

Integrating spatial context learning over contradictory signals: Recency effects in contextual cueing

Christian C. Luhmann

Stony Brook University, Stony Brook, NY, USA

Contextual cueing is a visual search phenomenon in which memory of global visual context guides spatial attention towards task-relevant portions of the search display. Recent work has shown that the learning processes underlying contextual cueing exhibit primacy effects; they are more sensitive to early experience than to later experience. These results appear to pose difficulties for associative accounts which typically predict recency effects; behaviour being most strongly influenced by recent experience. The current study utilizes trial sequences that consist of two contradictory sets of regularities. In contrast to previous results, robust recency effects were observed. In a second study it is demonstrated that this recency effect can be minimized, but not reversed by systematically manipulating task-irrelevant features of the search display. These results provide additional support for an associative account of contextual cueing and suggest that contextual cueing may, under some circumstances, be more sensitive to recent experience.

Keywords: Contextual cueing; Learning; Visual search.

Adaptive behaviour requires organisms to learn from past experience. It is thus unsurprising that processes dedicated to the detection of regularities are found across a variety of sensory modalities. The visual system, in particular, appears to extract a variety of spatial and temporal regularities from streams of input (Chun & Turk-Browne, 2008; Turk-Browne, Isola, Scholl, & Treat, 2008). Much of this processing occurs outside of our awareness, but must nonetheless be flexible enough to deal with a constantly changing environment. Recent work (Brockmole & Henderson, 2006; Jungé, Scholl, & Chun, 2007; Makovski & Jiang, 2010; Manginelli & Pollmann, 2009) has begun to

Please address all correspondence to Christian C. Luhmann, Department of Psychology, Stony Brook University, Stony Brook, NY 11794-2500, USA. E-mail: christian.luhmann@stonybrook.edu

The author would like to thank Evelyn Starosta for help with data collection and Arty Samuel and Greg Zelinsky for discussions pertaining to this study.

investigate how these learning processes operate over changing patterns of experience. In the current study, this line of investigation is extended and qualified.

ORDER EFFECTS IN CONTEXTUAL CUEING

Contextual cueing (Chun, 2000; Chun & Jiang, 1998) refers to a visual search phenomenon in which memory of visual context can guide spatial attention towards task-relevant portions of the search display. Cueing is measured by repeating search displays over the course of the experiment and comparing response times to these repeated displays with response times to novel, unfamiliar search displays. Faster responses to the repeated displays are taken as evidence that participants have retained and utilized information about the global visual context. This phenomenon has been linked to the broader literature on visual statistical learning (Fiser & Aslin, 2002).

Recently, several studies (Jungé et al., 2007; Manginelli & Pollmann, 2009) have investigated whether contextual cueing is sensitive to the presentation order of statistical regularities. To answer this question, these studies presented participants with trial sequences that consisted of two different sets of statistical regularities, one set presented early in the sequence and the other presented late in the sequence. By manipulating the presentation order, but equating the statistical properties of the input, these researchers have been able to isolate the influence of presentation order per se.

The first of these studies, reported by Jungé et al. (2007), presented a contextual cueing paradigm in which the learning sequence consisted of blocks of “signal” and “noise”. Signal trials were trials on which the global context was predictive of the target location (as in a typical contextual cueing experiment, e.g., Chun & Jiang, 1998), whereas noise trials were trials on which the global context had no systematic relationship with the target location. Participants were presented with either a set of signal trials followed by a set of noise trials (the signal-first condition), or a set of noise trials followed by a set of signal trials (the noise-first condition).

The results reported by Jungé et al. (2007) demonstrated that the signal-first condition led to normal contextual cueing effect, whereas the noise-first condition produced no evidence of learning. This pattern represents a primacy effect; learning was more heavily influenced by early experience than by later experience. These results were taken as evidence that, “implicit contextual learning may have an intrinsic endpoint—resulting in a type of functional hypothesis about regularities in displays . . . that then becomes resistant to further updating” (p. 9). That learning processes “turn off” after initial experience may indicate an efficiency designed to free learning-related resources.

Manginelli and Pollmann (2009) report a similar manipulation. In their study, participants completed a visual search task in which regularities were present throughout the experiment. However, for the first two-thirds of the sequence each context was associated with one target location. At that point in the trial sequence, each target was moved to a new, static location for the remaining third of the sequence. Over the course of the first two-thirds of the experiment, the repeated contexts came to elicit significantly faster responses than novel contexts. When the target was moved, this facilitation was eliminated. In the remaining third of the trial sequence (10 blocks), the new target location was static and thus predictable. However, despite the presence of regularities in the last third of the sequence, no cueing advantage was ever observed for the new target locations. Like the Jungé et al. (2007) study, these results are consistent with the idea that learning processes had “shut off” before encountering the second set of contradictory regularities.

These reports of primacy during contextual cueing represent an interesting divergence from traditional associative learning theory. For example, the classic Rescorla-Wagner model of conditioning (Rescorla & Wagner, 1972) suggests that associations are updated after every trial, resulting in robust recency effects (i.e., learning more heavily influenced by recent experience). This prediction is supported by a variety of work on both humans (LaBar, Gatenby, Gore, LeDoux, & Phelps, 1998) and nonhuman animals (Wagner, 1961). In contrast, there have been primacy effects observed in more deliberative learning tasks such as impression formation (Asch, 1946; Hendrick & Costantini, 1970), categorization (Duffy & Crawford, 2008), and contingency learning (Dennis & Ahn, 2001; Yates & Curley, 1986). The primacy effects in these more explicit tasks have typically been taken as evidence for a hypothesis-formation process that uses initial experience to generate beliefs which then dominate subsequently encountered information (e.g., Hogarth & Einhorn, 1992). Thus, the primacy effects observed in contextual cueing may actually pose a challenge for associative accounts of contextual cueing (e.g., Brady & Chun, 2007; Makovski & Jiang, 2010).

However, the generality of these primacy effects is tempered by the methods used. For example, Jungé et al. (2007) utilized sequences in which the training phase consisted of a combination of signal and noise trials. During the noise portions of the task, there was nothing to be learned. In contrast, previous investigations of order effects have utilized only signal trials and have instead manipulated the signal over the course of the sequence. For example, in a study by Labar et al. (1998), the learning sequence consisted of blocks in which a coloured shape systematically predicted the presence of an electric shock and blocks in which the same shape systematically predicted the absence of an electric shock. There was always a predictive signal to be learned, this signal simply changed over the course of the experiment. Manginelli and Pollmann (2009) did utilize such

sequences; search targets were presented in predictable locations and subsequently moved to new, predictable locations. However, because this study was not designed to investigate order effects per se, the two sets of regularities were not presented equally often. Instead, the regularities presented early in the sequence were presented on twice as many trials as the regularities presented later in the sequence.

The current study seeks to examine whether primacy is a general feature of contextual cueing and what factors may promote either recency or primacy effects. This is accomplished by utilizing a two-block sequence of trials like that used by Manginelli and Pollmann (2009). In the first half of the sequence, each global visual context reliably predicts the target location as in a traditional contextual cueing paradigm. In the second half of the sequence, each global visual context is identical to those used in the first half, but now reliably predicts a new, contradictory target locations. Participants then complete a transfer phase in which they are presented with displays with their targets located in the initial position (from the first half of the sequence) or final position (from the second half of the sequence). If participants are able to more quickly locate the target in its initial location than they are to locate the target in its final location, this pattern of results would be indicative of a primacy effect. This transfer phase also included new, novel contexts in order to provide further assessment of learning.

EXPERIMENT 1

Method

Participants. Fourteen Stony Brook University undergraduates participated for partial course credit.

Stimuli and task. Participants completed a standard visual search task in which they located and responded to a “T”-shaped target embedded amongst a field of “L”-shaped distractors. Shapes subtended approximately 2.5° of visual angle. The “T”-shaped targets were each rotated 90° or 270° such that the longer line segment of the “T” pointed either left or right. The “L”-shaped distractors were each rotated 0° , 90° , 180° , or 270° . In addition, the intersection of the two line segments was offset by approximately 0.15° to increase the difficulty of the search task. All shapes were black and were presented on a neutral grey background. Search arrays were constructed by placing 11 distractors and one target onto a 6×8 grid (the grid was not visible to participants). The placement of the 12 shapes was constrained such that three of the 12 shapes were presented in each of the four quadrants. The target was additionally restricted from appearing in the extreme corners of the display or in central locations leaving 28 possible target locations. Last,

search arrays always had an empty slot directly opposite the target location at an equivalent eccentricity (see Figure 1).

On each trial, a search array was displayed and participants were required to indicate which direction the target shape was pointing. Responses were made using the left and right arrow keys on the keyboard. After a correct response, the search array was removed from the screen and the next trial began 500 ms later. If an incorrect response was made or if no response was received within 7 s of stimulus onset, the search array was removed from the screen and the word “Incorrect” was displayed for 1000 ms. The screen was then cleared and the next trial began 500 ms later. Participants were offered a break after every 16 trials.

Design and procedure. At the beginning of the experiment, 16 new search arrays were constructed according to the constraints outlined above. During the construction of the array, each of the distractors was randomly mirrored (.5 probability of being mirrored) and rotated to a randomly selected orientation (0° , 90° , 180° , or 270°). The mirroring and orientation of a given distractor in a given search array did not change during the experiment. In contrast, the direction of the target (pointing left or pointing right) was randomly determined on each trial (50% left, 50% right).

The experiment consisted of a training phase and a transfer phase. The training phase consisted of 10 blocks of trials where a block consisted of a single presentation of each of the 16 search arrays (i.e., 160 trials in total). For the first half of training (the first five blocks), the arrays were presented




Figure 1. A sample search array utilized in the current study. The array on the left represents the initial array, shown throughout the first half of the training phase. The array on the right represents the change in target location made halfway through training and the resulting array that was presented throughout the second half of the training phase. The grid is shown for illustrative purposes only and was never shown to participants. In the actual experiment, the shapes were presented on a neutral grey background.

as initially constructed. For the second half of training (the final five blocks), the target shape was relocated to the slot directly opposite in the search array (i.e., its location was “mirrored” about the origin; see Figure 1). The two halves of the training sequence were contiguous and thus participants were given no indication that anything had changed. It should be noted that our sequence was perfectly symmetric with each set of regularities constituting exactly half of the training sequence (cf. Manginelli & Pollmann, 2009). Also note that all contexts presented during the training phase were predictive. Thus, the time course of learning during the training sequence itself cannot be assessed.

At the beginning of the transfer phase, eight of the 16 search arrays were randomly selected to have their target placed in the initial target location. The remaining eight search arrays had their targets placed in the final target location. A given search array was only tested in one of the two target locations during transfer. In addition, eight brand new search arrays were constructed under the same constraints used to construct the original arrays. Each of the 24 search arrays (eight initial target locations, eight final target locations, eight novel arrays) was presented five times (for a total of $24 \times 5 = 120$ trials).

After completing informed consent and receiving brief instructions and brief practice with the task (20 trials using search arrays not used in the actual experiment), participants completed the main search task. Each participant completed the training phase and the associated transfer phase. So as to maximize statistical power, participants then took a short break and then completed the entire experiment a second time (utilizing entirely new stimuli). The entire procedure took approximately 30 minutes.

Results

Response times for incorrect responses were excluded as were any response times more than three standard deviations above the participant’s average. Response times from the transfer phase were averaged separately for each repetition of the task and separately for the three conditions: Initial target location, final target location, and novel array. The pattern of response times did not vary across repetitions of the task, so each participant’s data was collapsed across the repetitions. Average transfer response times are presented in Figure 2A. Evaluating the critical target location manipulation, participants were significantly faster to respond when the target was located in its final location ($M = 1749$ ms, $SD = 307.8$) than when the target was located in its initial location ($M = 1860$ ms, $SD = 336.6$), $t(13) = 2.27$, $p < .05$. The trained arrays were next compared with the novel transfer arrays. Participants were significantly slower to respond to the novel transfer arrays ($M = 1910$ ms, $SD = 274.7$) than when the target was located in its

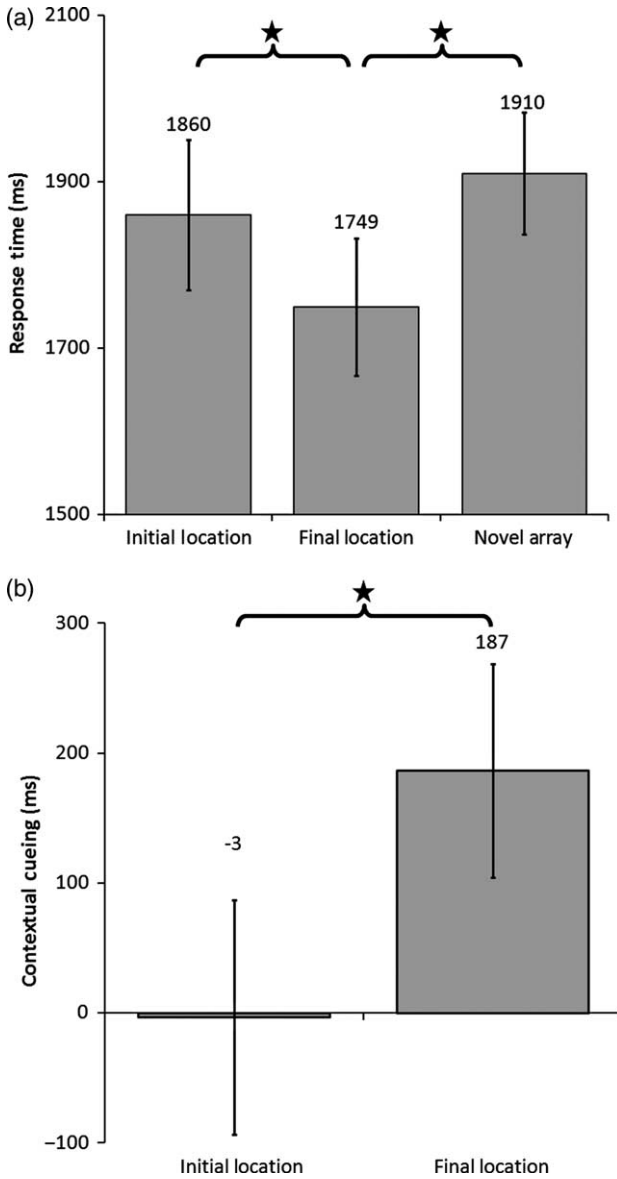


Figure 2. (A) Raw reaction times for the three conditions presented in the transfer phase of Experiment 1. (B) Contextual cueing computed as the transfer response time normalized by average response time during training.

final location, $t(13) = 3.45$, $p < .005$. In contrast, reaction times to targets located in their initial location did not differ from reactions times to the novel arrays, $t(13) = 1.20$, $p = .25$.

Reaction times during the transfer phase were next compared to reaction times during the training phase. This was done by computing the average reaction time for each search array across the entire training phase. Transfer reaction times were then subtracted from training reaction times for that search array during training. As can be seen in Figure 2B, placing the target in its final location resulted in a significant contextual cueing effect of 187 ms, $t(13) = 4.13$, $p < .005$. No such effect was shown when placing the target in its initial location (-3 ms, $t < 1$), indicating that transfer reaction times were no faster than training reaction times for these search arrays. Furthermore, the cueing advantage for targets in their final locations was significantly greater, $t(13) = 3.30$, $p < .01$, than for targets in their initial locations.

Summary and discussion

The current experiment presented repeated global visual contexts and paired them with two, conflicting target locations an equal number of times. When tested at each of the trained target locations, participants exhibited greater contextual cueing to the most recently observed target positions (a recency effect). These results demonstrate that the visual statistical learning processes that underlie contextual cueing do not always produce the primacy effects reported recently (Jungé et al., 2007; Manginelli & Pollmann, 2009). Instead, our results suggest that these processes may instead track recently encountered regularities, at least under certain circumstances.

Curiously, Experiment 1 provides no actual evidence that any learning took place during the initial portions of the training sequence. Given that previous manipulations of order effects have found evidence for strong initial learning effects, it seems unlikely that the current results reflect a complete absence of early learning. Instead, it seems more likely that early learning was simply not expressed during the transfer phase. I consider three related factors that may help to explain why such expression would be impaired. First, the representations of early regularities may be “overwritten” once new, incompatible regularities are encountered (so-called catastrophic interference). Second, because they are simply older, memory for older regularities may have decayed by the time participants reached transfer. Third, the representations of early regularities may be perfectly intact; however, the newer, more recent regularities may produce interference during transfer, making the retrieval of the older experience more difficult.

Experiment 2 is designed to begin distinguishing between these mechanisms. This is accomplished by providing greater differentiation between the two, incompatible sets of regularities. Specifically, target and distractor shapes are one colour during the first half of the task and a second colour during the second half. The strategy of manipulating task-irrelevant stimulus features during learning has been utilized in studies of both language learning (Gebhart, Aslin, & Newport, 2009; Weiss, Gerfen, & Mitchel, 2009) and perceptual learning in speech processing (Eisner & McQueen, 2005; Kraljic & Samuel, 2007). Changing incidental features of the input may allow learning processes to more easily distinguish between the two, incompatible sets of regularities. The colour change may also act as an occasion setter (Holland, 1992), a concept that has long been discussed in the associative learning literature. Occasion setters are not, themselves, predictive of anything (e.g., target location) but instead specify which cues the learner should use to guide behaviour. Thus, if colour acts as an occasion setter in the current experiment, participants could learn that a given context in one colour should be associated with one target location but that that same context in a different colour should be associated with a different target location. At transfer, the occasion setter would also help to facilitate retrieval of the older regularities in the face of interference from the newer, contradictory regularities. Given the substantial evidence of occasion setting in other associative learning tasks, this pattern would also provide further support for an associative account of contextual cueing.

EXPERIMENT 2

Method

Participants. Thirty-five Stony Brook University undergraduates participated for partial course credit.

Design. The task and stimuli were nearly identical to those used in Experiment 1 with the following differences. The training phase consisted of 10 blocks, each of which included a single presentation of the 16 search arrays. The entire sequence was again divided into two halves, each of which consisted of five blocks of trials. During the first half of training, when the target was in its initial location, all of the shapes in all the arrays were either black or white (counterbalanced across participants) on a neutral grey background. Halfway through the experiment, when the target was relocated across the search array (just as in Experiment 1), the colour of all the shapes in all the arrays was switched. That is, participants presented with black shapes during the first half of training saw white shapes during the second

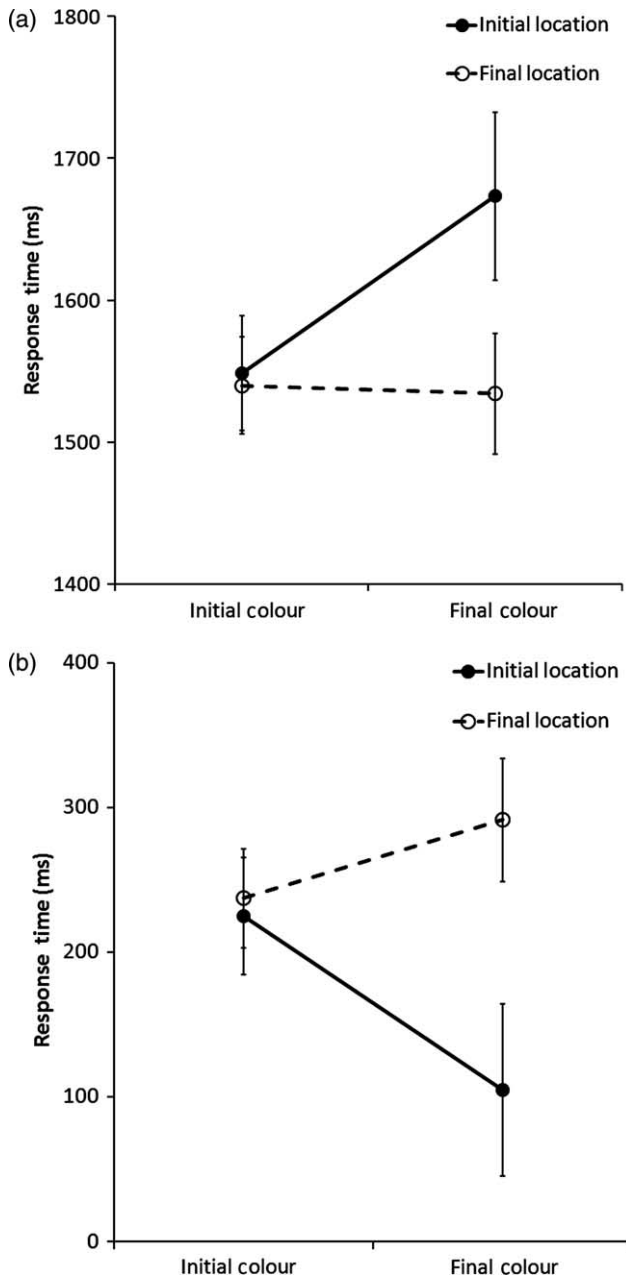


Figure 3. Results in Experiment 2 exhibited a significant interaction between target location and colour. (A) Raw transfer reaction times for the four conditions in Experiment 2. (B) Contextual cueing computed as the transfer response time normalized by average response time during training.

half of training, whereas those presented with black shapes during the first half of training saw white shapes during the second half of training.

During the transfer phase, the 16 search arrays were randomly divided into four sets of four arrays. Each set of four arrays was assigned to a specific combination of target location (initial/final) and array colour (initial/final). Thus, a given search array was only tested in one of the four combinations during transfer. The transfer phase consisted of eight blocks, each of which included a single presentation of each search array for a total of 128 trials.

RESULTS

Response times for incorrect responses were excluded as were any response times more than three standard deviations above the participant's average. To evaluate learning, participants' response times were submitted to a 2 (target location: Initial vs. final) \times 2 (array colour: Initial vs. final) repeated measures ANOVA. A marginal main effect of target location was observed, $F(1, 30) = 4.15$, $p = .0504$, with targets in their final locations eliciting significantly faster reaction times ($M = 1537$ ms, $SD = 276.5$) than targets in their initial locations ($M = 1611$ ms, $SD = 298.0$). A marginal interaction between factors was also observed, $F(1, 30) = 4.10$, $p = .052$. Post hoc tests revealed that when the array was presented in its final colour, reaction times were significantly faster, $t(34) = 2.34$, $p < .05$, when the target was presented in its final location ($M = 1534$ ms, $SD = 321.3$) than when the target was presented in its initial location ($M = 1673$ ms, $SD = 382.5$). In contrast, when the array was presented in its initial colour, reaction times did not differ by target location ($t < 1$).

As in Experiment 1, reactions times during the transfer phase were also compared to reaction times during the training phase by computing the cueing advantage measure described previously. These results can be seen in Figure 3. These results were again submitted to a 2 (target location: Initial vs. final) \times 2 (array colour: Initial vs. final) repeated measures ANOVA. A significant main effect of target location was observed, $F(1, 30) = 6.16$, $p < .05$, with targets in their final locations eliciting a significantly greater cueing effect ($M = 264$ ms, $SD = 183.1$) than targets in their initial locations ($M = 165$ ms, $SD = 234.9$). A significant interaction between factors was also observed, $F(1, 30) = 6.45$, $p < .05$. Post hoc tests revealed that when the array was presented in its final colour, the cueing effect was significantly greater, $t(34) = 3.02$, $p < .005$, when the target was presented in its final location ($M = 291$ ms, $SD = 251.1$) than when the target was presented in its initial location ($M = 104$ ms, $SD = 351.3$). In contrast, when the array was presented in its initial colour, the cueing effect did not differ by target location ($t < 1$).

SUMMARY

The results from Experiment 2 suggest that learning was operating continuously over the course of the sequence and that the learning produced during the early portions of experience is not necessarily discarded once new, contradictory regularities are encountered. Responses to more recent visual contexts were generally faster than responses to older contexts; however, manipulation of the task-irrelevant display colour was able to completely eliminate this effect. Clearly, if participants were not learning during the early portions of the task or if the output of that early learning were discarded, the colour of the array could not have had any effect. That being said, the colour manipulation was not able to elicit a primacy effect; under no conditions did our participants exhibit greater contextual cueing to the initial target locations than to the final target locations. This complete lack of primacy suggests that older representations were, in general, at an overall disadvantage. The colour of the array could only act to minimize this disadvantage. Furthermore, the results are broadly consistent with the idea that colour was acting as an occasion setter, allowing behavioural flexibility in the face of contradictory experience and generally suggesting that contextual cueing results from associative learning processes.

Our use of colour to modulate contextual cueing is similar to a recent study by Geyer, Shi, and Müller (2010). In their study, individual search arrays consisted of two subsets of items defined by colour. Each colour subset could be independently predictive or nonpredictive of the target location. Their results demonstrated facilitation by predictive distractors that matched the target colour and also by predictive distractors that conflicted with the target colour, though facilitation was weaker in the latter case. Their study was somewhat different from the current one in that the different colours used by Geyer et al. were simultaneously present on each trial. In addition, their study utilized a conjunction search in which the target was jointly defined by colour and orientation. Nonetheless, the colour manipulations in each study were conceptually similar, allowing the contribution of two different sets of regularities to be evaluated independently.

GENERAL DISCUSSION

The results of two experiments demonstrate, in contrast to recent reports (Jungé et al., 2007; Manginelli & Pollmann, 2009), that the visual statistical learning processes that underlie contextual cueing can operate continuously over the course of experience even when this experience is not homogenous. Furthermore, when in continuous operation, these learning processes do not simply aggregate over the cumulative experience. Instead, significantly

greater emphasis is placed on more recent experience. The observed recency effects appear to be driven by a combination of factors. First, recency appears to be driven, in part, by retrieval interference, in which later experience tends to inhibit the retrieval of older experience. Second, there appears to be a pure recency effect in which older experience is simply less accessible than more recent experience, perhaps partially due to the decay of the relevant memories.

A critical difference between the current study and those reporting primacy may lie in the way that regularities were manipulated. For example, Jungé et al. (2007) utilized a sequence consisting of blocks in which regularities were present and blocks in which regularities were absent. Given the current results, it may well be the case that the presence of changing regularities and the absence of regularities may exert differing influences on statistical learning processes. When regularities are absent from experience, learning processes may indeed terminate. The current results suggest that, when regularities are continuously present in the input, even when those regularities change in unpredictable ways, learning processes may attempt to adapt to the changing environment. The sequences utilized by Manginelli and Pollmann (2009), consisted of a continuous, but changing signal just as in the current study. However, Manginelli and Pollmann changed the target location two-thirds of the way through the trial sequence. This meant that the processes underlying cueing had twice as many trials to learn the initial target locations as they did to learn the new target locations. Thus, the sequences were not ideal to find strong evidence of relearning. However, their results do suggest that there was likely a gradual transition as participants were confronted and dealt with the new, contradictory regularities present in the latter portions of the current experiments.

Despite the ability to operate over disparate input, the learning processes that underlie contextual cueing appear to be subject to known constraints of memory. Given two sets of incongruent input, the visual statistical learning processes do not produce an “average” representation reflecting the aggregate experience. Instead, older experience appears to decay over time, reducing its ability to guide behaviour. In addition, newer contradictory experience appears to interfere with older experience hindering its influence even further. Note that the presence of these memory effects need not be interpreted as flaws. As has been noted (Jungé et al., 2007), if one’s environment is stable, then it may be cognitively efficient to build “hypotheses” based on initial experience and then terminate subsequent learning. However, in a changing environment, these initial experiences can quickly become outdated and lead to problematic behaviour. In dynamic environments, using one’s most recent experience as a basis for future behaviour seems like a more adaptive strategy. Thus, the memory constraints imposed on learning may actually be functional.

It is important to recognize that the current results provide evidence about how the learning that produces contextual cueing operates. They do not, however, provide good evidence for what is learned. This is in contrast to much of the previous work on contextual cueing which has focused on what features of the visual array are utilized to facilitate search. For example, there have been investigations into whether contextual cueing relies on associations between target locations and the holistic, global context, or individual distractor locations and identities (Brady & Chun, 2007; Brockmole, Castelano, & Henderson, 2006; Jiang & Wagner, 2004). Other work has explored whether contextual cueing relies on memory for specific visual features such as shape, size, colour, and conjunctions thereof (e.g., Ehinger & Brockmole, 2008; Geyer et al., 2010; Jiang & Wagner, 2004; Kunar, Flusberg, & Wolfe, 2006). The current results, in contrast, appear to place no strong constraints on which aspects of the visual context participants learned to guide their visual search. Instead, the current results speak directly to the learning process itself. Recency effects are inconsistent with learning accounts that posit the formation of perceptual hypotheses because these sorts of accounts predict primacy effects (e.g., Dennis & Ahn, 2001; Jungé et al., 2007; Luhmann & Ahn, in press). Instead, recency effects suggest that representations are continuously updated as experience is accumulated and that older, conflicting experience exerts a diminished influence on behaviour. Importantly, formal descriptions of the learning processes that predict recency are agnostic about what features are *predictive* and what features are *predicted*. For example, Brady and Chun (2007) have put forth a formal account of contextual cueing that likely produces recency effects. However, this prediction is not a consequence of learning the local context (the main thesis of that study), but of the nature of the updating procedure their model uses (i.e., the delta rule; Williams, 1992).

Similarly, the current results cannot directly speak to the recent controversy (Smyth & Shanks, 2008) pertaining to whether contextual cueing is truly an implicit phenomenon. Though the current studies did not employ an explicit memory test, traditionally, explicit learning processes have been strongly associated with primacy (e.g., Hogarth & Einhorn, 1992; Yates & Curley, 1986). Thus, the strong recency effects found here are at least consistent with the traditional idea that contextual cueing is governed by implicit processes. However, it should be noted that the current results do stand in contrast to several reports of primacy effects found utilizing implicit learning paradigms and nonvisual modalities. For example, Gebhart, Aslin, and Newport (2009) report an auditory statistical learning study in which the sequences consisted of nonsense words generated by one grammar followed by nonsense words generated by a second grammar. Participants in these studies consistently exhibited primacy, showing little or no learning of the second grammar. In one particularly powerful experiment, participants

received 5 minutes of exposure to the first grammar followed by 15 minutes of exposure to the second grammar. Despite the highly biased sequence, participants were still significantly better at recognizing items from the initial grammar. Similar primacy effects have been reported in other auditory experiments (e.g., Eisner & McQueen, 2005; Kraljic & Samuel, 2007). Thus, further work will be required to determine whether these reflect modality differences or more structural facets of the different tasks.

Nonetheless, for those concerned with the theoretical underpinnings of contextual cueing, the current results, particularly the apparent similarities between contextual cueing and traditional theories of learning, should give some comfort. Various theories of learning (e.g., Rescorla & Wagner, 1972), including those used to account for aspects of contextual cueing (Brady & Chun, 2007; Chun & Turk-Browne, 2008), have assumed that learned representations are more influenced by recent experience. This allows these theories to avoid some of the thornier issues associated with explaining primacy effects (Luhmann & Ahn, in press). These issues include specifying the conditions under which learning starts and stops, whether and how learning can be reinitiated, and the fate of any input encountered while learning is “turned off”.

That being said, the fact that contextual cueing can, under different circumstances, exhibit both recency and primacy suggests that additional work is necessary. In particular, these findings must be integrated both empirically and theoretically in order to paint a more coherent picture of the learning mechanisms that give rise to contextual cueing. The current results suggest that one fruitful future direction may be to investigate when and why learning processes treat the presence and absence of regularities differently.

REFERENCES

- Asch, S. E. (1946). Forming impressions of personality. *Journal of Abnormal and Social Psychology, 41*(3), 258–290.
- Brady, T. F., & Chun, M. M. (2007). Spatial constraints on learning in visual search: Modeling contextual cuing. *Journal of Experimental Psychology: Human Perception and Performance, 33*(4), 798–815.
- Brockmole, J. R., Castelano, M. S., & Henderson, J. M. (2006). Contextual cueing in naturalistic scenes: Global and local contexts. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 32*(4), 699–706.
- Brockmole, J. R., & Henderson, J. M. (2006). Recognition and attention guidance during contextual cueing in real-world scenes: Evidence from eye movements. *Quarterly Journal of Experimental Psychology, 59*(7), 1177–1187.
- Chun, M. M. (2000). Contextual cueing of visual attention. *Trends in Cognitive Sciences, 4*(5), 170–178.
- Chun, M. M., & Jiang, Y. (1998). Contextual cueing: Implicit learning and memory of visual context guides spatial attention. *Cognitive Psychology, 36*, 28–71.

- Chun, M. M., & Turk-Browne, N. B. (2008). Associative learning mechanisms in vision. In S. J. Luck & A. R. Hollingworth (Eds.), *Visual memory* (pp. 209–247). New York, NY: Oxford University Press.
- Dennis, M., & Ahn, W. (2001). Primacy in causal strength judgments: The effect of initial evidence for generative versus inhibitory relationships. *Memory and Cognition*, 29(1), 152–164.
- Duffy, S., & Crawford, L. E. (2008). Primacy or recency effects in forming inductive categories. *Memory and Cognition*, 36(3), 567–577.
- Ehinger, K. A., & Brockmole, J. R. (2008). The role of color in visual search in real-world scenes: Evidence from contextual cuing. *Perception and Psychophysics*, 70(7), 1366–1378.
- Eisner, F., & McQueen, J. M. (2005). The specificity of perceptual learning in speech processing. *Perception and Psychophysics*, 67(2), 224–238.
- Fiser, J., & Aslin, R. N. (2002). Statistical learning of new visual feature combinations by infants. *Proceedings of the National Academy of Sciences of the USA*, 99(24), 15822–15826.
- Gebhart, A. L., Aslin, R. N., & Newport, E. L. (2009). Changing structures in midstream: Learning along the statistical garden path. *Cognitive Science*, 33(6), 1087–1116.
- Geyer, T., Shi, Z., & Müller, H. J. (2010). Contextual cuing in multiconjunction visual search is dependent on color- and configuration-based intertrial contingencies. *Journal of Experimental Psychology: Human Perception and Performance*, 36(3), 515–532.
- Hendrick, C., & Costantini, A. (1970). Effects of varying trait inconsistency and response requirements on the primacy effect in impression formation. *Journal of Personality and Social Psychology*, 15(2), 158–164.
- Hogarth, R. M., & Einhorn, H. J. (1992). Order effects in belief updating: The belief-adjustment model. *Cognitive Psychology*, 24(1), 1–55.
- Holland, P. C. (1992). Occasion setting in Pavlovian conditioning. In D. L. Medin (Ed.), *The psychology of learning and motivation* (Vol. 28, pp. 69–125). New York, NY: Academic Press.
- Jiang, Y., & Wagner, L. C. (2004). What is learned in spatial contextual cuing-configuration or individual locations? *Perception and Psychophysics*, 66(3), 454–463.
- Jungé, J. A., Scholl, B. J., & Chun, M. M. (2007). How is spatial context learning integrated over signal versus noise? A primacy effect in contextual cuing. *Visual Cognition*, 15(1), 1–11.
- Kraljic, T., & Samuel, A. G. (2007). Perceptual adjustments to multiple speakers. *Journal of Memory and Language*, 56(1), 1–15.
- Kunar, M. A., Flusberg, S. J., & Wolfe, J. M. (2006). Contextual cuing by global features. *Perception and Psychophysics*, 68(7), 1204–1216.
- LaBar, K. S., Gatenby, J. C., Gore, J. C., LeDoux, J. E., & Phelps, E. A. (1998). Human amygdala activation during conditioned fear acquisition and extinction: A mixed-trial fMRI study. *Neuron*, 20(5), 937–945.
- Luhmann, C. C., & Ahn, W. (in press). Order effects during learning: Expectations and interpretations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*.
- Makovski, T., & Jiang, Y. V. (2010). Contextual cost: When a visual-search target is not where it should be. *Quarterly Journal of Experimental Psychology*, 63(2), 216–225.
- Manginelli, A. A., & Pollmann, S. (2009). Misleading contextual cues: How do they affect visual search? *Psychological Research*, 73(2), 212–221.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64–99). New York, NY: Appleton-Century-Crofts.
- Smyth, A. C., & Shanks, D. R. (2008). Awareness in contextual cuing with extended and concurrent explicit tests. *Memory and Cognition*, 36(2), 403–415.

- Turk-Browne, N. B., Isola, P. J., Scholl, B. J., & Treat, T. A. (2008). Multidimensional visual statistical learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34(2), 399–407.
- Wagner, A. (1961). Effects of amount and percentage of reinforcement and number of acquisition trials on conditioning and extinction. *Journal of Experimental Psychology*, 62(3), 234–242.
- Weiss, D. J., Gerfen, C., & Mitchel, A. D. (2009). Speech segmentation in a simulated bilingual environment: A challenge for statistical learning? *Language Learning and Development*, 5(1), 30–49.
- Williams, R. J. (1992). Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8(3), 229–256.
- Yates, J., & Curley, S. (1986). Contingency judgment: Primacy effects and attention decrement. *Acta Psychologica*, 62(3), 293–302.

Manuscript received July 2010

Manuscript accepted March 2011

First published online July 2011