

Auditory Temporal Contextual Cueing

By

Lori Anne Doan

A Thesis submitted to the Faculty of Graduate Studies of

The University of Manitoba

In partial fulfillment of the requirements of the degree of

DOCTOR OF PHILOSOPHY

Department of Psychology

University of Manitoba

Winnipeg, MB

Abstract

When conducting a visual search task participants respond faster to targets embedded in a repeated array of visual distractors compared to targets embedded in a novel array, an effect referred to as contextual cueing. There are no reports of contextual cueing in audition, and generalizing this effect to the auditory domain would provide a new paradigm to investigate similarities, differences, and interactions in visual and auditory processing. In 4 experiments, participants identified a numerical target embedded in a sequence of alphabetic letter distractors. The training phase (Epochs 1, 2, and 3) of all experiments contained repeated sequences, and the testing phase (Epoch 4) contained novel sequences. Temporal contextual cueing was measured as slower response times in Epoch 4 than in Epoch 3. Repeated context was defined by the order of distractor identities and the rhythmic structure of the portion of the sequence immediately preceding the target digit, either together (Experiments 1 and 2) or separately (Experiments 3 and 4). An auditory temporal contextual cueing effect was obtained in Experiments 1, 2, and 4. This is the first report of an auditory temporal contextual cueing effect and, thus, it extends the contextual cueing effect to a new modality. This new experimental paradigm could be useful in furthering our understanding of fundamental auditory processes and could eventually be used to aid in diagnosing language deficits.

Acknowledgements

I would like to thank the members of my thesis advisory committee, Dr. Jason Leboe-McGowan, Dr. Todd Mondor, Dr. Murray Singer, Dr. Malcolm Smith, and Dr. Aimée Surprenant, for their support and guidance during my thesis development process and completion. In particular, I would like to thank my advisor, Dr. Jason Leboe-McGowan, for his invaluable mentorship over the past several years. I would also like to thank the Natural Sciences and Engineering Research Council of Canada, the University of Manitoba, and the Department of Psychology at the University of Manitoba for the financial assistance they provided to support my research.

Dedication

I dedicate this thesis to my family and close friends who stood beside me every moment of my doctorate program. In particular, I want to thank my sister, Kristal, and my dear friend, Louisa, for their steadfast support over the years. You allowed me to lean on you when I needed encouragement, and you each spent many hours in person and over the phone reminding me that I could finish this thesis, I must finish this thesis, and I would finish this thesis. I also would like to thank my mother, who never stopped believing in me despite the grave personal struggles she endured over the course of my undergraduate and graduate training.

I especially want to acknowledge my husband and daughter for their unwavering support during the entire course of my academic training. I want to thank Rick for all the times he shouldered the burdens at home to allow me to concentrate on my studies, and for the many pep talks he gave me when I inevitably encountered challenges and obstacles along the way. Tess, I started university when you were in kindergarten and I continued with my studies throughout your entire elementary and secondary school years. Thank you for understanding when I had to study and for supporting my journey to achieve my highest goal. Everything I did was ultimately for you, and I look forward to watching as you travel on your own journey.

Table of Contents

Abstract	ii
Acknowledgments.....	iii
Dedication.....	iv
Table of Contents.....	v
List of Tables	vi
List of Figures.....	vii
Chapter I: Introduction	1
Chapter II: Sequence Learning and Contextual Cueing.....	3
Chapter III: The Current Study.....	23
Experiment 1	23
Experiment 2.....	30
Experiment 3.....	33
Experiment 4.....	35
Chapter IV: General Discussion	39
References	48
Appendix: Consent Form	53

List of Tables

Table 1. Mean Response Times in milliseconds and Error Proportions (Error) in Experiment 1 (with the Between-Participants Standard Error of the Mean in Parentheses) 29

Table 2. Mean Response Times in milliseconds and Error Proportions (Error) in Experiment 2 (with the Between-Participants Standard Error of the Mean in Parentheses) 32

Table 3. Mean Response Times in milliseconds and Error Proportions (Error) in Experiment 3 (with the Between-Participants Standard Error of the Mean in Parentheses) 35

Table 4. Mean Response Times in milliseconds and Error Proportions (Error) in Experiment 1 (with the Between-Participants Standard Error of the Mean in Parentheses) 37

Table 5. Mean RTs in milliseconds by Epoch for all experiments (with Standard Error of the Mean in Parentheses). 38

List of Figures

Figure 1. A schematic of trials for Experiments 1 - 4 27

Chapter I: Introduction

Statistical regularities in the environment exist across modalities (e.g., Conway & Christiansen, 2005; Fiser & Aslin, 2001; Fiser & Aslin, 2002; Kirkham, Slemmer, & Johnson, 2002), and an organism that can implicitly learn these regularities as it goes about its daily business arguably has an evolutionary advantage over an organism that cannot learn them. Specifically, implicit memory of events or stimuli that tend to co-occur in the environment can improve the speed (e.g., Chun & Jiang, 1998; Nissen & Bullemer, 1987) and accuracy (e.g., Conway & Christiansen, 2005; Reber, 1967) with which one responds to those events or stimuli. For example, a poisonous berry might look very similar to an edible and nutritious berry, but perhaps it grows only in the presence of particular soil conditions, specific plants, etc., and the presence of these conditions alerts the organism that there is something “not quite right” about the berry, even if the organism cannot explicitly state the problem. Similarly, subtle changes in the environment may signal that a storm or a predator is very close. Therefore, an organism’s survival may be affected by the perception, interpretation, and response to the sensory information received from the environment, even when the organism is not consciously aware of that information.

The current study was motivated by research within the visual domain that has revealed that implicit memory for spatial context interacts with perception to orient attention during visual search (*contextual cueing*; Chun & Jiang, 1998). Although there have been multiple demonstrations of contextual cueing within the visual domain, to my knowledge there has only been one demonstration of a cross-modal contextual cueing

effect and no demonstrations of an auditory contextual cueing effect. Nabeta, Ono, and Kawahara (2003) demonstrated that repeated contextual information implicitly learned in the visual modality aided search in the haptic modality. Kawahara (2007) examined whether repeatedly associating a specific auditory cue with a specific visual target location within a visual search array would produce a contextual cueing effect. Although Kawahara demonstrated that an auditory cue could provide contextual information that would orient attention during visual search and aid participants in locating a target, he did not demonstrate auditory contextual cueing during auditory search. In the current study, I examined whether a temporal contextual cueing effect analogous to the effect observed in vision (Olson & Chun, 2001) would also occur in audition.

The theoretical distinction between sequential learning and contextual cueing is an important component of my research. Thus, I will present a review of the literature pertaining to sequence learning and the contextual cueing effect in Chapter II. I will describe my experiments and results in Chapter III, followed by a General Discussion in Chapter IV.

Chapter II: Sequence Learning and Contextual Cueing

I examined whether the temporal contextual cueing effect observed in vision could also be observed in the auditory domain. To accomplish this, I adapted Olson and Chun (2001) temporal contextual cueing experimental design to the auditory modality.

The typical spatial contextual cueing design involves searching for a target embedded in a spatial configuration of distractors. However, the temporal contextual cueing design involves asking participants to detect a target embedded in a sequence of distractors. Thus, the issue of whether temporal contextual cueing is simply a variant of sequence learning bears some discussion. This chapter will begin with a discussion of sequence learning, followed by a review of the contextual cueing literature.

Sequence learning

Serial reaction time (SRT) tasks. Two areas of debate in the sequence learning literature that are relevant to the proposed experiments are a) what is being learned about the sequences (i.e., sequence structure or motor responses) and, b) whether the learning that occurs is implicit or explicit (Riedel & Burton, 2006). The serial reaction time (SRT; Nissen & Bullemer, 1987) paradigm is widely used to investigate implicit sequence learning (Riedel & Burton, 2006). In the typical visual SRT paradigm, a stimulus appears in various locations on the computer screen, and participants press a key on each trial according to the stimulus location. Some sequences repeat throughout the experiment and some sequences are novel. Implicit sequence learning is evident when participants respond faster to the repeated sequences than to the novel sequences, but there has been some controversy as to whether participants are learning the sequence configuration (e.g., the sequence of stimulus locations) or the sequence of motor responses (Riedel & Burton,

2006). Some researchers have found that visual sequences can be learned by observation alone and do not require a motor response (e.g., Howard, Mutter, & Howard, 1992). However, others propose that sequence learning by observation is a different process because it represents explicit sequence learning, whereas motor response sequence learning is acquired implicitly; therefore, implicit sequence learning *is* motor sequence learning (e.g., Kelly, Burton, Riedel, & Lynch, 2003).

To test these premises, Riedel and Burton (2006) applied Mayr's (1996) dual sequence paradigm to the auditory domain, reasoning that this domain is optimal for investigating implicit sequence learning by observation alone because auditory information is generally sequential in nature, occurs along the temporal dimension, and often does not require a motor response. In the dual sequence SRT (Mayr, 1996), participants are presented with a sequence that contains two unrelated patterns. For example, Mayr presented participants with sequences of objects that contained a repeating pattern of object identities and a repeating pattern of object locations. Participants responded only to the object identity sequences, but Mayr's results revealed that they implicitly learned the irrelevant object location sequences. Mayr posited that implicit learning without motor responses might be unique to spatial sequence information, but Riedel and Burton suggested that responding to one aspect of a sequence might be enough to facilitate implicit learning of irrelevant sequential information and that this premise may extend beyond the spatial dimension. In particular, they were interested in using Mayr's paradigm to examine whether implicit sequence learning would occur in the temporal dimension without associated motor responses.

Riedel and Burton (2006) presented participants with a series of words that included two unrelated auditory sequences. Specifically, there were four colour words (blue, green, red, and white) that were spoken by four different voices. Thus, one auditory sequence was related to the identity of the words, one auditory sequence was related to the identity of the voice, and neither sequence type contained a spatial component. The word colour sequence and voice identity sequence were presented concurrently. Participants were instructed to respond to each element in a sequence by pressing the computer keys that corresponded to those elements. Participants were exposed to both sequences, but one group of participants responded via keypress to the word-colour identity sequence, and another group of participants responded via keypress to the voice identity sequence (Riedel & Burton, 2006). Participants were unaware of any repetitions in the sequences. After several practice trials, participants heard 16 blocks that each contained 72 trials. Thirteen blocks contained repeated voice and word identity sequences, and two blocks (blocks 9 and 12) contained one randomized sequence type (i.e., in the voice identity condition, the voice identity was randomized in block 9 and the word identity was randomized in block 12, and vice versa for the word identity condition). Both sequence types were randomized for all participants in block 15. Results revealed a decrease in RTs only for the factors that were responded to in the experiment. For example, participants in the voice identity sequence learning condition showed a decrease in RTs for sequences that repeated the voice identities compared to the random sequences, but no corresponding decrease in RTs for repeated word identities. Similarly, participants in the word identity sequence learning condition showed a decrease in RTs for sequences that repeated the word identities but no corresponding decrease in RTs for

repeated voice identities. In particular, Riedel and Burton's (2006) results revealed that word identity sequences were responded to faster than voice identity sequences. The authors noted that although these results indicate that the colour word sequence was easier to learn than the voice identity sequence, this did not provide any additional benefit in terms of implicit sequence learning. Specifically, participants in the voice identity condition did not implicitly learn the colour word sequences by simply listening to them (Riedel & Burton, 2006).

Riedel and Burton (2006) suggested that the motor responses required in their experiments might have interfered with learning the task irrelevant sequence information (i.e., the motor responses include a spatial component, whereas the voice and word identity tasks both contain temporal components). As the authors noted, their findings are in contrast with other studies that have demonstrated non-spatial sequence learning (e.g., Heuer, Schmidtke, & Kleinsorge, 2001; Koch & Hoffmann, 2000), but these studies differed from Riedel and Burton's methodologically, suggesting that there are many factors that influence whether implicit sequencing learning will occur. Riedel and Burton's findings underscore the complexity of the issues of implicit learning and motor response associations in sequence learning. Although they found that auditory sequence learning was dependent upon the responses made to the sequence, their results were not straightforward in terms of whether the learning was implicit or explicit (Riedel & Burton, 2006). They obtained some evidence in questionnaire and generation data that participants were aware of some patterns in the sequences they provided a response to, but the data for the sequences they did not respond to is more difficult to interpret (Riedel & Burton, 2006). Specifically, although the questionnaire data suggested that participants

in the word identity condition were aware of patterns in the voice identity sequences, the generation test suggested that this was not the case (Riedel & Burton, 2006). Overall, Riedel and Burton's results suggested that there must be some relation between the motor responses and the contents of the sequences for sequence learning to occur (Riedel & Burton, 2006).

Artificial grammar learning (AGL). Artificial grammar learning (AGL; Reber, 1967) is another experimental procedure that has been used to study implicit sequence learning. The AGL paradigm differs from the SRT paradigm in that it uses accuracy to measure learning, and it does not require a response to every item in the sequence. In the AGL paradigm, a set of rules is used to generate an artificial grammar in which some sequences are legal and some sequences are illegal (e.g., the letter E can follow S and Y, but it cannot follow X). Participants are exposed to legal sequences during training, but the rules of the artificial grammar are not made explicit to them. After training, participants view novel sequences and classify these sequences as legal or illegal. Participants can accurately classify the sequences even though they cannot explicitly state the rules of the artificial grammar, leading to the conclusion that implicit learning has occurred (Reber, 1967).

Conway and Christiansen (2005) used the AGL paradigm to investigate how auditory, visual, and tactile sequences are learned. To facilitate direct comparisons across the modalities, Conway and Christiansen presented participants with sensory stimuli appropriate for each modality and associated each sensory stimulus with a number (i.e., the numbers 1 to 5) that was used by the experimenters to create the AGL sequences. Thus, to measure tactile AGL, each number was associated with a finger (i.e., the thumb

was 1 and the little finger was 5), and participants felt a vibration in the finger according to the sequence that was created. For example, if the sequence was 2-5-5-1, the participant would feel vibrations in the order of: index finger, little finger, little finger, thumb. To measure visual AGL, participants were presented with a black box in one of five locations on the computer screen, with each location representing the numbers 1 to 5 (i.e., the first location represented 1, the last location represented 5). The auditory stimuli were five tones, with the lowest tone assigned the number 1 and the highest tone assigned the number 5. A rule set was created that dictated which numbers could and could not follow one another; sequences that followed these rules were considered *legal* and sequences that did not follow these rules were considered *illegal*. Participants were exposed to a series of legal sequences, but were not informed that there were any regularities in those sequences. After the training phase, participants were told that a rule set dictated the stimulus order in the training phase sequences. During the test phase, participants were presented with legal-illegal pairs of sequences and their task was to judge which sequence was created from the same rule set as the sequences from the training phase.

Conway and Christiansen's (2005) results revealed a number of interesting findings. Participants accurately discriminated between legal and illegal sequences in all modalities, but discrimination was significantly better for auditory sequences (Experiment 1). However, AGL occurred only in the auditory modality when sequence complexity increased (Experiment 2). In addition, discrimination was greatest for information at the end of auditory sequences, whereas discrimination was greatest for

information at the beginning of the tactile sequences, and there was no temporal order bias for visual sequences (Conway & Christiansen, 2005).

Statistical learning. AGL and SRT paradigms are often used to study implicit statistical learning, a phenomenon where individuals extract multifaceted regularities from the environment (Conway & Christiansen, 2005) and use those regularities to aid them in various modes of cognitive processing, including language acquisition (e.g., Saffran, Aslin, & Newport, 1996), visual processing (e.g., Fiser & Aslin, 2001; Fiser & Aslin 2002), and auditory processing (e.g., Saffran, Johnson, Aslin, & Newport, 1999). Statistical learning is regarded as implicit because the computation of probabilities required for statistical learning to occur transpires without conscious intention or awareness on the part of the learner (Conway & Christiansen, 2005). Implicit statistical learning has been observed in infants (Saffran et al., 1996), adults (Conway & Christiansen, 2005), and non-human primate species (Goujon & Fagot, 2013). Contextual cueing is another form of statistical learning (Chun & Jiang, 1998) and, therefore, the contextual cueing paradigm provides another means of examining the way statistical probabilities in the environment are learned.

Contextual Cueing

The typical contextual cueing paradigm (e.g., Chun & Jiang, 1998) differs from the SRT and AGL paradigms in that participants search for a target embedded in a display of distractors. Most contextual cueing experiments have been conducted in the visual domain along the spatial dimension, but contextual cueing was obtained in the haptic modality (Nabeta, Ono, & Kawahara, 2003) and a temporal contextual cueing effect has been reported using sequences of visual stimuli (Olson & Chun, 2001).

Contextual cueing, SRT, and AGL paradigms are similar in that they are all used to measure how the cognitive system extracts statistical regularities from the environment (Chun & Jiang, 1998; Conway & Christiansen, 2005), but the nature of these regularities differ. In the SRT and AGL paradigms, participants learn sequence order information and use that information in responding to the stimuli according to the task demands (i.e., detection in SRT or illegality judgments in AGL). Thus, the statistical regularities extracted in these paradigms are related to the immediate and specific relation between the sequence elements. In contrast, the statistical regularities extracted in the contextual cueing paradigm are based on the global context of the search displays (see Chun & Jiang, 1998). In addition, unlike the SRT paradigm, participants do not respond to every item in the display in the contextual cueing paradigm and there is no association between repeated contexts and motor responses (Chun & Jiang, 1998; Olson & Chun, 2001). The following section will describe the typical contextual cueing paradigm along with a description of the contextual cueing studies that are particularly relevant to the proposed experiments.

Spatial contextual cueing. Chun and Jiang (1998) were the first to report that an association between global context and target location aids visual search performance; an effect they called *contextual cueing*. They argued that people implicitly learn statistical regularities in the environment and form implicit memory representations of global contexts that interact with perception and attention in facilitating effective visual search. Chun and Jiang proposed that implicit learning of complex visual scenes results in implicit memory representations for global context. Importantly, their research suggested

an interaction between memory and perception in guiding attention during subsequent encounters with repeated visual contexts (Chun & Jiang, 1998).

Chun and Jiang (1998) developed the contextual cueing paradigm to test their premise and control for confounding variables, such as previous experience and knowledge and semantic associations between stimuli. In the typical contextual cueing experiment, participants are presented with visual search displays in which a target is embedded in a spatial configuration of distractors, and their task is to find the target as quickly and accurately as possible. Half of the visual search displays are presented only once in the experimental session (*novel* configurations) and half of the displays are presented repeatedly throughout the experimental session (*repeated* configurations). Accuracy is usually fairly high in these search tasks and, thus, the contextual cueing effect is most often defined as faster response times in repeated visual displays than in novel visual displays.

In their first experiment, Chun and Jiang (1998) presented their participants with visual search arrays that consisted of one visual target (i.e., a σ rotated to the left or the right) embedded in a spatial configuration of distractors (δ presented in orientations of 0, 90, 180, or 270 degrees). Participants viewed 720 visual search displays. Twelve configurations were repeated 30 times each throughout the experiment (*repeated* configurations), and the remaining 320 configurations were presented only once in the experiment (*novel* configurations). The distractor location and identities were maintained in repeated configurations, and the target always appeared in the same location in repeated configurations. However, the target identity was randomized in repeated configurations to ensure that the association learned was between the global distractor

context and the target location and was not based on motor responses. As predicted, participants located the target faster when it was embedded in a repeated configuration than when it was embedded in a novel configuration.

In a second experiment, Chun & Jiang (1998) demonstrated that participants were learning the global context of distractors and not the specific surface features of the search display (i.e., distractor identities). To accomplish this, they changed the distractor identities between the first half and the second half of the experiment, but maintained the global repeated configurations and the target identities (i.e., an upright σ or δ) throughout the experiment. All other aspects of this experiment were the same as the first experiment. Participants located the target faster in the repeated configurations than in the novel configurations throughout the experiment, despite the change in distractor identities, suggesting that participants were learning the association between the global context of distractors and the target location. In addition, participants performed at chance levels in a forced-choice recognition test that contained an equal number of repeated and novel configurations, indicating that their learning and memory of the repeated configurations were implicit.

A contextual cueing effect was not observed when the target location was varied in repeated configurations (Experiment 3), supporting the premise that the contextual cueing effect reflects associative learning, rather than expertise (Chun & Jiang, 1998). Varying the set size (i.e., 8, 12, or 16 items) produced a reduction in target slope that was significantly shallower for repeated configurations than for novel configurations (Experiment 4). Increasing set size produced a corresponding increase in search times. Therefore, a shallower target slope represents a more efficient search process for repeated

configurations, indicating that contextual cueing reflects an interaction between memory and attention that increases the efficiency of locating the target. If contextual cueing reflected an influence on early perceptual processes or later response selection processes, there would be a difference between the intercepts of response times for repeated and novel configurations across epochs. However, there was no significant difference in intercept times between repeated and novel configurations (Chun & Jiang, 1998). The difference in the mean response times for intercepts across all epochs did not differ between repeated and novel configurations, suggesting that the repeated configuration was not providing a benefit for perceptual processing of the search array.

Temporal contextual cueing. The typical contextual cueing paradigm involves visual search in which repeated spatial contexts aid visual search for targets embedded in a configuration of distractors. However, it has been demonstrated that temporal context also guides attention. Olson and Chun (2001) designed a *temporal contextual cueing* paradigm to examine whether temporal context would guide attention to stimuli presented in a series of distractors in a similar manner with which spatial context interacted with perception and attention in previous contextual cueing studies. In this paradigm a series of letters was presented rapidly and participants reported whether the sequence contained a ðkð or an ðxð. Temporal structure was defined and studied in three ways: the rhythmic structure of a sequence, the distractor and target identities within the sequence, and a spatiotemporal structure that included both spatial and temporal information.

To study whether the presentation duration of distractors would create a temporal context that would predict target location, Olson and Chun (2001) created rhythmic

patterns of visual stimuli in which each letter in a particular sequence was presented for varying durations (e.g., 80 ms, 666 ms, 266 ms, 1066 ms, etc.). In repeated sequences, the rhythmic pattern was repeated, but the specific target identities were randomized. The sequences contained 1 target and 15 distractors, but the repeated temporal context of letters varied from 3 to 10 letters in length, with the target presented immediately after the repeated stimuli. Thus, the target could be presented anywhere from the 4th to the 11th location in the sequence. Both targets appeared randomly in each repeated sequence, and the letters following the target appeared in random order in terms of letter identity and presentation duration. The training phase consisted of 12 blocks. In each block, there were eight repeated temporal contexts, and each repeated context was presented twice per block, once with the ök target and once with the öx target. Thus, the training phase of the experiment did not contain any novel sequences. The test phase of the experiment consisted of 4 blocks, and there were 16 novel (i.e., both distractor identities and presentation durations were randomized for each trial) sequences in each block. Olson and Chun found faster reaction times in the training phase of the experiment (i.e., for the repeated sequences), suggesting that memory for temporal context was interacting with perception and guiding participants' attention to the target's temporal location within the sequence. In a subsequent experiment, Olson and Chun randomized the order of distractors presented before a repeated sequence containing eight distractors and one target. They demonstrated that it was the order of the presentation durations (i.e., the rhythmic structure of the sequence) and not the absolute timing between the beginning of the sequence and the target presentation that was providing the critical contextual information that was guiding participants' attention to the target.

Olson and Chun (2001) used the same procedure to study the influence of temporal structure as defined by target identities. Thus, they randomized the presentation durations of the letters, but maintained the repeated sequence of distractor identities. Participants responded faster to repeated sequences presented in the training phase, suggesting that distractor identity also provides temporal contextual information that aids task performance. However, the results were not as strong as those of the first experiment, and Olson and Chun suggested that perhaps the association between distractor temporal position location and target identity was not as easily formed and, thus, not as useful in providing contextual information as an association between distractor identity and target identity or distractor location and target location. In other words, they suggested that some associations between stimuli are formed more easily than others and, thus, participants are more likely to learn this information implicitly and use it to aid them in their task performance.

To make their study more ecologically valid, Olson and Chun (2001) presented participants with repeated visual sequences that contained both spatially and temporally predictive information. Thus, the sequence of distractors consisted of frames that contained two distractors that were defined by both their identities (L and a cross) and their location on the frame. In other words, the $\tilde{L}\tilde{o}$ and the $\tilde{c}\tilde{r}\tilde{o}\tilde{s}\tilde{s}$ distractors varied in their spatial positions on each frame. The target frames contained the target and one of the distractors. In this experiment, the rhythmic structure of the sequence was held constant on all trials. Repeated sequences were defined as the repetition of distractor locations on the visual stimuli leading up to the target frame (Experiment 3B). In an additional experiment, the distractor frame immediately preceding the target frame was

the same on all trials (Experiment 3C). Olson and Chun demonstrated that participants implicitly learned the spatiotemporal information and this improved their performance, accuracy, and speed. They also showed that, although the frame presented immediately before the target appeared to be the most critical in providing predictive information (Experiment 3B), the distractor frames presented earlier in the sequence were also providing the participants with contextual cues (Experiment 3C).

Visual-haptic transfer of contextual cueing. As previously noted, contextual cueing research is conducted primarily in the visual domain. This prompted Nabeta et al. (2003) to examine whether this effect was exclusive to the visual modality by investigating whether visual contextual cueing would transfer to the haptic modality. Previous research suggests that visual/haptic cross-modal information transfer is used to form implicit (Easton, Greene, and Srinivas, 1997) and explicit (Shelton & McNamara, 2001) memory representations, and this formed the basis of Nabeta et al.'s hypothesis that implicit memory for the configuration of visual search arrays would transfer to haptic search arrays.

The training phase of the experiment was a typical contextual cueing experiment in which participants searched for a target (a rotated σ Tö) in an array of distractors (Lø rotated at 0, 90, 180, or 270 degrees). There were 320 trials in the training phase, 160 *novel* search configurations and 160 *repeated* search configurations. Participants indicated the orientation of the targets by pressing a left or right foot switch. In the transfer phase of the experiment, participants were presented with plastic sheets that had raised distractors and targets. A curtain was positioned between the participants and the plastic sheet to prevent the participants from viewing the spatial array of distractors.

Participants were instructed to use both hands to feel the raised items on the plastic sheet and to indicate the orientation of the target with the appropriate right or left foot switch. They conducted 32 haptic searches, with 16 *repeated* search arrays from the visual training phase of the experiment and 16 *novel* search arrays. A contextual cueing effect was observed in the training phase of the experiment and this effect transferred to the haptic modality in that participants were faster at locating the targets embedded in the *repeated* visual-haptic search arrays compared to the *novel* haptic search arrays that were presented only in the transfer phase of the experiment. They obtained the same pattern of results in a second experiment, in which the only difference was that the training phase was reduced to 192 trials. In addition, a recognition test conducted at the end of the first experiment indicated that participants were not explicitly aware of the *repeated* search arrays.

This experiment was the first experiment to demonstrate that implicit memory for global context in the visual modality can guide attention in another modality (Nabeta et al., 2003). In addition, the authors note that cross-modal contextual cueing effects provide evidence that contextual cueing is not a form of perceptual learning because perceptual learning does not generalize across conditions and stimuli (Nabeta et al., 2003).

Auditory-visual contextual cueing effect. To my knowledge, only one study has examined how the auditory and visual modalities interact to form memories for context that influence perception and guide attention during a visual search task. Kawahara (2007) paired an auditory stimulus with a specific target location within a visual search array. He used the same procedure as the one employed by Olson and Chun (2001) to study the temporal contextual cueing effect in that the first five blocks

contained repeated visual-auditory associations and the last block contained novel repeated visual-auditory associations. Kawahara's experiments consisted of a training phase and a testing phase. In the training phase, participants were presented with a target (rotated $\delta T\delta$) embedded in a configuration of distractors (rotated $L\delta$), and their task was to indicate if the $\delta T\delta$ was pointing to the right or to the left. Unlike previous visual contextual cueing experiments, the distractors were located randomly in each configuration, and Kawahara created an auditory-visual context by pairing a specific auditory sound with each target location. The auditory stimulus was a recording of a female voice reading Japanese, but played backwards and divided into sound files so that the voice sounded speech-like but the words did not contain any semantic information.

Each trial began with the auditory stimulus that predicted the target location, followed 2 seconds later by the visual search array. Participants were told to ignore the auditory stimulus and respond to the target $\delta T\delta$ as quickly and accurately as possible. Participants completed 400 visual searches in the training phase, all of which were preceded by an auditory stimulus that predicted the target location. In the testing phase, participants were presented with 80 visual search displays. The auditory stimulus-target location pairing was maintained for one group of participants (the "*consistent transfer*" group), but for another group of participants (the "*inconsistent transfer*" group) this pairing was disrupted so that the auditory stimulus no longer predicted the correct target location. A contextual cueing effect was observed in the training phase for both groups, with participants responding faster to the target over the course of the experiment. Most importantly, search performance continued to improve for the final 80 trials for the group in which the auditory stimulus continued to predict the target location. However, search

for the target was slower for the final 80 trials for the group in which the auditory stimulus no longer predicted the target location. Thus, a contextual cueing effect was observed in visual search arrays in which the configuration of distractors was randomized, but an auditory stimulus predicted the location of the target in the display.

Kawahara (2007) also found that participants performed above-chance in a recognition test performed at the end of the first experiment, suggesting that participants explicitly recognized the auditory stimuli. However, when participants were asked to indicate the target location predicted by the specific auditory stimuli, they were unable to do so, suggesting that the auditory stimuli-target location pairings continued to be inaccessible to conscious awareness. In other words, although participants were able to explicitly recognize the auditory stimuli in a post-experiment recognition test, they remained unaware of the auditory-visual pairings and of the predictive nature of the auditory stimuli. Thus, implicit memory for consistent auditory-visual pairings improved search efficiency (Kawahara, 2007).

Kawahara (2007) examined the boundary conditions of the auditory-visual contextual cueing effect. In the earlier experiment, participants recognized the auditory stimuli at above-chance levels, which was a finding contrary to the chance-level recognition typically demonstrated in visual contextual cueing experiments. Thus, Kawahara (2007) concentrated on this difference to ascertain whether auditory-visual contextual cueing would occur without explicit recognition of the auditory stimuli. To reduce participants' ability to recognize the auditory stimuli, Kawahara used tones from a touch-tone telephone, as they are less distinctive than the human voice, but each tone is distinct from the other tones used in the experiment (Kawahara, 2007). The telephone

tones were associated with specific target locations. A contextual cueing effect was not observed, and a recognition test completed immediately after the experiment revealed that participants were unable to explicitly recognize the auditory stimuli. Kawahara interpreted this as a boundary condition of the auditory-visual contextual cueing effect and argued that explicit recognition of the auditory stimuli is necessary to produce a contextual cueing effect with the procedure employed in his experiment. Kawahara proposed that auditory-visual contextual cueing is more difficult to obtain than visual contextual cueing and, thus, requires the additional support of explicit recall of the auditory stimuli to reduce the statistical noise and facilitate sensitivity to the statistical regularities (i.e., the predictive audio-visual pairing) in the search conditions. When the additional support is not available, a contextual cueing effect does not occur.

Kawahara's (2007) auditory-visual contextual effect is consistent with results obtained in other auditory-visual cross-modal studies. In particular, Mondor and Amirault (1998) demonstrated that an auditory cue that consistently predicts the location of a single visual target produces faster RTs to that target compared to a non-predictive auditory cue or a non-predictive visual cue. However, Kawahara's experiment extended these findings to the contextual cueing domain by demonstrating that an auditory cue can be incorporated into the memory representation of a visual search display and can subsequently enhance perception of, and responses to, the visual target embedded in the distractor configuration upon subsequent encounters with that search display.

Contextual cueing, SRT, and AGL paradigms reveal how the cognitive system extracts statistical regularities from the environment (Chun & Jiang, 1998; Conway & Christiansen, 2005). Researchers using SRT and AGL procedures have demonstrated that

people can extract information pertaining to the immediate association between sequence components, and to the associations between repeated contexts and motor responses (e.g., Riedel & Burton, 2006; Conway & Christiansen, 2005). Chun and Jiang (1998) utilized contextual cueing procedures to demonstrate that people could also extract the statistical regularities in a global context. In all three paradigms, statistical regularities facilitate response times and accuracy, and participants are often unaware of having learned these statistical regularities. Many studies examine exposure to visual regularities, but there is evidence that learning visual regularities can transfer to other modalities, such as Nabeta et al. (2003) findings that the regularities extracted from a visual search array facilitate performance when that array is transferred to the haptic domain (Nabeta et al., 2003). There is also evidence that people extract statistical regularities in the auditory modality. Conway and Christiansen (2005) revealed that discrimination between legal and illegal sequences in an AGL paradigm was significantly better for complex auditory sequences than for similar visual or tactile sequences. Discrimination was also greatest for information at the end of auditory sequences, whereas it was greatest for information at the beginning of the tactile sequences, and there was no temporal order bias for visual sequences. Predictive auditory cues also facilitate performance in visual detection and search tasks (Mondor and Amirault, 1998; Kawahara, 2007). Although the aforementioned studies provide evidence to support the premise that people extract statistical regularities in the auditory modality, to my knowledge, there are no previous studies that have examined whether exposure to temporal regularities facilitate detection of a target embedded within a series of distractors in a global auditory context. Thus, I examine this question in the current study. Chun and Jiang (1998) demonstrated that

attention and perception interact with memory in a visual search task. They proposed that global memory for context is an instance-based memory representation that, when matched to newly received perceptual information, guides attention in detecting a target embedded in a visual array of distractors. In the current study, I demonstrate that memory for context interacts with perception and attention in the auditory domain. I also provide a new experimental method to examine the cross-modal transfer and integration of global contextual information between audition and vision.

Chapter III: The Current Study

As previously noted, there is only one study reported in the literature that has used auditory stimuli in a contextual cueing paradigm (Kawahara, 2007), and it involved an auditory stimulus that predicted a target location in an array of randomly located visual distractors. The current study differs significantly from Kawahara's. Kawahara demonstrated that an auditory stimulus can cue the location of a target in a visual search display, whereas the intent of the current experiments is to demonstrate a temporal contextual cueing effect that is analogous to the temporal contextual cueing effect observed in the visual domain (Olson & Chun, 2001). In the following experiments, the training phase consisted of 12 blocks of trials containing auditory sequences with invariant information. A portion of each sequence in the training phase contained repetitions in the order of distractor identities and the rhythmic structure of the sequence (Experiments 1 and 2), rhythmic structure alone (Experiment 3), or order of distractor identities alone (Experiment 4). In all experiments, the testing phase consisted of four blocks of trials containing auditory sequences with variant information, in which neither the order of distractor identities nor the rhythmic structure of the sequences contained systematic repetitions. An auditory temporal contextual cueing effect is defined as significantly slower response times in the test phase of each experiment.

Experiment 1

Method

Participants

Twenty-two participants were recruited from the Introductory Psychology Participant Pool at The University of Manitoba. Participants were between 17 and 30

years of age, and had self-reported normal hearing. All participants received partial course credit for their participation.

Design

Participants searched for a target digit (2 or 9) presented within an auditory sequence of letters for all trials in the practice, training, and testing phases of the experiment. The experiment began with 6 practice trials, followed by a training phase that consisted of 12 invariant sequence blocks, and concluding with a test phase that consisted of four novel variant sequence blocks. Four invariant sequences were repeated across all training phase blocks. Each invariant sequence was presented twice per block, once with a 2 target and once with a 9 target, for a total of 8 trials in each training block. The sequences were presented in random order in each block. The test phase consisted of four blocks containing 8 trials each, with each target digit presented four times in each block. All sequences in the test phase trials were novel sequences that did not contain any invariant information. The invariant sequences were generated from new random permutation of the ISIs and the letters for each participant.

Apparatus and Stimuli

Computer and sound system. The experiment was conducted using a Dell computer connected to a 17-inch colour monitor. The E-Prime software system (Psychology Software Tools, Inc., 2003) was used to present stimuli and record responses, and sounds were presented using headphones.

Training phase. The distractor stimuli were recordings of alphabetic letters read in a male voice. All distractor and target stimuli were presented for 300 ms, but the intersimulus intervals (ISIs) were filled with buzzes that were presented for durations of

5, 105, 205, 305, or 405 ms. The total duration of each sequence varied from approximately 7 to 8 seconds each. The alphabetic letters were distinctive from each other in that no letter in the sequence rhymed with any other letter in the sequence. Thus, the following alphabetic letters were used as distractors: *A, B, F, I, L, M, O,* and *R*. The targets were the numbers 2 and 9. The õbuzzö sound was constructed using a square-wave and included a fundamental frequency of 200 Hz plus the first, second, third, and fourth harmonics (400 Hz, 600 Hz, 800 Hz, and 1000 Hz, respectively).

Following Olson and Chunø (2001) vision-based temporal contextual cueing procedure, only a portion of each 15 item sequence contained the invariant sequence, with eight invariant sequences that were three to ten letters in length; each sequence began with the invariant sequence and continued to the target item. Thus, the serial position of the target item varied from the 4th to 11th item in the sequence. The serial positions of the target were counterbalanced, with one group hearing sequences with the targets in the 4th to 7th positions, and one group hearing sequences with the targets in the 8th to 11th positions. The ISIs and the identities of the distractors that followed the target were randomized. The invariant temporal structure of the sequences was defined both by the rhythmic structure (i.e., sequence of inter-distractor intervals) of the alphanumeric sequence and by the distractor identities.

Testing phase. For the sequences in the four variant sequence blocks, the distractors and ISIs were completely randomized and, thus, those sequences did not systematically overlap with invariant sequences, either in the order of distractors or in the rhythmic structure of distractor presentations. Within sequences, each target digit was presented with equal frequency within each block.

Procedure

The training phase began with 6 practice trials, followed by 12 invariant blocks and 4 variant blocks. There was no overt division between the training and the testing phases of the experiment. Sound sequences were presented to participants via headphones.

Instructions were presented on the computer screen, and the research assistant also read the instructions aloud to all participants. Participants were instructed that they would hear a series of alphanumeric sound sequences and their task was to respond as quickly and accurately as possible when they heard the target digit by pressing the associated number key on the number keys above the alphabet keyboard. Participants were told that they should not wait until the end of the sequence to make their response, but should make their response as quickly as possible as soon as they heard the target digit. Participants pressed the space bar to begin each trial. A fixation cross appeared on the screen and remained there for the duration of the trial. The alphanumeric sound sequence began 1000 ms after the initial appearance of the fixation cross. Participants responded by pressing the appropriate target key. For each trial, RTs were calculated from the onset of the target digit until the participant pressed either the 2 or 9 key. A prompt reading *“Press the space bar to begin the next trial”* appeared on the screen 1000 ms after the completion of the alphanumeric sound sequence. Participants were not informed that any aspect of the sound sequences was being repeated during the experiment, and there were no overt breaks between the blocks to cue participants that they were beginning a new experimental block. Participants received feedback on their

accuracy on all practice and experimental phase trials. A schematic of the trials used for all experiments is displayed in Figure 1.

At the end of the experiment, participants were asked to report anything that they may have noticed about the experimental sequences.

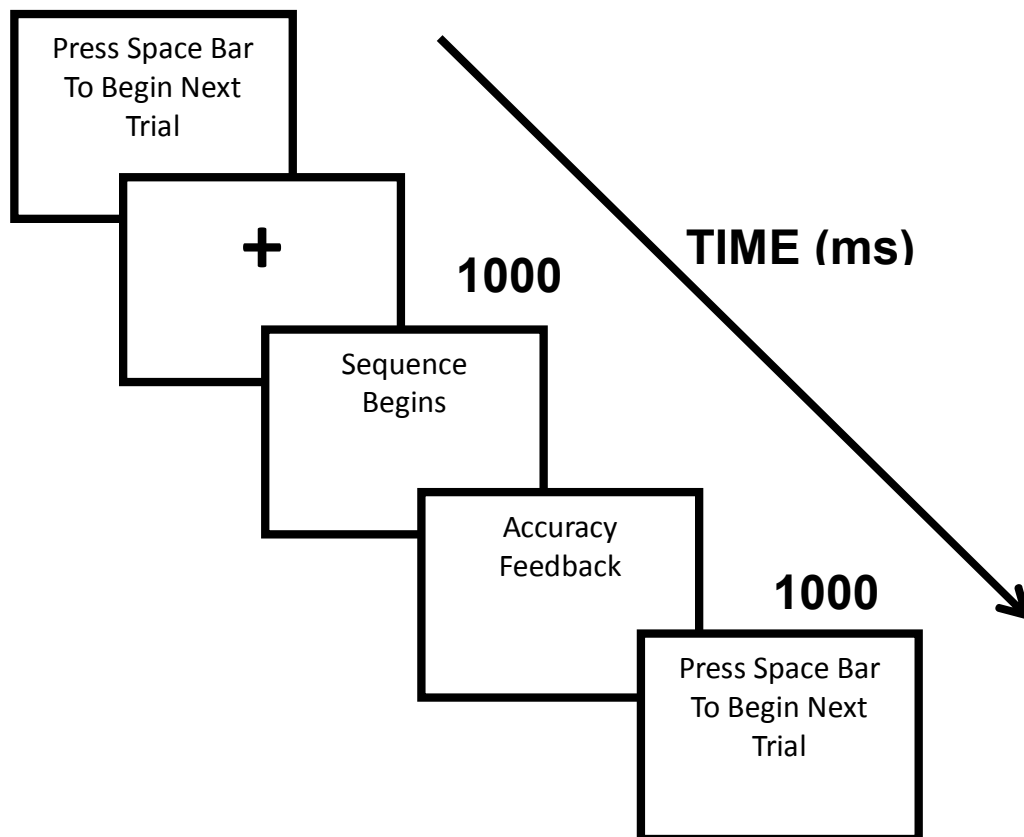


Figure 1. A schematic of trials for Experiments 1 to 4.

Results

To increase statistical power, the experimental blocks were grouped into larger units referred to as *epochs*; each epoch contained four blocks. Thus, the first three epochs consisted of the invariant sequences and the last epoch contained the variant (i.e., novel)

sequences. The means of the response times (RTs) for each epoch and condition were calculated on correct response trials only. RTs that indicated that a participant responded before the target digit was presented were not included in the analysis (less than 1% of the trials), and three participants whose RTs indicated that they had waited until the end of the sequences to respond were excluded from the analysis. Participants' mean response times for each epoch were submitted to a repeated measures Analysis of Variance (ANOVA), treating epoch (Epoch 3 vs. Epoch 4) as a within-participant factor. Auditory contextual cueing was measured as the difference in RTs between the third and fourth epochs.

Mean accurate response times were computed for each participant and were submitted to a repeated measures Analysis of Variance (ANOVA), treating Epoch (*Epoch 3 vs. Epoch 4*) as a within-participant factor. This analysis revealed that participants' response times were 31 ms slower in *Epoch 4* than in *Epoch 3*, $F(1, 21) = 4.731$, $MSe = 2274.143$, $p < .05$. The mean RTs for the first three Epochs (i.e., the training phase) were submitted to a repeated measures ANOVA. This analysis revealed a nonsignificant effect of Epoch, $F(2, 42) = .692$, $MSe = 2116.226$, $p > .05$.

Error effects will be reported for all four experiments, but these effects were small and inconsistent across experiments and do not provide a reasonable basis of interpretation for the results obtained for the RTs. Specifically, the error effects do not suggest that there was a speed-accuracy trade-off in any of the experiments. In Experiment 1, mean errors were computed for each participant and were submitted to a repeated measures Analysis of Variance (ANOVA), treating Epoch (*Epoch 3 vs. Epoch 4*) as a within-participant factor. There was no difference in error rates between *Epoch 3*

and *Epoch 4*, $F(1, 21) = 2.392$, $MSe = .000$, $p > .05$. The mean errors for the first three Epochs (i.e., the training phase) were submitted to a repeated measures ANOVA. This analysis also revealed a nonsignificant effect of Epoch, $F(2, 42) = 1.657$, $MSe = .000$, $p > .05$. Response times, error proportions, and between-participants standard error of the mean are displayed in Table 1.

Table 1

Mean Response Times in milliseconds and Error Proportions (Error) in Experiment 1 (with the Between-Participants Standard Error of the Mean in Parentheses)

Epoch	Response Time	Error
1	576 (28)	.008 (.004)
2	560 (22)	.013 (.004)
3	572 (24)	.018 (.004)
4	603 (23)	.009 (.004)

Note. Response times and error proportions are given as a function of epoch. The last epoch is the testing phase with variant sequences. The first three epochs are the training phases with invariant sequences. The between-participant standard error of mean RTs and error rates are provided in parentheses.

To assess whether outliers were influencing the results, response times less than 250 ms and more than 2000 ms were eliminated for all experiments, resulting in an elimination of 1% of all trials. Median analyses conducted on all experiments revealed the same pattern of results across experiments as those obtained with the reported mean analyses and, thus, this will not be discussed further in subsequent experiments.

These results are consistent with an auditory temporal contextual cueing effect. Participants responded slower to the target in Epoch 4, which consisted of novel variant sequences, suggesting that participants acquired knowledge of the invariant distractor identity and rhythmic structure preceding the target digit in the first three Epochs, and the

absence of this information during Epoch 4 impaired their auditory search performance. At the conclusion of the testing phase, all participants were asked to report anything they noticed about the experimental sequences. None of the participants reported that there were repeated sequences in the experiment, despite the fact that there were only four repeated sequences (presented an equal number of times with each target digit) across the first 96 experimental trials during the training phase of the experiment. Participants' unanimous failure to express awareness of any regularity in the presentation of the sound sequences occurred across all four of my experiments.

Although the results of Experiment 1 were consistent with an auditory temporal contextual cueing effect analogous to the visual temporal contextual cueing effect reported by Olson and Chun (2001), there is an alternative explanation for these results. It is possible that participants were using information regarding the absolute time from the start of the sequence to the target to guide their auditory search. Experiment 2 was designed to investigate this possibility.

Experiment 2

Method

Participants

Twenty-two participants were recruited from the Introductory Psychology Participant Pool at The University of Manitoba. Participants were between 17 and 30 years of age, and had self-reported normal hearing. All participants received partial course credit for their participation.

Design

The design was the same as Experiment 1.

Apparatus and Stimuli

Computer and sound system. The computer, software, and sound system were the same as in Experiment 1.

The stimuli were the same as in Experiment 1 except that the repeated portion of the sequences were preceded by a randomized order of distractor identities and ISIs. In addition, invariant sequences contained five to eight distractors, with targets in the 6th, 7th, 8th, or 9th positions in the sequences.

Procedure

The procedure for the training, testing, and recognition phases of the experiment was the same as Experiment 1.

Results

Mean accurate response times were computed for each participant and were submitted to a repeated measures ANOVA, treating Epoch (*Epoch 3* vs. *Epoch 4*) as a within-participant factor. This analysis revealed that participants' response times were 54 ms slower in *Epoch 4* than in *Epoch 3*, $F(1, 21) = 5.465$, $MSe = 5897.906$, $p < .05$. The mean RTs for the first three Epochs (i.e., the training phase) were submitted to a repeated measures ANOVA. This analysis revealed a nonsignificant effect of Epoch, $F(2, 42) = 2.199$, $MSe = 1809.723$, $p > .05$.

Mean errors were computed for each participant and were submitted to a repeated measures Analysis of Variance (ANOVA), treating Epoch (*Epoch 3* vs. *Epoch 4*) as a within-participant factor. There was no difference in error rates between *Epoch 3* and *Epoch 4*, $F(1, 21) = .000$, $MSe = .001$, $p > .05$. The mean errors for the first three Epochs

(i.e., the training phase) were submitted to a repeated measures ANOVA. There was a significant effect of Epoch, $F(2, 42) = 4.188$, $MSe = .000$, $p < .05$. Post-hoc tests using the Bonferroni correction revealed a trend toward more errors in *Epoch 1* than in *Epoch 2* ($p = .072$) and *Epoch 3* ($p = .088$), but neither of these comparisons were statistically significant. *Epoch 1* contained .019 higher proportion of errors than *Epoch 2*, and .014 higher proportion of errors than *Epoch 3*. Response times, error proportions, and between-participants standard error of the mean are displayed in Table 2.

Table 2

Mean Response Times in milliseconds and Error Proportions (Error) in Experiment 2 (with the Between-Participants Standard Error of the Mean in Parentheses)

Epoch	Response Time	Error
1	539 (15)	.036 (.008)
2	524 (19)	.017 (.005)
3	551 (18)	.021 (.006)
4	605 (30)	.021 (.005)

Note. Response times and error proportions are given as a function of epoch. The last epoch is the testing phase with variant sequences. The first three epochs are the training phases with invariant sequences. The between-participant standard error of mean RTs and error rates are provided in parentheses.

The results obtained for RTs indicate that participants did not rely on the absolute time from the onset of the sequence to target onset to guide them in their auditory search. This suggests that participants acquired knowledge of the invariant sequence of distractor identities and rhythmic structure preceding the target digit in the first three Epochs, and the absence of this information during Epoch 4 impaired their auditory search performance. Therefore, the results of Experiment 2 are consistent with an auditory temporal contextual cueing effect.

The results of Experiments 1 and 2 revealed an auditory temporal contextual cueing effect, with context defined as the invariant sequence of distractor identities and the invariant rhythmic structure within the portion of the auditory sequences leading up to the target digit. In both experiments, participants responded slower to the target in Epoch 4 than in Epoch 3. Experiments 3 and 4 will investigate if repetitions are required in both distractor identities and rhythmic structures within the auditory sequences to produce an auditory temporal contextual cueing effect, or if either the distractor identities or the rhythmic structures alone are enough to produce the effect.

Experiment 3

Experiment 3 tested the premise that invariant auditory temporal structure, as defined by the rhythmic structure of an auditory sequence, will produce an auditory temporal contextual cueing effect.

Method

Participants

Twenty-four participants were recruited from the Introductory Psychology Participant Pool at The University of Manitoba. Participants were between 17 and 30 years of age, and had self-reported normal hearing. All participants received partial course credit for their participation.

Design

The design was the same as Experiment 1.

Apparatus and Stimuli

Computer and sound system. Computer equipment and sound system were the same as Experiment 1.

The stimuli were the same as Experiment 1 except that the invariant sequences consisted of repetitions in the rhythmic structure of the sequences (i.e., the pattern of ISIs), and all distractor identities were randomized within the sequences.

Procedure

Experimental phases. Training and test phases of Experiment 3 were identical to Experiment 1.

Results

Mean accurate response times were computed for each participant and were submitted to a repeated measures ANOVA, treating Epoch (*Epoch 3* vs. *Epoch 4*) as a within-participant factor. This analysis revealed that participants response times were not significantly slower in *Epoch 4* than in *Epoch 3*, $F(1, 23) = 2.339$, $MSe = 5437.215$, $p > .05$. The mean RTs for the first three Epochs (i.e., the training phase) were submitted to a repeated measures ANOVA. This analysis revealed a nonsignificant effect of Epoch, $F(2, 46) = .914$, $MSe = 4821.256$, $p > .05$.

Mean errors were computed for each participant and were submitted to a repeated measures Analysis of Variance (ANOVA), treating Epoch (*Epoch 3* vs. *Epoch 4*) as a within-participant factor. There was a significant difference in error rates between *Epoch 3* and *Epoch 4*, $F(1, 21) = 4.285$, $MSe = .000$, $p = .05$. *Epoch 4* contained .009 higher proportion of errors than *Epoch 3*. The mean errors for the first three Epochs (i.e., the training phase) were submitted to a repeated measures ANOVA. This analysis revealed a nonsignificant effect of Epoch, $F(2, 42) = 1.512$, $MSe = .000$, $p > .05$. Response times, error proportions, and between-participants standard error of the mean are displayed in Table 3.

Table 3

Mean Response Times in milliseconds and Error Proportions (Error) in Experiment 3 (with the Between-Participants Standard Error of the Mean in Parentheses)

Epoch	Response Time	Error
1	560 (22)	.021 (.004)
2	577 (27)	.016 (.004)
3	581 (32)	.012 (.004)
4	614 (45)	.021 (.004)

Note. Response times and error proportions are given as a function of epoch. The last epoch is the testing phase with variant sequences. The first three epochs are the training phases with invariant sequences. The between-participant standard error of mean RTs and error rates are provided in parentheses.

These results are not consistent with an auditory temporal contextual cueing effect, suggesting that invariant rhythmic structure alone is insufficient to produce the effect.

Experiment 4

Experiment 4 tested the premise that invariant auditory temporal structure, as defined by the order of distractor identities in the sequence, will produce an auditory temporal contextual cueing effect. Only the distractor identities were repeated in the training phase and, thus, the predictive information in the current experiment did not include a temporal dimension.

Method

Participants

Twenty-seven participants were recruited from the Introductory Psychology Participant Pool at The University of Manitoba. Participants were between 17 and 30

years of age, and had self-reported normal hearing. All participants received partial course credit for their participation.

Design

The design was the same as Experiment 1.

Apparatus and Stimuli

Computer and sound system. The computer, software, and sound system were the same as the previous experiments.

The stimuli was the same as in Experiment 1 except that the invariant sequences used in the training phase were composed of a repeated sequence of distractor identities and the duration of the ISIs were randomized across all trials.

Procedure

The procedure for the training, testing, and recognition phases of the experiment was the same as the previous three experiments.

Results

Mean accurate response times were computed for each participant and were submitted to a repeated measures ANOVA, treating Epoch (*Epoch 3* vs. *Epoch 4*) as a within-participant factor. This analysis revealed that participants response times were 54 ms slower in *Epoch 4* than in *Epoch 3*, $F(1, 26) = 5.360$, $MSe = 7425.566$, $p < .05$. The mean RTs for the first three Epochs (i.e., the training phase) were submitted to a repeated measures ANOVA. This analysis revealed a nonsignificant effect of Epoch, $F(2, 52) = .533$, $MSe = 14527.773$, $p > .05$. Response times and the standard error of the mean for all four Epochs are displayed in Figure 4.

Mean errors were computed for each participant and were submitted to a repeated measures Analysis of Variance (ANOVA), treating Epoch (*Epoch 3* vs. *Epoch 4*) as a within-participant factor. There was no difference in error rates between *Epoch 4* and *Epoch 3*, $F(1, 26) = 1.638$, $MSe = .002$, $p > .05$. The mean errors for the first three Epochs (i.e., the training phase) were submitted to a repeated measures ANOVA. There was a significant effect of Epoch, $F(2, 52) = 5.416$, $MSe = .001$, $p < .05$. Post-hoc tests using the Bonferroni correction revealed a significant difference in errors between *Epoch 1* and *Epoch 3*, ($p < .05$), with *Epoch 1* containing .019 higher proportion of errors than *Epoch 3*. Response times, error proportions, and between-participants standard error of the mean are displayed in Table 4.

Table 4

Mean Response Times in milliseconds and Error Proportions (Error) in Experiment 4 (with the Between-Participants Standard Error of the Mean in Parentheses)

Epoch	Response Time	Error
1	636 (41)	.027 (.007)
2	670 (31)	.027 (.008)
3	657 (35)	.008 (.003)
4	711 (48)	.022 (.012)

Note. Response times and error proportions are given as a function of epoch. The last epoch is the testing phase with variant sequences. The first three epochs are the training phases with invariant sequences. The between-participant standard error of mean RTs and error rates are provided in parentheses.

Experiment 4 results revealed an auditory temporal contextual cueing effect with the temporal context defined as the invariant order of distractor identities. Participants were significantly slower in Epoch 4 than in Epoch 3. This suggests that distractor

identity alone is sufficient for participants to acquire knowledge of the invariant pattern in repeated sequences.

In summary, a contextual cueing effect was observed in Experiments 1, 2, and 4.

Table 5 displays the mean RTs and the standard error of the mean for all four Epochs and all experiments.

Table 5

Mean RTs in milliseconds by Epoch for all experiments (with the Between-Participants Standard Error of the Mean in Parentheses)

Experiment	Epoch			
	1	2	3	4
1*	576 (28)	560 (22)	572 (24)	603 (23)
2*	539 (15)	524 (19)	551 (18)	605 (30)
3	560 (22)	577 (27)	581 (32)	614 (45)
4*	636 (41)	670 (31)	657 (35)	711 (48)

Note. The last Epoch is always from the testing phase. The first three Epochs are always from the training phase. Experiments with an asterisk denote an auditory contextual cueing effect, as measured by the difference in response times between Epoch 3 and Epoch 4. The between-participant standard error of mean RTs are provided in parentheses.

CHAPTER IV: General Discussion

The present study introduces a new experimental paradigm to the contextual cueing literature. To my knowledge, this study is the first to reveal an auditory temporal contextual cueing effect that is analogous to the effect reported in the visual modality (Olson & Chun, 2001). Temporal contextual cueing occurred when the repeated context was defined as distractor identity and rhythmic structure (Experiment 1) and distractor identity alone (Experiment 4). Contextual cueing did not occur when the repeated context was defined by rhythmic structure alone (Experiment 3). Experiment 2 revealed that participants were extracting information regarding the global temporal context of the repeated portion of the sequence, and were not using the absolute time from sequence onset to the target. The current study did not test whether memory for the sequences was implicit or explicit. However, at the end of each experiment participants were asked if they had noticed anything specific about the sequences and all participants failed to identify the repetitions in the sequences.

It is important to note that while I found a contextual cueing effect in Experiments 1, 2, and 4, there was no evidence that the repetitions in the first three epochs were facilitating performance in any of those experiments. In all three experiments, there was no main effect of Epoch for the first three repeated sequence epochs, indicating that RTs did not differ between these epochs. If the repetitions were aiding participants in detecting the target, one would expect to find a gradual improvement in performance (i.e., decreasing RTs) across the first three epochs. Thus, it would appear that replacing the invariant sequences of the first three epochs with variant sequences in the last epoch impaired performance; the invariant sequences did not improve performance across the

epochs containing those sequences. This is, to some extent, consistent with the findings of Olson and Chun (2001), but there are a few important differences between their results and my results.

Olson and Chun (2001) observed a general learning effect across training phase epochs (i.e., invariant sequences) in Experiments 1A and 1B (repetition in the visual sequences was defined by the invariant rhythmic structure of the sequence). However, they did not find a general learning effect across training phase epochs in Experiment 2 (repetition in the sequences was defined by the invariant order of distractor identities). In Experiments 1A and 1B, the stimuli (i.e., the rhythmic structure of the sequences) matched the global context of the sequence along the temporal dimension. The facilitated performance in these experiments is consistent with the principles of encoding specificity (Tulving & Thompson, 1973) and transfer-appropriate processing (Morris, Bransford, & Franks, 1977), which propose that a match in context and processing between an initial encounter with a stimulus and subsequent encounters with the same stimulus benefits remembering. In Experiment 2, the invariant context of the sequences was defined by the distractor identities and, thus, the stimuli did not match the global context of the sequence along the temporal dimension. This mismatch may have reduced the size of the temporal contextual cueing effect in Experiment 2 (Olson & Chun, 2001).

Consistent with Olson and Chun's (2001) Experiment 2, I did not find a general learning effect in any of the experiments in which contextual cueing was observed (Experiments 1, 2, and 4). Analyses of the early blocks of the experiment revealed a similar lack of a general learning effect, eliminating the possibility that the learning effect was masked by grouping the blocks into epochs. Furthermore, in contrast to Olson and

Chun's Experiments 1A and 1B, I did not observe a contextual cueing effect in Experiment 3, where there was a match along the temporal dimension between the stimuli (i.e., the rhythmic structure of the sequence) and the global context of the sequence. Furthermore, I found a contextual cueing effect in Experiment 4, where the principles of match and the findings of Olson and Chun would predict an absence of contextual cueing, or a reduction in the magnitude of the effect, due to the mismatch between the stimuli (i.e., distractor identities) and the global context of the sequence. This could indicate a difference in visual and auditory processing of global temporal context. However, it is also possible that a contextual cueing effect was not observed in Experiment 3 due to a lack of power in that experiment.

Limitations and Future Directions

In Olson and Chun's (2001) visual displays, the rhythmic structure was defined by the duration of time that the distractors were presented on the screen. In the auditory modality, presenting alphanumeric stimuli for varying lengths of time produces an awkward and unrealistic sounding sequence and, therefore, I defined the invariant rhythmic structures in the current experiments by the ISIs in the sequence. The ISIs contained a buzz sound, as opposed to silence, to highlight the rhythmic structure of the sequence. However, this may not have provided enough information for participants to develop associations between the repeated rhythm of each sequence and the associated target position. As a result, the distractor identities may have been more salient than the rhythmic structure in my sequences and, thus, people were able to extract the statistical regularities of the invariant order of distractor identities (Experiment 4), but they were unable to do so with the invariant rhythmic structure of the sequences (Experiment 3).

Although Experiment 4 represents a mismatch in the temporal dimension between stimuli and context, it may be that the distractor identities were more distinctive (see Von Restorff, 1933) than the global rhythmic structure in an auditory sequence. If this was the case, the benefit of the distinctiveness of the distractor identities may have outweighed the disadvantage of the mismatch between the stimuli and the global context in Experiment 4. In addition, people typically listen with the goal of identifying what they are hearing and, therefore, they selectively attend to stimuli identities and ignore other irrelevant features of an auditory context. Prior experience of this nature may have resulted in participants attending to stimuli identities within the sequences and ignoring the rhythmic structure of the sequences in the current experiments. In particular, the rhythmic structure was in the ISIs, not the distractors, and this manipulation may have been too weak to allow participants to extract and learn the statistical regularities in the repeated sequences in Experiment 3. In the future, researchers should examine the roles of match, distinctiveness, and prior experience in the auditory temporal contextual cueing effect by using stimuli other than alphanumeric distractors and targets, such as tones and glides. Although the stimuli in the current study were inspired by the stimuli used in Olson and Chun's study, tones and glides are fundamentally less meaningful than alphanumeric stimuli and participants would rely less on prior experience with the stimuli when searching for a glide target in an auditory sequence of tones. As well, using tones as distractors would allow researchers to employ a stronger manipulation of the rhythmic structure of the sequences because the structure could be defined by the presentation duration of the tones instead of by the ISIs. Thus, using tones and glides as stimuli within an auditory temporal contextual cueing paradigm could be better suited for investigating

the boundaries of the auditory temporal contextual cueing effect. The stimuli in the current study would work well for experiments investigating cross-modal interactions between vision and audition.

There was strong reason to expect that an auditory temporal contextual cueing effect would be observed in the current study, as the auditory system is particularly adept at temporal processing (e.g., Rubinstein & Gruenberg, 1971; Saffran, 2002). The auditory modality appears to be particularly sensitive to statistical regularities in temporal sequences (Conway & Christiansen, 2005), and auditory cues can aid perception of, and responses to, visual target locations (Kawahara, 2007; Mondor & Amirault, 1998). However, there is also a strong motor response component to implicit sequence learning (Riedel & Burton, 2006), and it could be the association between the motor response and the auditory stimuli that is learned, and not the global context of the sequence. The motor responses were not associated with the repeated sequences in the current experiments and, therefore, the current study reveals that the auditory system is also sensitive to the global context of auditory sequences. However, although I found a contextual cueing effect in Experiments 1, 2, and 4, I did not find evidence that invariant sequences provide a benefit to detecting a target in an auditory sequence. This represents a null effect and, as such, it precludes definitive conclusions, but it does suggest possible differences in statistical learning constraints between the visual and the auditory modalities. Future research should examine this issue further to determine what conditions produce a benefit for invariant sequences in auditory search.

One of the most important contributions of this study is that of an auditory contextual cueing paradigm that researchers can use to explore the cross-modal exchange

of information between audition and vision. There has been a great deal of research into the cross-modal integration of perceptual information, and to the blended or emergent percepts this transfer creates (de Gelder & Bertelson, 2003). At its most fundamental level, the ability to receive and interpret information from multiple sources facilitates survival by supplementing perceptual gaps in one modality with sensory information from another modality (Shams & Seitz, 2008; Thompson et al., 2008; Van den Stock et al., 2008). Visual cues augment audition in the interpretation of emotion (e.g., de Gelder & Vroomen, 2000; Van den Stock et al., 2008) and auditory cues supplement vision in interpreting the location of a sound (e.g., McDonald, Teder-Sälejärvi, & Hillyard, 2000). Perhaps the most obvious interaction between vision and audition is in human communication.

Although the ability to communicate per se is not unique to the human species, some suggest that the ability to acquire complex and flexible language structures is a defining feature that sets humans apart from other species (see Hauser, Chomsky, & Fitch, 2002). Humans use language to communicate through sounds and through symbols, and cross-modal transfer and integration between audition and vision operates in both spoken and written communication. For example, research suggests that deficits in auditory processing are implicated in reading difficulties, such as dyslexia (e.g., Habib, Rey, Daffaure, Camps, Espesser, Joly-Pottuz, & Dermonet, 2002; Rey, De Martino, Espesser, & Habib, 2002; Tallal, Miller, & Fitch, 1993; but see Mody, 2003), and the influence of multimodal integration of perceptual information is not limited to the written language. Although it is obvious that perceiving spoken language requires auditory processing, the McGurk Effect (McGurk & MacDonald, 1976) and the ventriloquism

effect (Choe, 1975) demonstrate how visual cues may influence the perception and interpretation of spoken language. Cross-modal integration also influences the perception of emotion through interactions between facial expression and voice tone (de Gelder & Vroomen, 2000) and between body language and non-verbal auditory stimuli (Van den Stock et al., 2008). Multisensory integration contributes to various aspects of human communication, but the transfer of information between modalities is also an important component. For example, cross-modal transfer enhances speech perception when an initial visual context predicts a subsequent auditory event (Sánchez-García, Alsius, Enns, & Soto-Faraco, 2011). Thus, cross-modal integration and transfer of information between vision and audition is an important component of language comprehension, and the development of an auditory temporal contextual cueing paradigm provides a new means of investigating this aspect of language in future research.

Conclusion

In the current study, I found evidence of temporal contextual cueing in the auditory modality that is analogous to the temporal contextual cueing effect reported in the visual modality (Olson & Chun, 2001). My findings give rise to questions regarding the similarities and differences in temporal processing between the visual and auditory modalities, but they also lead to another question regarding the interaction between vision and audition when perceiving and remembering a global temporal context. It is plausible to suggest that the cognitive system is flexible in perceiving and interpreting information from multiple modalities and that the strategy that is most efficient for the task will be the one employed. If cross-modal *integration* is most efficient, such as when one is listening to a speaker while simultaneously lip-reading, then visual and auditory information will

be integrated as it is processed. Conversely, if cross-modal *transfer* is the most efficient strategy for the task, then instance-based memory traces from one modality will transfer to another modality. The auditory temporal contextual cueing experimental procedure presented in these studies provides a new means to investigate the mechanisms underlying the transfer and integration of auditory and visual information. Researchers could use this procedure to examine the interaction between memory, perception, and attention in the exchange of global contextual information between the auditory and visual domains. This information could eventually be useful in the diagnosis of language comprehension deficits and in the development of subsequent treatments for these deficits.

As previously noted, contextual cueing involves an interaction between memory, perception, and attention. As such, an auditory contextual cueing effect that is analogous to the visual contextual cueing effect provides a means of exploring similarities and differences between visual and auditory attention. Some argue for a common attentional process governing vision and audition (see Shinn-Cunningham, 2008), but if this were true, one would expect that attending to information in one modality would reduce the ability to attend to information in the other modality. Studies show conflicting results on this matter, suggesting that top-down processing modulates attention (Fritz, Elhilali, David, & Shamma, 2007). Chun and Jiang (1998) suggested that encounters with visual search layouts lead to the formation of context maps, and these context maps operate in a top-down fashion to guide attention on subsequent encounters with those layouts. In the auditory modality, Bregman (1990) proposed that auditory stimuli are separated into auditory streams. Thus, one stream is attended to while the other stream is ignored.

Certainly, the alphanumeric sequences and the buzzes used in the current experiments may have been perceived as separate streams, and this may explain the null effect in Experiment 3. This leads to speculation about how top-down processing modulates attention in vision and audition. Is there a common attentional process that operates in both modalities and, therefore, the observed differences in auditory and visual attention are the result of modulating factors? Alternatively, are there unique attention processes for each modality? Auditory contextual cueing provides a means to examine this question more closely.

In conclusion, this study is a novel and important contribution to the research literature, as there are no reports in the literature of an auditory temporal contextual cueing effect that is analogous to the effect demonstrated in vision (Olson & Chun, 2001). My findings suggest that memory for global temporal context interacts with perception and attention in the auditory domain. In addition, I have presented an auditory temporal contextual cueing paradigm that could be used in future research to explore the boundaries of the interaction of memory, perception and attention in the auditory modality. Moreover, researchers could use this paradigm to examine the fundamental processes that underlie cross-modal processing of information between audition and other modalities.

References

- Bregman, A. S. (1990) *Auditory scene analysis: The perceptual organization of sounds*. London: MIT Press.
- Choe, C. S. *et al.* (1975) The 'ventriloquist effect' Visual dominance or response bias? *Perception and Psychophysics*, 18, 55-60.
- Chun, M. M., & Jiang, Y. (1998). Contextual cueing: Implicit learning and memory of visual context guides spatial attention. *Cognitive Psychology*, 36, 28-71.
- Conway, C. M., & Christiansen, M. H. (2005). Modality-constrained statistical learning of tactile, visual, and auditory sequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 31(1), 24.
- de Gelder, B. & Bertelson, P. (2003). Multisensory integration, perception and ecological validity. *Trends in Cognitive Sciences*, 7(10), 460- 467.
- de Gelder, B. & Vroomen, J. (2000). The perception of emotion by ear and by eye. *Cognition and Emotion*, 14, 289-311.
- Easton, R. D., Greene, A. J., Srinivas, K. (1997). Transfer between vision and haptics: Memory for 2-D and 3-D objects. *Psychonomic Bulletin & Review*, 4, 403-410.
- Sánchez-García, C., Alsius, A., Enns, J. T., & Soto-Faraco, S. (2011). Cross-modal prediction in speech perception. *PloS one*, 6(10), e25198.
- Fiser, J. & Aslin, R. N. (2001). Unsupervised statistical learning of higher-order spatial structures from visual scenes. *Psychological Science*, 12, 499-504.
- Fiser, J. & Aslin, R. N. (2002). Statistical learning of higher order temporal structure from visual shape-sequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 28, 458-467.

- Fritz, J. B., Elhilali, M., David, S. V., & Shamma, S. A. (2007). Auditory attention - focusing the searchlight on sound. *Current opinion in neurobiology*, 17(4), 437-455.
- Goujon, A. & Fagot, J. (2013). Learning of spatial statistics in non-human primates: Contextual cueing in baboons (*Papio papio*). *Behavioural Brain Research*, 247, 101-109.
- Habib, M., Rey, V., Daffaure, V., Camps, R., Espesser, R., Joly Pottuz, B., & Dermonet, J.-F. (2002). Phonological training in children with dyslexia using temporally modified speech: A three-step pilot investigation. *International Journal of Language and Communication Disorders*, 37, 289-308.
- Hauser, M. D., Chomsky, N., & Fitch, W. T. (2002). The faculty of language: What is it, who has it, and how did it evolve? *Science*, 298, 1569-1579.
- Heuer, H., Schmidtke, V., & Kleinsorge, T. (2001). Implicit learning of sequences of tasks. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 27(4), 967-983.
- Howard, J. H., Mutter, S. A., & Howard, D. V. (1992). Serial pattern learning by event observation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18(5), 1029-1039.
- Kawahara, J. (2007). Auditory-visual contextual cuing effect. *Perception & Psychophysics*, 69(8), 1399-1408.
- Kelly, S. W., Burton, A. M., Riedel, B., & Lynch, E. (2003). Sequence learning by observation: Evidence for separate mechanisms. *British Journal of Psychology*, 94(3), 355-372.

- Kirkham, N. Z., Slemmer, J. A., & Johnson, S. P. (2002). Visual statistical learning in infancy: Evidence for a domain-general learning mechanism. *Cognition*, *83*, B35-B42.
- Koch, I. & Hoffmann, J. (2000). The role of stimulus-based and response-based spatial information in sequence learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *26*(4), 863-882.
- McDonald, J. J., Teder-Sälejärvi, W. A., & Hillyard, S. A. (2000). Involuntary orienting to sound improves visual perception. *Nature*, *407*, 906-908.
- McGurk, H. and MacDonald, J. (1976). Hearing lips and seeing voices. *Nature*, *264*, 746-748.
- Mayr, U. (1996). Spatial attention and implicit sequence learning: Evidence for independent learning of spatial and nonspatial sequences. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *22*, 350-364.
- Mody, M. (2003). Rapid auditory processing deficits in dyslexia: A commentary on two differing views. *Journal of Phonetics*, *31*, 529-539.
- Mondor, T.A. & Amirault, K. J. (1998). Effect of same- and different- modality spatial cues on auditory and visual target identification. *Journal of Experimental Psychology: Human Perception and Performance*, *24*, 745-755.
- Morris, C. D., Bransford, J. D. & Franks, J. J. (1977). Levels of processing versus transfer-appropriate processing. *Journal of Verbal learning and Verbal Behavior*, *16*, 519-533.
- Nabeta, T., Ono, F., & Kawahara, J. (2003). Transfer of spatial context from visual to haptic search. *Perception*, *32*, 1352-1358.

- Nissen, M. J. & Bullemer, P. (1987). Attentional requirements of learning: Evidence from performance measures. *Cognitive Psychology, 19*, 1-32.
- Olson, I. R., & Chun, M. M. (2001). Temporal contextual cuing of visual attention. *Journal of Experimental Psychology: Learning, Memory and Cognition, 27*, 1299-1313.
- Psychology Software Tools Incorporated. (2003). *E-Prime Software System*. Pittsburg, PA.
- Reber, A. S. (1967). Implicit learning of artificial grammars. *Journal of Verbal Learning and Verbal Behavior, 6*, 855-863.
- Rey, V., De Martino, S., Espesser, R., & Habib, M. (2002). Temporal processing and phonological impairment in dyslexia: Effect of phoneme lengthening on order judgment of two consonants. *Brain and Language, 80*, 576-591.
- Riedel, B. & Burton, A. M. (2006). Auditory sequence learning: Differential sensitivity to task relevant and task irrelevant sequences. *Psychological Research, 70*, 337-344.
- Rubinstein, L. & Gruenberg, E. M. (1971). Intramodal and crossmodal sensory transfer of visual and auditory temporal patterns. *Perception and Psychophysics, 9*, 385-390.
- Saffran, J. R. (2002). Constraints on statistical language learning. *Journal of Memory and Language, 47*, 172-196.
- Saffran, J. R., Aslin, R. N. & Newport, E. L., (1996). Statistical learning by 8-month-old infants. *Science, 274*, 1926-1928.
- Saffran, J. R., Johnson, E. K., Aslin, R. N., & Newport, E. L. (1999). Statistical learning of tone sequences by human infants and adults. *Cognition, 70*, 27652.

- Shams, L. & Seitz, A. R. (2008). Benefits of multisensory learning. *Trends in Cognitive Sciences, 12*(11), 411-417.
- Shelton, A. M. & McNamara, T. P. (2001). Visual memories from nonvisual experiences. *Psychological Science, 12*, 343-347.
- Shinn-Cunningham, B. G. (2008). Object-based visual and auditory attention. *Trends in Cognitive Science, 12* (8), 182-186.
- Tallal, P., Miller, S., & Fitch, R. H. (1993). Neurobiological bases of speech: A case for the pre-eminence of temporal processing. In P. Tallal, A. M. Galaburda, R. R. Llinas, & C. von Euler (Eds.), *Temporal information processing in the nervous system* (pp. 22-47). New York: New York Academy of Sciences.
- Thompson, W. F., Russo, F. A., & Quinto, L. (2008). Auditory-visual integration of emotional cues in song. *Cognition and Emotion, 22*(8), 1457-1470.
- Tulving, E. & Thompson, D. M. (1973). Encoding specificity and retrieval processes in episodic memory. *Psychological Review, 80*, 352-373.
- Van den Stock, J. Grèzes J., & De Gelder, B. (2008). Human and animal sounds influence recognition of body language. *Brain Research, 1242*, 185-190.
- Von Restorff, H. (1933). Über die wirkung von bereichsbildungen im spurenfeld. *Psychologische Forschung, 18*(1), 299-342.

Appendix

University of Manitoba, Department of Psychology
(*this letter was on university letterhead*)

LETTER OF INFORMATION/INFORMED CONSENT

Research Project Title: *Auditory Contextual Cueing*

Principle Investigator (P.I.): Lori Doan **Advisor:** Dr. Jason Leboe-McGowan

P.I. Contact Information:

doan@cc.umanitoba.ca

Room P217, Duff Roblin Building

190 Dysart Road

University of Manitoba

Winnipeg, MB R3T 2N2

This consent form, a copy of which will be left with you for your records and reference, is only part of the process of informed consent. It should give you the basic idea of what the research is about and what your participation will involve. If you would like more detail about something mentioned here, or information not included here, you should feel free to ask. Please take the time to read this carefully and to understand any accompanying information.

Purpose: The general purpose of this research is to further our understanding of human memory and perception.

Description: Participants will make button presses on a keyboard in response to the presentation of a target embedded within an auditory sequence. The study will take approximately 45 minutes to 1 hour to complete.

Recording Devices: Speed and accuracy of button press responses on the keyboard will be recorded by the computer program.

Risks: Although minimal, there is some risk that participants will become bored or irritated by the repetitive procedure. Accuracy feedback is given after each trial and there is some minimal risk that participants' self-esteem may be temporarily negatively affected when receiving feedback indicating an inaccurate response.

Benefits: There will be no direct benefits to participants from their involvement in the study.

Confidentiality: Confidentiality will be respected. Data will be kept in a secure office to which only the researchers will have access. No information that discloses your identity

