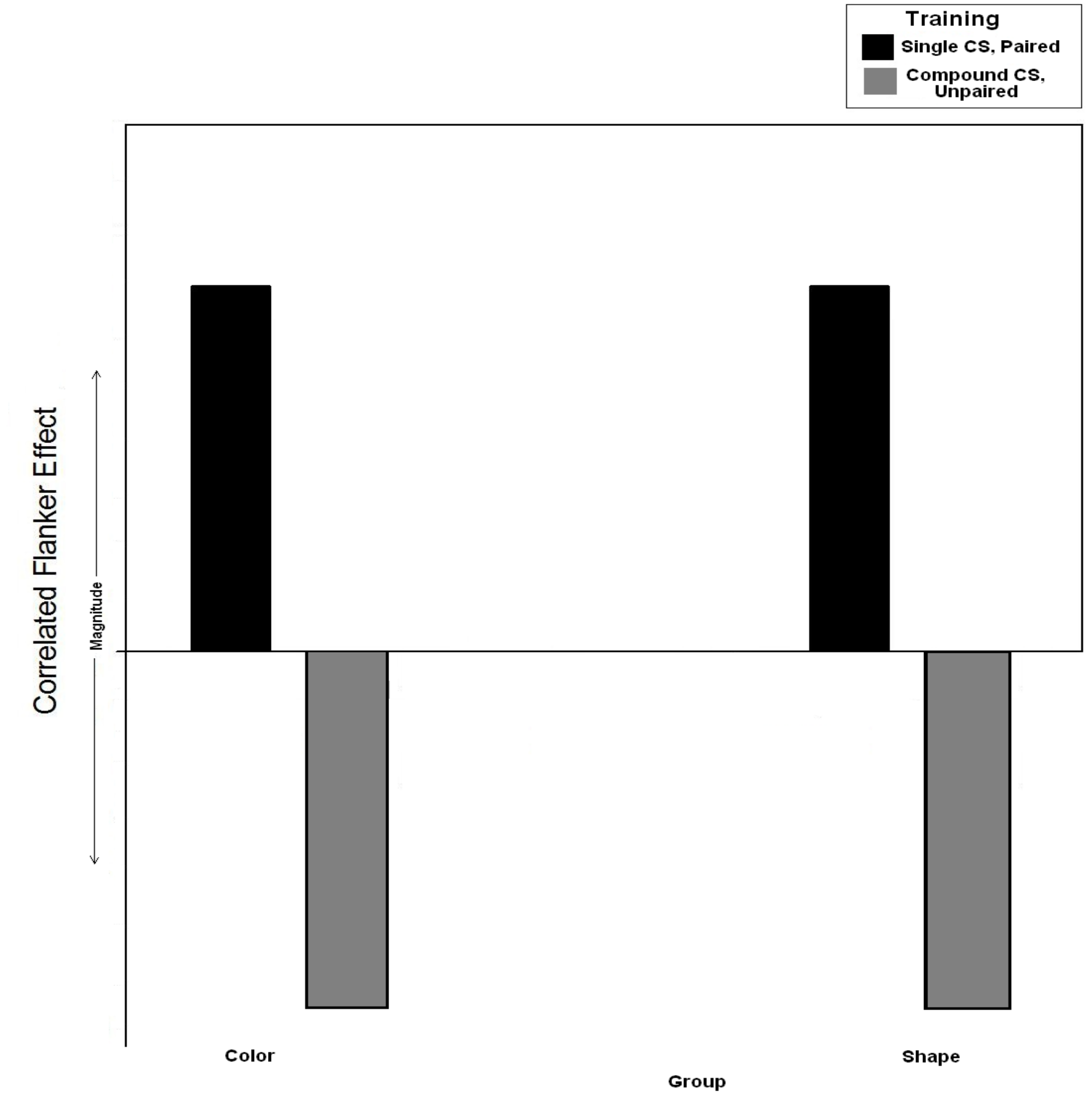

Figure 13. Predicted results for conditioned inhibition experiment. The presented data are idealized data so there are not specific numbers for the CFE plotted. Instead the focus is on the relative magnitudes of the conditions.

Experiment 4

In order to look at conditioned inhibition in the correlated flankers task, training will be conducted with the flankers in order to produce a conditioned inhibitor. Unlike the animal literature which must use one of the three tests described previously in order to show the absence of a behavior, conditioned inhibition can be tested in the correlated flankers task by looking at the RT data. If a flanker has become a conditioned inhibitor then it should show the opposite pattern of responding. If a flanker is presented in compound with the previously contingent "go" flanker and now the contingency is removed, then the single flanker when tested alone should behave as though it is negatively contingent with producing a response. So a reversal of the typical correlated flanker effect should be observed. A reversal of the typical CFE would result in the RT for the positive flanker trials being slower than the RT for the negative flanker trials.

Design

The design is illustrated in the following tables (for abbreviations refer to General Method, pp. 53-55 or Appendix B, p. 112).

Table 13. Experimental design for the color group in Experiment 4.

Phase 1	**Phase 2**	**Test Phase**
C-> R	C -> R	C
	C,S	S

Table 14. Experimental design for shape group in Experiment 4.

Phase 1	**Phase 2**	**Test Phase**
S-> R	S -> R	C
	C,S	S

Participants

Ninety undergraduates participated in a single session lasting about 45 minutes. All of the participants provided informed consent but were naïve as to the study's design and purpose. All reported normal or corrected-to-normal, full color vision. Random assignment was used to assign 45 participants to the Color Control Group (26 women, 19 men; ranging in age from 18 to 21 years) and to assign the 45 participants to the Shape Control Group (27 women, 18 men; ranging in age from 18 to 21 years). The majority of participants in both groups were right-handed according to self-report (6 left-handers, 39 right-handers for the color group and 3 left-handers, 42 right-handers for the shape group).

Procedure

The procedure for this experiment was similar to that described in the General Method except for the training blocks. Participants in this experiment were assigned to either the color or shape group. Those in the color group received single flanker trials containing colors that were positively or negatively valenced while those in the shape group received shape flanker trials that were positively or negatively valenced. The practice and testing blocks were identical to those described in the General Method.

Training Blocks. As with the practice blocks, each trial began with the presentation of a warning stimulus which alerted the participant that the trial was starting. The warning stimulus was presented for 350 ms. In the training blocks immediately following the offset of the warning stimulus, the flankers onset for 150 ms, and then the target was presented. This final display remained on the screen until the participant responded or the 750 ms deadline had passed (for "no-go" trials), at which time feedback was given only on error trials for 1 second. The next stimulus display appeared about 1.5 seconds after the offset of the feedback for the previous trial. Any error trials were discarded and not rerun in the block. Participants saw in the first phase of training saw either color or shape flankers on every trial (depending on group assignment). In the second phase of training in which the compound is introduced and no contingency is present, participants saw both single color or shape trials (depending on the group assignment) and compound trials in which both color and shape flankers were presented, but occurred equally often, such that the compound had a net result of no contingency with a response. The flankers with the same dimension (i.e., color or shape) appeared in random locations on the imaginary grid, but always appeared in opposite corners to each other.

Results and Discussion

The color group was given single color trials in which contingencies between the color dimension and a given response were present, and color/shape compound trials which occurred equally often with each response (i.e., no contingency was present). The pattern of results predicted by theories assuming SWS is one in which a positive CFE for the color dimension and a negative CFE for the shape dimension would be present (i.e., RT for the negative flanker trials would be faster than for the positive flanker trials; a reversal of the typical finding). This was not observed. Additionally, the typical magnitude of the CFE for the trained dimension (color) was not observed. The color group produced a mean CFE for the color dimension of 8.96 ms, t(44) = 3.10, p = .003 and a mean CFE for the shape dimension of -0.38 ms, t(44) = -0.115, p = .909. Mean RTs for the Color Group on the color dimension were 378 ms for the negative flanker trials and 369 ms for the positive flanker trials. While the mean RTs for the Color Group on the shape dimension were 372 ms for the negative flanker trials and 372 ms for the positive flanker trials. Mean RTs for the Shape Group on the color dimension were 364 ms for the negative flanker trials and 364 ms for the positive flanker trials. While the mean RTs for the Shape Group on the shape dimension were 370 ms for the negative flanker trials and 354 ms for the positive flanker trials. The overall error rate for both color and shape groups in this experiment was less than 1%, making any the false alarms too infrequent to be analyzed.

The shape group was given single shape trials in which contingencies were present and again the color/shape compound trials which did not have a relationship with either response. The predicted pattern for this group would be a positive CFE for the shape dimension and a negative CFE for the color dimension. Only a portion of this pattern was observed for the shape group. The mean CFE for the shape dimension was 16.14 ms, t(44) = 5. 58, p < .001 and the mean CFE for the color dimension was 0.185 ms, t(44) = 0.084, p = .934. For the shape group, the shape dimension did produce the expected positive CFE at approximately its typical magnitude, but the color dimension did not produce a significant CFE or a trend toward a negative CFE. The results are presented in the following figure. The data from this experiment are reported because it was the most successful design for the observation of conditioned inhibition. A previous experiment was conducted as well, but the data from that experiment are only reported as text because the results were not statistically significant.

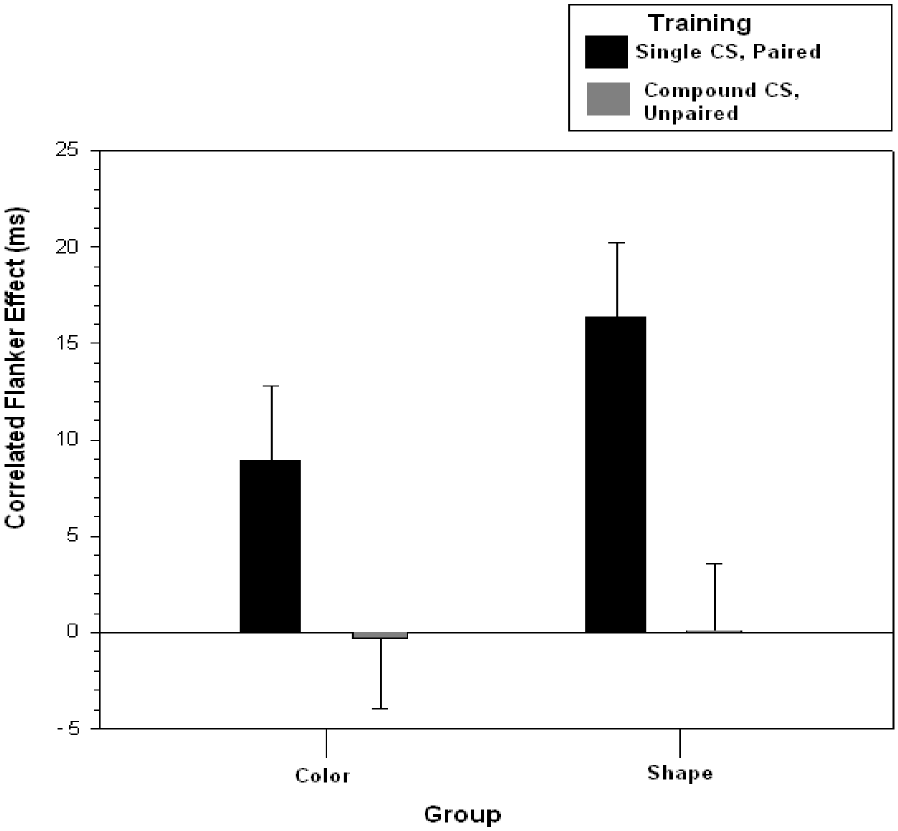

Figure 14. Results for conditioned inhibition experiment. The data presented here are from two groups of participants, the Color Group and the Shape Group. The single CS, paired data refers to the dimension that was trained for that group (color for the Color Group and shape for the Shape Group; presented as black bars). The other data are from the dimension that appeared in the unpaired compound for that group (shape for the Color Group and color for the Shape Group; presented as gray bars). The key finding to focus on is that the magnitude for the trained dimension is close to the magnitude typically observed while the magnitude for the untrained dimension is not the opposite magnitude.

As discussed previously, this experiment is the second attempt to conduct a conditioned inhibition experiment with the correlated flankers task. In the first version of this experiment (hereafter referred to as Experiment 4A), a design which might be considered more of a typical design for a conditioned inhibition experiment was conducted. In Experiment 4A, the two phases of training were identical. Therefore there were six training blocks (three in each training phase) which consisted of half the single flanker trials (which had a contingency) and half of the compound flanker trials (which did not have a contingency). The design is illustrated below.

Table 15. Experimental design for the color group in Experiment 4A.

Phase 1	**Phase 2**	**Test Phase**
C-> R	C -> R	C
C,S	C,S	S

Table 16. Experimental design for shape group in Experiment 4A.

Phase 1	**Phase 2**	**Test Phase**
S-> R	S -> R	C
C,S	C,S	S

The findings from Experiment 4A were similar to the findings reported for Experiment 4, but they were not statistically significant, even for the dimensions that had the contingency (i.e., color for the Color Group and shape for the Shape Group). Mean RTs for the Color Group on the color dimension were 378 ms for the negative flanker trials and 376 ms for the positive flanker trials. While the mean RTs for the Color Group on the shape dimension were 383 ms for the negative flanker trials and 374 ms for the positive flanker trials. Mean RTs for the Shape Group on the color dimension were 355 ms for the negative flanker trials and 357 ms for the positive flanker trials. While the mean RTs for the Shape Group on the shape dimension were 360 ms for the negative flanker trials and 350 ms for the positive flanker trials. The color group produced a CFE of 1.41 ms, t(17) = .298, p = .769 for the color dimension and a -1.77 ms CFE, t(17) = -0.523, p = .608 for the shape dimension. The shape group produced a CFE of 9.95 ms, t(17) = 2.53, p = .022 for the shape dimension and a CFE of -2.82 ms, t(17) = -0.669, p = .512 for the color dimension. These findings are similar to the finding from the current experiment (Experiment 4). However, the findings from Experiment 4A were not close to approaching significance. One possibility is that more participants needed to be run in Experiment 4A. However, at the time of conducting the experiment, this was not the first possibility considered. Instead the design of Experiment 4 was an attempt to strengthen the training for the first attribute on which responding was contingent before introducing the compound trials in which no contingencies were present. Retrospectively looking at the data from the first 36 participants in each group of Experiment 4 reveals that the trained color and shape dimensions were exhibiting a much stronger trend toward producing typically observed CFE, even with the smaller number of participants. Most likely then, the issue was not just one of power, but was one of training. A possible explanation for the specific problem with training is discussed below.

There are a few possible reasons that conditioned inhibition was not demonstrated with this task. First, it's a possibility that conditioned inhibition cannot be demonstrated with the correlated flankers task because it is not using the same mechanism underlying learning in the animal literature. However, the results from previous experiments suggest a more complicated explanation of the total package of findings.

As with the overexpectation experiment, evidence from this experiment suggests that presenting the single flankers prior to presenting the compound flankers allows the two flankers to be processed separately as opposed to as a conjunction stimulus. Again it appears that when both stimuli are then put into a compound, one of them is processed more quickly and affects the response, while the other does not. When color is trained on separate trials it obtains some associative strength commensurate with the amount of training that is occurring. When shape is then presented in the compound with color and not trained participants process the color stimulus, but do not process this second stimulus, and so when shape is tested alone evidence of learning is not obtained. Similarly, when color is presented in the compound with shape and not trained alone, participants process the shape stimulus, but do not process the second stimulus, so that no evidence of learning is obtained for the color flankers. Much like the blocking experiment, prior experience changes the processing of the stimuli. The stimulus which is trained alone gains associative strength, while the other stimulus does not. One might expect that the CFE would be weakened for the dimension that was trained alone because it is also being processed in the compound trials in which the flankers occur equally often with all the targets. Assuming that only the stimulus which has received prior training continues to be processed when presented in the compound the result is approximately half of the trials for the trained stimulus do not contain contingencies. In combination with the trials in which a contingency is present the overall contingency is weakened. This was observed for the color flanker, but not for the shape flanker, however, one hesitates to draw strong conclusions from the magnitude of the effect in this case, given that the CFE is highly variable. Though this weakening of the CFE for when the single dimensions were trained as compounds might be an explanation for why Experiment 4A did not produce a statistically significant effect. If there is not a phase of training with the single dimension in order to strengthen the processing of the single dimension it becomes greatly reduced and does not produce a CFE (as found in Experiment 4A).

CHAPTER 5

GENERAL DISCUSSION

The ultimate goal of this line of research is to understand contingency mechanisms in humans. The experiments here imported one of the major advancements in explaining generalized learning mechanisms from animal conditioning, the assumption of shared weight space (SWS), into an information processing task in order to determine if the mechanism underlying the correlated flankers task shares the same assumption. Contemporary learning theories derive much of their explanatory power from the assumption that all stimuli presented vie for associative strength, the assumption of SWS. Theories based on this assumption have proven successful in explaining many of the observed conditioning phenomena in animals. However, work with humans has proven more complex due to outside knowledge, biases and heuristics (see, e.g., Chapman, 1991; Msetfi et al., 2005; Perales et al., 2004; Tversky & Kahneman, 1974; Viken et al., 2005; Waldmann, 2000 & 2001). The present series of experiments sought to test the assumption of SWS in a task that is less susceptible to the influence of "top-down" factors. Four compound conditioning phenomena were studied in order to test this theory. Evidence for the simple predictions coming from theories assuming SWS was mixed. However, a slightly more complex version of these theories can explain the entire pattern of data quite elegantly.

The results from the overshadowing experiment provided evidence for a theory which assumes SWS. It was demonstrated that overshadowing would occur if the displays prompted an elemental processing strategy (i.e., treating each set of flankers as a separate dimension) rather than as a conjunction stimulus. This was accomplished by introducing a single set of "neutral" flankers for one phase of training before viewing the compound trials. When the single set of flankers was included prior to the trained flankers participants demonstrated overshadowing: color flankers produced the typical CFE, while shape flankers did not produce differences in RT for positive versus negative flanker trials.

The experiment to demonstrate blocking also produced an effect predicted by theories assuming SWS. Shape flankers were presented prior to color flankers in an attempt for the weaker stimulus (the shape flanker; as determined from the overshadowing experiment) to block the stronger stimulus. In this experiment, shape produced the typical CFE when it was given a phase of training in which it was the only item presented. Though the CFE for color in the blocking experiment was smaller than it was in the control experiment, there was not a significant reduction in the CFE for the blocked dimension. This finding is in-line with findings from previous research in which less salient stimuli fail to completely block a more salient stimulus (see Denton & Kruschke, 2006).

The results from the overexpectation experiment can be explained from a SWS theory as well. The extinction testing phase for the experiments was not enough to completely eliminate the CFE. If the strength of learning was still close to 1 (as evidenced by the CFE), then the predicted pattern of data would resemble the first part of an overexpectation experiment and would produce an increased CFE for the compound stimulus in relation to the single stimuli. This result was not obtained with the current experiment. However, it is possible that there are two processes affecting the observed results. One process uses an associative learning mechanism such as that described by those learning theories which assume that SWS is occurring. A second process may be occurring as well. This second process which affects perceptual processing is similar to what is found when other words are added to a Stroop display.

In a typical Stroop experiment a participant must verbally respond with the color of ink a word is presented in. If the word presented is the name of the color the of the ink the word is written in, response time is significantly less than when the word is a different color name from its ink. (Stroop, 1935/1992). However, if there are two colored words that appear in the display, one a color name and the other a neutral word, the Stroop effect is weakened – a “dilution of the Stroop effect” occurs (Kahneman & Chajczyk, 1983). One explanation for this observation is that there is a 50% chance of processing the color name first and a 50% chance of processing the neutral word first. If the color word is processed, then the typical Stroop effect is observed. If the neutral word is processed, this does not have an effect on the ability to name the color of the ink and no Stroop interference occurs. This means that the Stroop effect is approximately half of its typical strength.

In the overexpectation experiment, a variation of this pattern was observed. The compound stimulus was composed of a shape flanker and a color flanker. If only one of the dimensions was processed on a given trial then the compound CFE should be the average of the CFEs for the color and shape dimensions. On half the trials the CFE would reflect the color dimension being processed, and on half the trials the CFE would come from the shape dimension being processed. The overall CFE for the compound would then be an average of these two CFEs.

Similar to the overexpectation experiment, the conditioned inhibition experiment produced a pattern of data suggestive of two processes. The stimulus which was presented alone and was paired with responding produced a CFE that was similar to a CFE produced in the pilot experiment in which only one dimension was ever trained. There was not a negative effect for the stimulus dimension that appeared only as a part of the unpaired compound. It is as though the single trials shifted processing of the stimuli to be exclusively on the attribute that was trained alone, while the second attribute was not processed at all. Then the second attribute occurred in the test phase and still did not have a contingent relationship with responding. Given that the second attribute was likely not processed in the compound trials because prior experience with the single flanker trials biased processing to the trained attribute, there was not a chance to observe the effects of SWS in the current experiment though it could have been occurring.

Mechanism Underlying the Correlated Flanker Effect

Are these findings evidence against the idea of as associative mechanism which assumes SWS operating in humans? The body of evidence suggests that it is not. Rather, it appears that there is an associative mechanism operating as well as the differential processing of stimuli based on experience. If participants are given a biased experience with stimuli (by presenting one stimulus first or by reinforcing one stimulus over another) it changes the focus of processing until only one dimension (the stronger) is processed, which then masks the evidence of the associative processing that is occurring. This finding is similar to that of other researchers who postulate that attention plays a significant role in learning. For example, Kruschke relies on the notion of attentional processes which alter learning about cues to explain many of the observed conditioning phenomena in animals and humans (e.g., Denton & Kruschke, 2006; Kruschke, 2006; Kruschke, in press; Kruschke & Hullinger, in press). As he states,

> “Learners exhibit many apparently irrational behaviors in their use of cues, sometimes learning to ignore relevant cues or to attend to irrelevant ones…[which] can be accounted for by recent connectionist models in

> which both attentional shifting and associative learning are driven by the rational goal of rapid error reduction." (Kruschke, 2003, p. 171).

The current research also recognizes the potential role for attention to influence evidence of learning. One proposal from this work is that there is an associative mechanism operating as well as the differential processing of stimuli based on attention. This theory was specifically proposed in order to explain the findings from the overexpectation and conditioned inhibition experiments. But the idea of associative mechanisms operating in concert with attentional processing can explain the findings from the overshadowing and blocking experiments as well. Recall that the findings from the overshadowing and blocking experiments are consistent purely with associative theories, but these findings are also consistent with an additional attentional component as well. The definition of salience is often the item which is most "attention-getting" and what determines if an item is attention-getting? The salience of the item. The role of attention in the overshadowing and blocking experiments is now quite evident. Saliency determines which flankers are attended to and therefore processed. Those flankers that are processed are then the ones that demonstrate learning. This saliency can be determined by the properties of the stimulus (i.e., color being more salient than shape in the overshadowing experiment) or by experience biasing salience toward a particular dimension (i.e., experience guiding attention to the shape dimension in the blocking experiment). Thus, all findings are consistent with the proposal that for the correlated flankers task associative mechanisms are operating as well as an attentional processing component.

This task has proven to have potential to be a new and powerful technique to study contingency mechanisms in humans without the influence of explicit top-down mechanism. There have been many demonstrations of human contingency learning tasks which have been altered by outside knowledge, biases and heuristics which people are using to make the judgments (e.g., Alloy & Tabachnik, 1984; Brewer, 1974; Chapman & Chapman, 1967; Chapman, 1991; Msetfi et al., 2005; Perales et al., 2004; Smedslund, 1963; Tversky & Kahneman, 1974; Viken et al., 2005; Waldmann, 2000; Waldmann, 2001). The benefits of the correlated flankers task for studying contingency mechanisms in humans comes from the fact that the central task of the participant is to respond to a target letter, any of the learning that is done about the correlated flankers is incidental and not subject to outside knowledge, biases or heuristics. The central task (i.e., respond to the target) is used to get a response time so that it can be determined if learning is taking place and it also serves as a "distractor" task for the participants because it is a task that uses cognitive resources, and according to the instructions, is the only information on which to base a response. It is harder for participants to import outside knowledge, biases or heuristics into learning about the contingencies between the flankers and responses because there is nothing innate about the relationship any of the given colors or shapes and a particular response as evidenced by the finding that there was no difference in RT for any of the flankers.

Currently, what can we say about the mechanism underlying the correlated flanker effect? First, it does seem to be a different mechanism than that which underlies the processing of task-irrelevant information that has an arbitrary, instruction-based relationship with either stimuli or responses. This has been shown by the double dissociation that exists for standard flankers (which have an instruction-based relationship with the response) and correlated flankers. Standard flankers are able to cross case that is, if the targets are uppercase versions of letters, but the flankers are lowercase versions of those letters, the flankers still have an effect. This is not true for correlated flankers, even though the flankers share the same identity and only differ in perceptual elements. A correlated flanker effect is not obtained when training occurs with the uppercase versions of the flankers and testing is conducted with the lowercase versions (Mordkoff & Danek, in prep). Standard flankers, on the other hand, must be single features (i.e., color *or* shape), as conjunction flankers (i.e., colored shapes) do not produce an effect on RT. However, correlated flankers can be conjunctions of features and still produce an effect (Mordkoff & Halterman, 2008). This double dissociation makes it likely that the pathway for information processing of irrelevant information that has a contingent relationship with the response is different from the pathway for processing information that has an arbitrary, instruction-based relationship with the response. There has been much research into this pathway that processes explicit information. This series of experiments is the first to examine more closely the mechanism underlying the alternative pathway that processes contingency information.

There are three conclusions which come from the current experiments. Probably the most parsimonious explanation for the data is that multiple processes are occurring at the same time: attentional processes which allow focus to be restricted to a subset of the information presented and general associative properties which allow learning to occur and seem to involve SWS. Kruschke has argued that both humans and animals exhibit the influence of attention on learning (for discussion see Kruschke, 2009a). He argues that the role of attention in learning is critical and that it influences the information that is processed, not because of the inherent capacity limitation of processing but because it is advantageous to use attention in learning (Kruschke, 2006; in press, Kruschke & Hullinger, in press). The experiments presented here lead to a similar conclusion. Attentional properties might allow for a subset of the stimuli to be attended, whether this is because of something inherent about the stimulus (e.g., salience) or because the particular stimulus has been biased by experience and only this subset will have access to influencing the response.

Second, the contingency mechanism underlying the correlated flankers task does seem to be the same as the generalized learning mechanism that has been studied in animals. When similarities in behavior or conditions of learning are produced from different paradigms it strengthens the possibility that both are being served by the same mechanism. Research conducted thus far suggests that the mechanism underlying the correlated flanker effect is the same as the generalized learning mechanism which underlies other types of learning, and which has been studied in detail in the animal learning literature. The opportunity to import all of the knowledge that we have from this area is a benefit due to the large number of studies conducted in developing a theory about the generalized learning mechanism in animals.

Finally, in addition to being able to bring what is known from the animal literature to bear on the correlated flankers task, the correlated flankers task can also be used to investigate the generalized learning mechanism in humans without having to be concerned with outside knowledge, biases and heuristics which might be interfering with the true results from human contingency learning tasks. These findings are a boon for the animal learning literature as well because the correlated flankers task can be used to investigate a learning mechanism in humans.

Research Applications

Beyond the basic research applications of the present work which include the bridging of two areas of research in an effort to understand more clearly the types of questions that this research should be asking and answering, as well as providing an additional method for looking at learning mechanisms in humans, the current research has other applications as well. Though the current research discusses the topic of “learning”, there are not clear cut implications about improving how people learn, for example, in structured training situations, that come from this research. Instead the current research is informative about the nature of processing of contingency information in the environment. This research points to the idea that redundant environmental information may not be helpful in improving people’s reaction to the information. As an example, there are multiple stoplights on each pole in the intersection. These stoplights each act as a cue about the correct driving maneuver to execute. The current research would suggest that there is not an advantage (e.g., speed up in reaction time) for the redundant lights because most likely only one of these cues is being processed and affecting response time. The suggestion from the current research is to improve cues in the environment which indicate the appropriate behavior. For example, when driving a car or engaging military targets. Instead of including multiple cues (e.g., warning lights and buzzers) it is more important to focus on providing a single salient cue which most accurately indicates the appropriate behavior.

Most importantly this research offers new insights into the mechanism underlying processing contingency information in the environment. Every day, humans are bombarded with many different cues that may or may not be predictive of upcoming events. Being able to understand how these cues are processed is a building block for designing more efficient computer interfaces, training scenarios, and environments conducive to optimum performance.

APPENDIX A

DETAILS OF PILOT EXPERIMENT

Flanker Stimulus	**Mean RT (ms)**	**Error Rate (proportion)**
Red	380.27	0.009
Green	381.07	0.007
Yellow	381.02	0.008
Blue	378.51	0.008
Pound Sign	384.22	0.008
Tilted pound sign	382.26	0.007
Square	376.60	0.009
Diamond	378.55	0.003

There was no difference among the means, F(1, 7) = 0.15, p = .994. The average response time was 380 ms across all of the dimensions.

APPENDIX B

NOTATIONS FOR EXPERIMENTAL DESIGNS

The following notations will be used to capture the relationships between the flankers and the responses in each experiment. These relationships are illustrated in the various design tables for each of the experiments. **C** stands for color, **S** stands for shape and **R** stands for response. **C -> R** will be used to indicate when the color flankers have a contingent relationship with the response. **S -> R** indicates the shape flankers have a contingent relationship with the response. **C,S -> R** will be used to indicate a compound trial in which both shape and color flankers are presented simultaneously and have a contingent relationship with the response. **C->R, S->R** is used to indicate that there are single trials which have color flankers and single trials which have shape flankers. Both flankers have contingent relationships with the response though each is presented on its own. Finally, a **C** or **S** not followed by an **R** indicates that there is not a contingent relationship between the flankers and the response. This will occur primarily in the testing phase though it can occur in the training phase if there is not a relationship between the flankers and the responses.

APPENDIX C

MEAN RESPONSE TIMES BY CONDITION FOR EXPERIMENTS

	Mean RTs by Condition	
Experiment	**Negative Flankers (ms)**	**Positive Flankers (ms)**
Pilot Data (Control 11) -Color group/color dimension	386	372
Pilot Data (Control 11) -Color group/shape dimension	383	377
Pilot Data (Control 11) -Shape group/color dimension	380	380
Pilot Data (Control 11) -Shape group/shape dimension	387	369
Experiment 1 (Pilot/ Overshadow 1) - Color	383	378
Experiment 1 (Pilot/ Overshadow 1) - Shape	378	376
Experiment 1A (Pilot / Block Control) - Color	381	368
Experiment 1A (Pilot / Block Control) - Shape	379	375
Experiment 1B (Block Control/Overshadow Control) - Color	376	365
Experiment 1B (Block Control/Overshadow Control) - Shape	374	361
Experiment 2 (Block Control/ Blocking - Shape Group - Color)	389	381
Experiment 2 (Block Control/ Blocking -Shape Group - Shape)	391	376
Experiment 3 (Overexpectation -Color)	384	371
Experiment 3 (Overexpectation -Shape)	380	369
Experiment 3 (Overexpectation -Compound)	381	370
Experiment 4 (Conditioned Inhibition -Color group / color)	378	369
Experiment 4 (Conditioned Inhibition -Color group/ Shape)	372	372
Experiment 4 (Conditioned Inhibition -Shape group/ color)	364	364
Experiment 4 (Conditioned Inhibition -Shape group/ Shape)	370	354

APPENDIX D

SAMPLE COMMENTS REGARDING AWARENESS

The following are a representative sample of the comments that came from awareness surveys. These questions were asked for all experiments and a representative sample of those responses (as opposed to each response) is reported here. Overall, most people could identify at least some of the flankers presented both colors and shapes. People did not seem to be aware of any contingencies in the training phases. Some people seemed to notice that there were additional flankers or a change to the presentation of the flankers for the testing phase (recall that in the testing phase only color flankers or shape flankers were presented on each trial). Though there was not a relationship to report for the flankers in the testing phase, participants were not able to identify any relationships. Most reported that they didn't notice any relationship (just as they reported for the training phase.

1. For most of the experiment, what items appeared diagonally to the letter in the center (name them all)?
    - "Daimonds[sic], pound key (#), colors"
    - "Colored blobs, # signs, circles and squares"
    - "Diamond, square, #, colored blobs"
    - [Accurate drawing of stimuli]
    - "Boxes, colored shapes"

2. Did you notice that any of these outside items were associated with either of the two responses in particular (either responding or not responding)? In other words, when certain outside items appeared the response was

almost always the same. Do you know which outside items these were and which response went with each one?

   - “No”
   - “No, I did not see a pattern for response and no response”
   - “I do not know”
   - “I didn’t notice”
   - “I don’t remember seeing any sign-letter correlation”

3. In the very last few blocks of the experiment did you notice any changes to the outside items? If you did, what were these changes?

- “ I didn’t notice any change”
- “No”
- “They were less colorful, mostly circles and squares”
- “Sometimes the colored blobs appeared and sometimes they didn’t”
- “Colored spots were added”

4. In the last few blocks did the outside items appear more often with one response or the other? If they did, please describe the relationship.
   - “Not sure”
   - “No”
   - “No, I didn’t notice”
   - “I don’t know”
   - “ I saw no item appear more often than others”

APPENDIX E

CALCULATIONS FOR PHENOMENA

Overshadowing Demonstration (Different Salience Assumed) with Rescorla-Wagner Theory

Assume β = learning rate parameter = 1.0
λ = asymptote of learning = 1.0
α = stimulus salience for stimuli are different = 0.5 for stronger stimulus; 0.25 for weaker

Trials present the compound of both stimuli, CS_{AX}

Trial 1: $\Delta V_A = \alpha_A \beta (\lambda - V_{AX})$
$\Delta V_A = 0.5(1.0)(1.0\text{-}0)$
$\Delta V_A = 0.5$
$V_A = 0 + 0.5 = 0.5$

$\Delta V_X = \alpha_X \beta (\lambda - V_{AX})$
$\Delta V_X = 0.25(1.0)\ (1.0 - 0)$
$\Delta V_X = 0.25$
$V_X = 0 + 0.25 = 0.25$

Trial 2: $\Delta V_A = \alpha_A \beta (\lambda - V_{AX})$
$\Delta V_A = 0.5(1.0)(1.0\text{-}0.75)$
$\Delta V_A = 0.125$
$V_A = 0.5 + 0.125 = 0.625$

$\Delta V_X = \alpha_X \beta (\lambda - V_{AX})$
$\Delta V_X = 0.25(1.0)\ (1.0 - .75)$
$\Delta V_X = 0.0625$
$V_X = 0.25 + 0.0625 = 0.3125$

Trial 3: $\Delta V_A = \alpha_A \beta (\lambda - V_{AX})$
$\Delta V_A = 0.5(1.0)(1.0\text{-}0.9375)$
$\Delta V_A = 0.03125$
$V_A = 0.625 + 0.03125 = 0.65625$

$\Delta V_X = \alpha_X \beta (\lambda - V_{AX})$
$\Delta V_X = 0.25(1.0)\ (1.0 - .9375)$
$\Delta V_X = 0.015625$
$V_X = 0.3125 + 0.015625 = 0.328125$

Trial 4: $\Delta V_A = \alpha_A \beta (\lambda - V_{AX})$
$\Delta V_A = 0.5(1.0)(1.0\text{-}0.984375)$
$\Delta V_A = 0.0078$
$V_A = 0.65625 + 0.0078 = 0.664$

$\Delta V_X = \alpha_X \beta (\lambda - V_{AX})$
$\Delta V_X = 0.25(1.0)\ (1.0 - 0.984375)$
$\Delta V_X = 0.0039$
$V_X = 0.328125 + 0.0039 = 0.332$

$V_A = 0.664$
$V_X = 0.332$

In this way the more salient stimulus presented overshadows learning to the less salient stimulus presented. This difference would continue to grow as the learning asymptotes (when there is no longer a change in the associative strength; ΔV_A and $\Delta V_X = 0$).

Overshadowing Demonstration (Same Salience Assumed) with Rescorla-Wagner Theory

Assume β = learning rate parameter = 1.0
λ = asymptote of learning = 1.0
α = stimulus salience for stimuli = 0.25

Trials present the compound of both stimuli, CS_{AX}

Trial 1: $\Delta V_A = \alpha_A \beta (\lambda - V_{AX})$
$\Delta V_A = 0.25(1.0)(1.0-0)$
$\Delta V_A = 0.25$
$V_A = 0 + 0.25 = 0.25$

$\Delta V_X = \alpha_X \beta (\lambda - V_{AX})$
$\Delta V_X = 0.25(1.0)\ (1.0 - 0)$
$\Delta V_X = 0.25$
$V_X = 0 + 0.25 = 0.25$

Trial 2: $\Delta V_A = \alpha_A \beta (\lambda - V_{AX})$
$\Delta V_A = 0.25(1.0)(1.0-0.5)$
$\Delta V_A = 0.\ 25$
$V_A = 0.25 + 0.25 = 0.5$

$\Delta V_X = \alpha_X \beta (\lambda - V_{AX})$
$\Delta V_X = 0.25(1.0)\ (1.0 - 0.5)$
$\Delta V_X = 0.25$
$V_X = 0.25 + 0.25 = 0.\ 5$

Trial 3: $\Delta V_A = \alpha_A \beta (\lambda - V_{AX})$
$\Delta V_A = 0.25(1.0)(1.0-1.0)$
$\Delta V_A = 0$
$V_A = 0.5 + 0 = 0.5$

$\Delta V_X = \alpha_X \beta (\lambda - V_{AX})$
$\Delta V_X = 0.25(1.0)\ (1.0 - 1.0)$
$\Delta V_X = 0$
$V_X = 0.5 + 0.0 = 0.5$

$V_A = 0.5$
$V_X = 0.5$

In this way both stimuli exhibited mutual overshadowing.

Blocking Demonstration with Rescorla-Wagner Theory

Assume β = learning rate parameter = 1.0
λ = asymptote of learning = 1.0
α = stimulus salience for all stimuli = 0.5 for both

Trials 1 & 2 present only CS_A, Trials 3 & 4 present the compound CS_{AX}

Trial 1: $\Delta V_A = \alpha_A \beta (\lambda - V_{AX})$
$\Delta V_A = 0.5(1.0)(1.0\text{-}0)$
$\Delta V_A = 0.5$
$V_A = 0 + 0.5 = 0.5$

Trial 2: $\Delta V_A = \alpha_A \beta (\lambda - V_{AX})$
$\Delta V_A = 0.5(1.0)(1.0\text{-}0.5)$
$\Delta V_A = 0.25$
$V_A = 0.5 + 0.25 = 0.75$

Trial 3: $\Delta V_A = \alpha_A \beta (\lambda - V_{AX})$
$\Delta V_A = 0.5(1.0)(1.0\text{-}0.75)$
$\Delta V_A = 0.125$
$V_A = 0.75 + 0.125 = 0.875$

$\Delta V_X = \alpha_X \beta (\lambda - V_{AX})$
$\Delta V_X = 0.5(1.0)\ (1.0 - .75)$
$\Delta V_X = 0.125$
$V_X = 0 + 0.125 = 0.125$

Trial 4: $\Delta V_A = \alpha_A \beta (\lambda - V_{AX})$
$\Delta V_A = 0.5(1.0)(1.0\text{-}1)$
$\Delta V_A = 0$
$V_A = 0.875 + 0 = 0.875$

$\Delta V_X = \alpha_X \beta (\lambda - V_{AX})$
$\Delta V_X = 0.5(1.0)\ (1.0 - 1)$
$\Delta V_X = 0$
$V_X = 0.125 + 0 = 0.125$

$V_A = 0.875$
$V_X = 0.125$

In this way the first stimulus presented blocks learning to the second stimulus presented.

Overexpectation Demonstration with Rescorla-Wagner Theory

Assume β = learning rate parameter = 1.0
λ = asymptote of learning = 1.0
α = stimulus salience for all stimuli = 0.5 for both

Assuming that CS_A and CS_X have been previously presented on single trials and have obtained V_A and V_X of 1 the compound, CS_{AX}, is presented.

Trial 1: $\Delta V_A = \alpha_A\beta\,(\lambda - V_{AX})$
$\Delta V_A = 0.5(1.0)(1.0 - 2.0)$
$\Delta V_A = -0.5$
$V_A = 1 + -0.5 = 0.5$

$\Delta V_X = \alpha_X\beta\,(\lambda - V_{AX})$
$\Delta V_X = 0.5(1.0)\,(1.0 - 2.0)$
$\Delta V_X = -0.5$
$V_X = 1 + -0.5 = 0.5$

Trial 2: $\Delta V_A = \alpha_A\beta\,(\lambda - V_{AX})$
$\Delta V_A = 0.5(1.0)(1.0-1.0)$
$\Delta V_A = 0$
$V_A = 0.5 + 0 = 0.5$

$\Delta V_X = \alpha_X\beta\,(\lambda - V_{AX})$
$\Delta V_X = 0.5(1.0)\,(1.0 - 1.0)$
$\Delta V_X = 0$
$V_X = 0.5 + 0 = 0.5$

$V_A = 0.5$
$V_X = 0.5$
$V_{AX} = 1.0$

In the example presented here, both stimuli have the same salience and it is assumed that the learning rate is rapid. This allows for the quick drop of the associative strength from 2 to 1 which can be supported by the US.

Conditioned Inhibition Demonstration with Rescorla-Wagner Theory

Assume β = learning rate parameter = 1.0
λ = asymptote of learning = 1.0
α = stimulus salience for all stimuli = 0.5 for both

Stimulus CS_A is used to indicate the presentation of the single stimulus when it is paired with the US. $CS_{A'}$ is used to indicate when there is the presentation of the CS_A without the US presentation. CS_X is the second stimulus which always appears in a compound and is never paired with the US.

Trial 1: $\Delta V_A = \alpha_A \beta (\lambda - V_{AX})$
$\Delta V_A = 0.5(1.0)(1.0 - 0)$
$\Delta V_A = 0.5$
$V_A = 0 + 0.5 = 0.5$

$\Delta V_{A'} = \alpha_{A'} \beta (\lambda - V_{AX})$
$\Delta V_{A'} = 0.5(1.0)\ (0 - 0)$
$\Delta V_{A'} = 0$
$V_{A'} = 0 + 0 = 0$

$\Delta V_X = \alpha_X \beta (\lambda - V_{AX})$
$\Delta V_X = 0.5(1.0)\ (0 - 0)$
$\Delta V_X = 0$
$V_X = 0 + 0 = 0$

Trial 2: $\Delta V_A = \alpha_A \beta (\lambda - V_{AX})$
$\Delta V_A = 0.5(1.0)(1.0 - 0.5)$
$\Delta V_A = 0.25$
$V_A = 0.5 + 0.25 = 0.75$

$\Delta V_{A'} = \alpha_{A'} \beta (\lambda - V_{AX})$
$\Delta V_{A'} = 0.5(1.0)\ (0 - 0.5)$
$\Delta V_{A'} = -0.25$
$V_{A'} = 0 + -0.25 = -0.25$

$\Delta V_X = \alpha_X \beta (\lambda - V_{AX})$
$\Delta V_X = 0.5(1.0)\ (0 - 0.5)$
$\Delta V_X = -0.25$
$V_X = 0 + -0.25 = -0.25$

Trial 3: $\Delta V_A = \alpha_A \beta (\lambda - V_{AX})$
$\Delta V_A = 0.5(1.0)(1.0 - 0.25)$
$\Delta V_A = 0.375$
$V_A = 0.25 + 0.375 = 0.625$

$\Delta V_{A'} = \alpha_{A'} \beta (\lambda - V_{AX})$
$\Delta V_{A'} = 0.5(1.0)\ (0 - 0.25)$
$\Delta V_{A'} = -0.125$
$V_{A'} = 0.25 + -0.125 = 0.125$

$\Delta V_X = \alpha_X \beta (\lambda - V_{AX})$
$\Delta V_X = 0.5(1.0)\ (0 - 0.25)$
$\Delta V_X = -0.125$
$V_X = -0.25 + -0.125 = -0.375$

Trial 4: $\Delta V_A = \alpha_A \beta (\lambda - V_{AX})$
$\Delta V_A = 0.5(1.0)(1.0 - 0.375)$
$\Delta V_A = 0.3125$
$V_A = 0.625 + 0.3125 = 0.625$

$\Delta V_{A'} = \alpha_{A'} \beta (\lambda - V_{AX})$
$\Delta V_{A'} = 0.5(1.0)\ (0 - 0.375)$
$\Delta V_{A'} = -0.1875$
$V_{A'} = 0.125 + -0.1875 = -0.0625$

$\Delta V_X = \alpha_X \beta (\lambda - V_{AX})$
$\Delta V_X = 0.5(1.0)(0 - 0.375)$
$\Delta V_X = -0.1875$
$V_X = -0.375 + -0.1875 = -0.5625$

$V_{A\ total} = 0.625 + -0.0625 = 0.5625$
$V_X = -0.5625$

This procedure would continue until the change in the associative strength reached asymptote. This produces a positive value for the stimulus that is presented alone and paired with the US and a negative value for the stimulus that always appears in the compound and is never paired with the US.

APPENDIX F
JUSTIFICATIONS FOR DESIGNS

As this is a new area of research (or new instantiation of two types of research), it is incumbent upon me to explain not only the successes of the research, but also the failures. This appendix will detail some of the areas of the research which proved to be quite difficult to accomplish so that others seeking to do this research can avoid these pitfalls.

Preliminary work with the correlated flanker effect was conducted using the original design as proposed by Miller (1987). Miller (1987) used a two-alternative-forced-choice task in which three targets were mapped onto a left button press response and three other targets were mapped onto the right button press response. This design was optimal because it provided a way to minimize other known effects. For example, it has been shown in previous research that increasing the probability of an expected display speeds response time to the display (i.e., display frequency bias; Miller & Pachella, 1973).

Because of the display frequency bias, Miller's (1987) design used two types of targets. One set of targets were "inducing targets" which appeared with a specific flanker on 30 out of 32 trials per block. For example, Flanker 1 might appear with Targets 1 and 2 on 30 out of 32 trials in a block whereas Flanker 2 would appear with Targets 4 and 5 on 30 out of 32 trials. These targets are inducing targets because they were not used for data analysis, but were only used to induce the contingency. The other two targets appeared equally often with both flankers (16 trials with each flanker). These targets were the "test targets" because they did not have any display frequency bias. Additionally, this design ensured there was a contingency between the flanker and the response, not a particular flanker and the test targets. In order to determine the effect of the flankers, the correlated flanker effect (CFE) was calculated. The CFE is a comparison of the RT from trials on which the flanker correlated with the correct response was present in the display and the RT from trials on which the flanker correlated with the opposite response was present in the display (Miller, 1987).

Miller's (1987) original design allowed for data from both test targets to be used for analysis providing him with many observations per participant. In an effort to make the correlated flankers task more parallel to the types of tasks used in animal conditioning (i.e., detecting presence or absence of behaviors), I used a modified task in which three targets were mapped onto one response, while the other three targets signaled that no response was to be made (i.e., a "go/no-go" task). This meant that from the onset I was working with half the data that was possible in the original experiment (because RTs are only collected for the "go" trials which are half of the experimental trials). In pilot studies it was found that reducing the number of trials on which RT was collected caused the correlated flanker effect to be weak and often not significant because of the high variability due to the small number of observations.

In an attempt to decrease the variability it was necessary to optimize the conditions needed for detecting the CFE. One could either run a large number of people in a given experiment or collect more observations from each participant. In the interest of efficiency I went with collecting more observations from each participant. Thus the design was changed to an experiment in which all targets were essentially "test targets" in that each had a biased display frequency- a particular flanker was paired with a particular target set. However, in the critical testing phase each display (combination of flankers and targets) occurred equally often. In this way any problems with display frequency effects were avoided (see Miller & Pachella, 1973 for discussion) in the data being analyzed. Data from the training blocks in which there were display frequency biases were not kept.

Another way to increase the probability of detecting a CFE was to increase the efficacy of the correlated flankers. It has been shown that providing irrelevant information slightly prior to the relevant information increases the effect of the irrelevant information (e.g., Eriksen & Shultz, 1979; Halterman, 2006). This has been shown both for irrelevant information that has an instruction-based relationship with the response (i.e., standard flankers; Eriksen & Shultz, 1979) and for irrelevant information that has a contingent relationship with the response (Halterman, 2006). One plausible explanation of the increased efficacy of the irrelevant information when presented slightly prior to the critical information is that increases in the viewing time of the irrelevant stimulus allow it to overcome deficits in visual acuity. These processing deficits are a product of being presented to either side of the fovea (i.e., off of fixation). Response time increases as the distance of a stimulus from foveal fixation increases (Eriksen & Schultz, 1977). In the case of the flankers that have a contingent relationship with the response presenting the flankers ahead of the target is more akin to typical conditioning procedures in which the presentation of the CS prior to the US is more likely to produce conditioning than if the CS-US pair is presented simultaneously (Schwartz et al., 2002). In the original correlated flankers task the target and flankers onset simultaneously, but in the current experiments the flankers preceded the target by 150 ms.

In the first experiment conducted in this series it was logical to present the sets of flankers in either a horizontal or vertical strip in relation to the target. This was done because typical displays generally consist of two identical flankers horizontally flanking the target item. Because later experiments would involve the presentation of multiple flanker sets simultaneously it was important to include both axes. However, when the data were analyzed there was an effect of presentation axis such that the flankers presented horizontally had a greater effect than those presented vertically. One explanation for this effect comes from the general bias for readers of western languages to read from left to right which makes the horizontal axis more salient. This bias has been shown to exist in other research as well (see Abed, 1991 for review). Because there was a difference depending on the axis of presentation the data had to be split which worked counter to getting more observations in a given cell for analysis.

The next experiment presented the flankers on diagonal axes removing the horizontal bias to the flankers. This experiment was also changed to include a perfect contingency for a given flanker and a given target/response in an effort to increase the observed CFE, but this produced "awareness" in participants as evidenced by incredibly short RTs. Participants initiated responding prior to the onset of the target as was evidenced from RTs of 50 ms or so which does not happen even in the simplest cases of responding (i.e., a simple-RT task) (e.g., Brebner & Welford, 1980; Galton, 1889; Welford, 1980). This means that they were not responding to the targets, but to the flankers which defeats the purpose of the task to involve incidental learning.

In an attempt to decrease awareness in the participants for the irrelevant information coming from the flankers, a 3:1 contingency was used. This meant that a given flanker appeared with a given set of targets 3 out of 4 times in a block. Interestingly, this experiment did not produce any evidence of the correlated flanker effect. This is strange because Miller's contingencies were approximately the same strength in some experiments, and he was still able to obtain a correlated flanker effect. However, in the current experiment the 3:1 contingency did not produce an effect.

Finally, the design changed to include longer blocks with forced midblock breaks. This allowed a 7:1 contingency to be used. This means a given flanker appeared with a given target set 7 out of 8 trials in a block. This change in the design produced a CFE, but did not seem to affect awareness as the participants produced typical RTs, and when probed with a questionnaire afterward were unable to identify the contingencies present. These failures to find the CFE without awareness of the participant beg a couple of questions.

Why Don't 3:1 Contingencies Work?

It is possible that participants did learn the contingencies of the 3:1 design, but the contingency was so weak that when tested in a phase in which every combination of targets and flankers occurred equally often, learning was extinguished too rapidly to be detected. This doesn't seem plausible though because when looking at the training blocks, the participants did not demonstrate a correlated flanker effect either. Perhaps there is some critical threshold that exists which must be passed in order for the system to detect contingencies and the 3:1 contingency does not meet this threshold.

Why Do Perfect Contingencies Produce Awareness in Participants?

Interestingly, using a design in which there was a perfect contingency between the flankers and targets seemed to produce "awareness" of this contingency in the participants. It has been argued elsewhere that awareness of contingencies can change a participant's willingness to make responses (e.g., Brewer, 1974). It seems obvious that a perfect contingency would produce awareness because the flanker was perfectly predictive of the target. We can tell from the data that it did affect the participants because their RTs were much faster than is typical for this task or even for a simple-RT task. Was this effect at the level of "consciousness"? Were participants consciously aware while performing the task of the contingency between the flankers and the target set /response or were participants simply able to reflect back when proved on questionnaires? Most likely awareness was present consciously while they were performing the task which presents a problem for this type of investigation (as previously discussed).

Recommendations for Conducting Correlated Flanker Effect Studies

One recommendation is to use a high contingency relation (7:1). This appears to be critical for the detection of a correlated flanker effect. If the contingency is not high enough then no conditioning occurs. This is likely due to the fact that the flankers are irrelevant information which participants are told to ignore, and the contingency must be strong enough to be learned incidentally in conjunction with the task at hand. Another recommendation is to use a task in which more observations per condition can be made, for example, using a two-alternative- forced-choice task which has the added advantage of not having an inhibitory component (see for example, Verbruggen & Logan, 2008 for review; and also Ulrich, Mattes, & Miller, 1999). Although this inhibitory component does not seem to produce differences in responding between go/no-go and two-alternative-forced choice tasks in humans (A. Poremba, personal communication, April 14, 2009). Finally, another recommendation is to reduce the repetition effect for the flankers by always alternating the flankers from trial N-1 to N. This change removes any confounding effects of repetition from the design.

REFERENCES

Abed, F. (1991). Cultural influences on visual scanning patterns. *Journal of Cross-Cultural Psychology, 22,*(4), 525 – 534.

Allan, L. G., Siegel, S., & Hannah, S. (2007). The sad truth about depressive realism. *The Quarterly Journal of Experimental Psychology, 60*(3), 482 – 495.

Alloy, L. B., & Abramson, L. Y. (1979). Judgment of contingency in depressed and nondepressed students: Sadder but wiser?. *Journal of Experimental Psychology: General, 108*(4), 441 – 485.

Alloy, L. B., & Tabachnik, N. (1984). Assessment of covariation by humans and animals: The joint influence of prior expectations and current situational information. *Psychological Review, 91*(1), 112 – 149.

Arcediano, F., Matute, H., & Miller, R. R. (1997). Blocking of Pavlovian Conditioning in Humans. *Learning and Motivation, 28,* 188-199.

Benson, H., Beary, J. F., & Carol, M. P. (1974). The Relaxation Response. *Psychiatry, 37,* 37- 46.

Brebner, J. T., & Welford, A. T. (1980). Introduction: an historical background sketch. In A. T. Welford (Ed.), *Reaction Times* (pp. 1 - 23). New York, NY: Academic Press.

Brembs, B., Baxter, D. A., & Byrne, J. H. (2004). Extending in vitro conditioning in *Aplysia* to analyze operant and classical processes in the same preparation. *Learning and Memory, 11,* 412 – 420.

Brembs, B., Lorenzetti, F. D., Reyes, F. D., Baxter, D. A., & Byrne, J. H. (2002). Operant reward learning in *Aplysia*: Neuronal correlates and mechanisms. *Science, 296*, 1706 – 1709.

Brewer, W. F. (1974). There is no convincing evidence for operant or classical conditioning in adult humans. In W. B. Weimer (Ed.), *Cognition and the symbolic processes* (pp.1-35). Hillsdale, N.J: Lawrence Erlbaum Associates.

Chapman, G. B. (1991). Trial order affects cue interaction in contingency judgment. *Journal of Experimental Psychology: Learning, Memory & Cognition, 17*(5), 837 – 854.

Chapman, L. J., & Chapman, J. P. (1967). Genesis of popular but erroneous psychodiagnostic observations. *Journal of Abnormal Psychology, 72*(3), 193 – 204.

Chapman, G. B., & Robbins, S. J. (1990). Cue interaction in human contingency judgment. *Memory & Cognition, 18*(5), 537 – 545.

Cohen, A., & Shoup, R. (1997). Perceptual dimensional constraints in response selection processes. *Cognitive Psychology, 32,* 128 – 181.

De Houwer, J., & Beckers, T. (2002). A review of recent developments in research and theories on human contingency learning. *Quarterly Journal of Experimental Psychology, 55B*, 289-310.

Denton, S. E., & Kruschke, J. K. (2006). Attention and salience in associative blocking. *Learning and Behavior, 34*(3), 285 - 304.

Dickinson, A., Shanks, D. R., & Evenden, J. L. (1984). Judgment of act-outcome contingency: The role of selective attribution. *Quarterly Journal of Experimental Psychology, 36A*, 29 – 50.

Eriksen, B. A. & Eriksen, C. W. (1974). Effects of noise letters upon identification of a target letter in a nonsearch task. *Perception & Psychophysics 16*, 143-149.

Eriksen, C. W., & Schultz, D. W. (1977). Retinal locus and acuity in visual information processing. *Bulletin of the Psychonomic Society*, *9*, 81 – 84.

Eriksen, C. W., & Schultz, D. W. (1979). Information processing in visual search: A continuous flow conception and experimental results. *Perception & Psychophysics, 25*(4), 249 – 263.

Eriksson, D., & Wallin, L. (1986). Male bird song attracts females – a field experiment. *Behavioral Ecology and Sociobiology, 19*, 297 - 299.

Galton, F., (1899). On instruments for (1) testing perception of differences of tint and for (2) determining reaction time. *Journal of the Anthropological Institute,19*, 27 - 29.

Garcia, J., & Koelling, R. A. (1966). Relation of cue to consequence in avoidance learning. *Psychonomic Science, 4*(3), 123-124.

Halterman, Rose "Testing Models of the Flanker Preview Effect" master's thesis, Pennsylvania State University.

Kahneman, D., & Chajczyk, D. (1983). Tests of the automaticity of reading: Dilution of Stroop effects by color-irrelevant stimuli. *Journal of Experimental Psychology: Human Perception and Performance, 9*(4), 497 – 509.

Kamin, L. J. (1968). "Attention-like" processes in classical conditioning. In M. R. Jones (Ed.), *Miami Symposium on the Prediction of Behavior 1967: Aversive Stimulation* (pp. 9 -31). Coral Gables, FL: University of Miami Press.

Kerka, S. (2000). Incidental learning. *Trends and Issues, 18,* 3 – 4.

Kornblum, S. (1992). Dimensional overlap and dimensional relevance in stimulus-response and stimulus-stimulus compatibility. In G. E. Stelmach & J. Requin (Eds.), *Tutorials in motor behavior H,* (pp. 743-777). Amsterdam, Netherlands: Elsevier Science Publishers.

Kruschke, J. K. (2003). Attention in learning. *Current Directions in Psychological Science, 12*(5), 171 - 175.

Kruschke, J. K. (2006, June). Learned Attention. Paper presented at the International Conference on Learning and Development, Bloomington, IN.

Kruschke, J. K. (in press). Models of attentional learning. In E. M. Pothos & A.J. Wills (Eds.), *Formal approaches in categorization*. New York, NY: Cambridge University Press.

Kruschke, J. K., & Hullinger, R. A. (in press). Evolution of attention in learning. In: N. A. Schmajuk (Ed.), *Computational models of conditioning*. New York, NY: Cambridge University Press.

Levin, I. P., Wasserman, E. A., & Kao, S.-F. (1993). Multiple methods for examining biased information use in contingency judgments. *Organizational Behavior and Human Decision Processes, 55,* 228 – 250.

Lindsay, D. S., & Jacoby, L. L. (1994). Stroop process dissociations: The relationship between facilitation and interference. *Journal of Experimental Psychology: Human Perception and Performance, 20*(2), 219- 234.

Logan, G. D., & Zbrodoff, N. J. (1979). When it helps to be misled: Facilitative effects of increasing the frequency of conflicting stimuli in a Stroop-like task. *Memory & Cognition, 7*(3), 166 – 174.

Mackintosh, N. J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review, 82*, 276 – 298.

Mehta, R., & Russell, E. (2009). Effects of pre-training on acquisition of novel configural discriminations in human predictive learning. *Learning and Behavior, 37*(4), 311 – 324.

Miller, J. (1987). Priming is not necessary for selective-attention failures: Semantic effects of unattended, unprimed letters. *Perception & Psychophysics, 41*(5), 419-434.

Miller, J. O., & Pachella, R. G. (1973). Locus of the stimulus probability effect. *Journal of Experimental Psychology, 101*(2), 227 – 231.

Mordkoff, J. T. (1996). Selective attention and internal constraints: There is more to the flanker effect than biased contingencies. In A.F. Kramer, M.G.H. Coles, and G.D. Logan, (Eds.), *Converging operations in the study of visual selective attention* (pp. 483–502). Washington, D.C.: American Psychological Association.

Mordkoff, J. T. (1998). Between dimension flanker effects: A clarification with encouraging implications. *Psychonomic Bulletin & Review, 5,* 670-675.

Mordkoff, J. T., & Barth, H. (2001). Using pre-pulse inhibition to study attentional capture: A warning about pre-pulse correlations. In C. L. Folk & B. S. Gibson (Eds.), *Attraction, distraction, and action: Multiple perspectives on attentional capture. Advances in psychology*, (pp. 177-189). New York, NY: Elsevier.

Mordkoff, J. T. & Danek, R. H. (2009). What's case got to do with it?: Another difference between correlated and standard Flankers. Manuscript in preparation.

Mordkoff, J. T. & Danek, R. H. (2010). (in revision). Dividing attention between color and shape revisited: Evidence of coactivation when redundant targets are within the same object. *Perception & Psychophysics.*

Mordkoff, J. T. & Halterman, R. (2008). Feature integration without visual attention: Evidence from the correlated flankers task. *Psychonomic Bulletin & Review, 15* (2), 385 – 389.

Mordkoff, J. T., & Miller, J. (1993). Redundancy gains and coactivation with two different targets: The problem of target preferences and the effects of display frequency. *Perception & Psychophysics, 53*(5), 527-535.

Mordkoff, J. T., Miller, J., & Roch, A.-C. (1996). Absence of coactivation in the motor component: Evidence from psychophysiological measures of target detection. *Journal of Experimental Psychology: Human Perception and Performance, 22*(1), 25 - 41.

Mordkoff, J. T., & Yantis, S. (1991). An interactive race model of divided attention. *Journal of Experimental Psychology: Human Perception and Performance, 17*, 520-538.

Msetfi, R. M., Murphy, R. A., Simpson, J., & Kornbrot, D. E. (2005). Depressive realism and outcome density bias in contingency judgments: The effect of the context and intertrial interval. *Journal of Experimental Psychology: General, 134*(1), 10 – 22.

Neely, J. H. (1977). Semantic priming and retrieval from lexical memory: Roles of inhibitionless spreading activation and limited-capacity attention. *Journal of Experimental Psychology: General, 106*(3), 226 – 254.

Palef, S. R., & Olson, D. R. (1975). Spatial and verbal rivalry in a stroop-like task. *Canadian Journal of Psychology, 29*(3), 201 – 209.

Pashler, H., & Baylis, G. (1991). Procedural learning: 2. Intertrial repetition effects of speeded-choice tasks. *Journal of Experimental Psychology: Learning, Memory and Cognition, 17*(1), 33-48.

Pavlov, I. P. (1927). Conditioned Reflexes. London, England: Oxford University Press.

Pearce, J. M. (1987). A model for stimulus generalization in pavlovian conditioning. *Psychological Review, 94*(1), 61 – 73.

Pearce, J. M. (1994). Similarity and discrimination: A selective review and a connectionist model. *Psychological Review, 101*(4), 587 – 607.

Pearce, J. M., & Bouton, M. E. (2001). Theories of associative learning in animals. *Annual Review of Psychology, 52,* 111 – 139.

Perales, J. C., Catena, A., & Maldonado, A. (2004). Inferring non-observed correlations from causal scenarios: The role of causal knowledge. *Learning & Motivation, 35,* 115-135.

Pineño, O., Denniston, J. C., Beckers, T., Matute, H., & Miller, R. R. (2005). Contrasting predictive and causal values of predictors and of causes. *Learning & Behavior, 33*(2), 184 – 196.

Price, P. C., & Yates, J. F. (1993). Judgmental overshadowing: Further evidence of cue interaction in contingency judgment. *Memory & Cognition, 21*(5), 561 – 572.

Reber, A. S. (1989). Implicit learning and tacit knowledge. *Journal of Experimental Psychology: General, 118*(3), 219 – 235.

Rescorla, R. A. (1968). Probability of shock in the presence and absence of CS in fear conditioning. *Journal of Comparative and Physiological Psychology, 66*(1), 1-5.

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black and W. F. Prokasy, (Eds.), *Classical Conditioning II* (pp. 64-99). New York, NY: Appleton-Century-Crofts.

Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology, 8*(3), 382 - 439.

Rumelhart, D. E., & McClelland, J. L. (1986). On learning the past tenses of English verbs. In J. L. McClelland, D. E. Rumelhart and the PDP Research Group, (Eds.), *Parallel Distributed Processing, Vol. 2: Psychological and Biological Models* (pp. 216 – 271). Cambridge, MA: MIT Press.

Schwartz, B., Wasserman, E. A., & Robbins, S. J. (2002). *Psychology of learning and behavior (5th Ed.).* New York: Norton.

Sebe, N., Cohen, I., Cozman, F. B., Geveres, T., & Huang, T. S. (2005). Learning probabilistic classifiers for human-computer interaction applications. *Multimedia Systems, 00,* 1 – 15.

Shanks, D. R. (1985). Forward and backward blocking in human contingency judgment. *Quarterly Journal of Experimental Psychology, 37B,* 1 – 21.

Siegel, S. (2005). Drug tolerance, drug addiction, and drug anticipation. *Current Directions in Psychological Science, 14*(6), 296 – 300.

Siegel, S., & Allan, L.G. (1996). The widespread influence of the Rescorla-Wagner model. *Psychological Bulletin & Review, 3*(3), 314-321.

Siegel, S. & Ramos, B. M. C. (2002). Applying laboratory research: Drug anticipation and the treatment of drug addiction. *Experimental and Clinical Psychopharmacology, 10*(3), 162 – 183.

Smedslund, J. (1963). The concept of correlation in adults. *Scandinavian Journal of Psychology, 4,* 165 – 173.

Stroop, J. R. (1992). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology: General, 121*(1), 15- 23. (Reprinted from *Journal of Experimental Psychology: General, 18*, 643 - 662).

Thomas, G. & O'Callaghan, M. (1981). Pavlovian principles and behavior therapy. In G. Davey (Ed.), *Applications of Conditioning Theory* (pp. 129-148). London & New York: Methuen.

Thorndike, E. L. (1898). Animal Intelligence: An experimental study of the associative processes in animals. *Psychological Review, Monograph Supplements,(8):* entire issue.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science, 185*(4157), 1124 – 1131.

Ulrich, R., Mattes, S., & Miller, J. (1999). Donders's assumption of pure insertion: An evaluation on the basis of response dynamics. Acta Psychologica, 102, 43 - 75.

Van Hamme, L. J., Kao, S.-F., & Wasserman, E. A. (1993). Judging interevent relations: From cause to effect and from effect to cause. Memory & Cognition, 21, 802-808.

Verbruggen, F., & Logan, G. D. (2008). Automatic and controlled response inhibition: Associative learning in the go/no-go and stop-signal paradigms. Journal of Experimental Psychology: General, 137, 649-672.

Viken, R. J., Treat, T. A., Bloom, S. L., & McFall, R. M. (2005). Illusory correlation for body type and happiness: Covariation bias and its relationship to eating disorder symptoms . *International Journal of Eating Disorders, 38,* 65-72.

Wagner, A. R. (1981). SOP: A model of automatic memory processing in animal behavior. In N. E. Spear & R. R. Miller (Eds.), *Information Processing in Animals: Memory Mechanisms.* (pp. 5 - 47). Hillsdale, NJ: Lawrence Erlbaum Associates.

Waldmann, M. R. (2000). Competition among causes but not effects in predictive and diagnostic learning. *Journal of Experimental Psychology: Learning, Memory & Cognition, 26*(1), 53 – 76.

Waldmann, M. R. (2001). Predictive versus diagnostic causal learning: Evidence from an overshadowing paradigm. *Psychonomic Bulletin & Review, 8*(3), 600- 608.

Wasserman, E. A., Franklin, S. R., & Hearst, E. (1974). Pavlovian appetitive contingencies and approach vs. withdrawal to conditioned stimuli in pigeons. *Journal of Comparative and Physiological Psychology, 86*(4), 616 – 627.

Wasserman, E.A., Kao, S., Van Hamme, L., Katagiri, M., & Young, M.E. (1996). Causation and association. In D.R. Shanks, K.J. Holyoak, & D.L. Medin (Eds.), *The Psychology of Learning and Motivation, Vol. 34: Causal learning*. (pp. 207-264). San Diego, CA: Academic Press.

Welford, A. T., (1980). Choice reaction time: Basic concepts. In A. T. Welford (Ed.), *Reaction Times.* (pp. 73 - 128). New York, NY: Academic Press.

Williams, D. A., Sagness, K. E., & McPhee, J. E. (1994). Configural and elemental strategies in predictive learning. *Journal of Experimental Psychology: Learning, Memory & Cognition 20*(3), 694 – 709.