Research Report

TOP-DOWN ATTENTIONAL GUIDANCE BASED ON IMPLICIT LEARNING OF VISUAL COVARIATION

Marvin M. Chun and Yuhong Jiang

Yale University

Abstract—*The visual environment is extremely rich and complex, producing information overload for the visual system. But the environment also embodies structure in the form of redundancies and regularities that may serve to reduce complexity. How do perceivers internalize this complex informational structure? We present new evidence of visual learning that illustrates how observers learn how objects and events covary in the visual world. This information serves to guide visual processes such as object recognition and search. Our first experiment demonstrates that search and object recognition are facilitated by learned associations (covariation) between novel visual shapes. Our second experiment shows that regularities in dynamic visual environments can also be learned to guide search behavior. In both experiments, learning occurred incidentally and the memory representations were implicit. These experiments show how top-down visual knowledge, acquired through implicit learning, constrains what to expect and guides where to attend and look.*

Top-down perceptual knowledge helps prescribe what visual objects to expect and which objects demand more attention in a particular behavioral context. For example, the visual context of a scene (e.g., kitchen) facilitates recognition of objects (such as toaster or refrigerator) relevant to that context (Biederman, Mezzanotte, & Rabinowitz, 1982; Palmer, 1975; but see Hollingworth & Henderson, 1998), while irrelevant features and objects tend to receive less processing priority (Rensink, O'Regan, & Clark, 1997). Imagine the outdoor scenery you encounter every day on your way to work. The spatial layout of landmarks in the environment tends to be stable over time and provides useful cues for navigation. As you drive, you may notice that even moving objects (such as other cars) move around in somewhat predictable ways. If not for such predictability, how could one maintain sanity while switching lanes toward the off-ramp on a highway commute?

These examples illustrate that the visual world is highly structured, such that objects and events tend to covary in predictable, invariant ways (Biederman, 1972). Knowledge of this structure may serve to reduce the large amount of uncertainty and complexity that exists in the stimulus input (Gibson, 1969, 1991) and can provide useful constraints on visual processes such as object recognition or search. Given the richness of visual structure in the environment (Garner, 1974), how do perceivers utilize this information? It seems likely that environmental structure must be represented within the brain in order to constrain perceptual processing. Hence, it is important to consider how such visual structure is learned by observers through interactions with the visual world.

We recently introduced a new paradigm to examine how useful top-down knowledge is extracted from visual experience (Chun & Jiang, 1998). We showed that perceivers are sensitive to the context of visual targets, so that implicit learning and memory of novel visual context information can guide attentional deployment, facilitating visual behaviors such as search. We described this process as *contextual cuing.* Subjects were asked to locate and identify a rotated-T target presented among rotated-L distractor objects. This is a serial search task that requires attentional scrutiny of the display (Duncan & Humphreys, 1989; Treisman & Gelade, 1980; Wolfe, Cave, & Franzel, 1989). A role for context influencing search was examined as follows. First, global context was defined as the spatial layout of the distractor objects in the visual search array. Second, we generated a set of these displays and made them invariant by repeating them across blocks throughout the experimental session. We refer to these as *old* contexts. Finally, targets appeared in consistent locations within their invariant contexts. Hence, the invariant configurations were predictive of the location of embedded targets. The question was whether subjects would be sensitive to this contextual information. If so, over time, search would be facilitated for targets appearing in old contexts, relative to targets appearing in new contexts, randomly generated in each block to serve as a control.

In several experiments, search performance was faster for targets appearing in old contexts than for targets appearing in new contexts (Chun & Jiang, 1998). This benefit was termed contextual cuing because visual context information served to cue spatial attention to target locations. The context was encoded implicitly during visual search; subjects were never instructed to encode the array, nor did anyone report ever having tried to. The resulting memory representations that guided attention were also implicit. An explicit recognition test revealed that subjects could not consciously discriminate between old and new contexts. Hence, memory and attention interact in important ways, with perceptual processing being influenced by implicit memory traces of past perceptual interactions. In other words, visual learning allows for useful and specific top-down knowledge to be derived and applied to subsequent acts of perception.

The goal of the present study was to show that visual learning occurs for a wide variety of qualitatively different types of visual information important for perception and behavior. Our previous study (Chun & Jiang, 1998) focused on cuing of target location on the basis of implicit learning of spatial layout. Spatial layout is an ecological variable because landmarks and global structures, as well as their configurations in the environment, tend to be stable over time. However, visual learning mechanisms must be able to encode other types of visual regularities and structure, too. Here, Experiment 1 tested whether associations between novel objects can be learned to facilitate perceptual processing. Such object association learning would be fundamental for constructing perceptual schemas. Experiment 2 examined whether dynamic regularities are encoded. Sensitivity to such regularities is important because observers interact with a dynamic visual environment in which objects move about and change in a regular manner.

EXPERIMENT 1

Perceptual schemas specify how objects covary. For example, ovens and refrigerators are typically found in kitchens, barns and

Address correspondence to Marvin M. Chun, Department of Psychology, Yale University, P.O. Box 208205, New Haven, CT 06520-8205; e-mail: marvin.chun@yale.edu.

chickens on farms, desks and computers in offices. This information is acquired over a lifetime of visual experience, based on detection of invariant covariation between related objects. Experiment 1 introduced a new paradigm to examine how covariation between visual objects may be learned and how these associations may support top-down expectancies.

We asked subjects to search for novel target objects presented in visual arrays of other novel distractors. We wished to examine whether the shape of targets could be cued, or primed, by the set of distractor shapes that covaried with the presentation of the target over the course of visual experience (the experimental session). The question was, can observers learn the association between novel target shapes and novel distractor shapes? We also examined whether such learning can occur without explicit instruction or intention.

Method

A sample search array is shown in Figure 1. The target of each trial was defined as the single object whose shape was symmetric around the vertical axis. This criterion allowed us to define a target for the search task without specifying the actual shape of the target. Each distractor was symmetric around an axis that was oriented away from the vertical axis.

Ninety-six different novel objects were generated. Sixteen of the objects were symmetric around the vertical axis (0°) and were used as targets. The remaining 80 objects were used as distractors. The distractors formed five groups of 16 different objects each; these five groups were symmetric around the 30°, 60°, 90°, 120°, and 150° axis of global orientation, respectively. The objects were colored white and presented on a gray background. For each search trial, 11 items (1 target, 10 distractors) were randomly arrayed in an invisible 8 × 6 grid that subtended 40° × 30° of visual angle. The 10 distractors formed a distractor set, which consisted of five different distractor types (shapes); each type was randomly replicated zero to two times to produce a total of 10 distractors on each trial. The location of the target and the configuration of distractor items were randomized for every trial, in contrast to the fixed target locations and configurations used in previous studies of contextual cuing (Chun & Jiang, 1998).

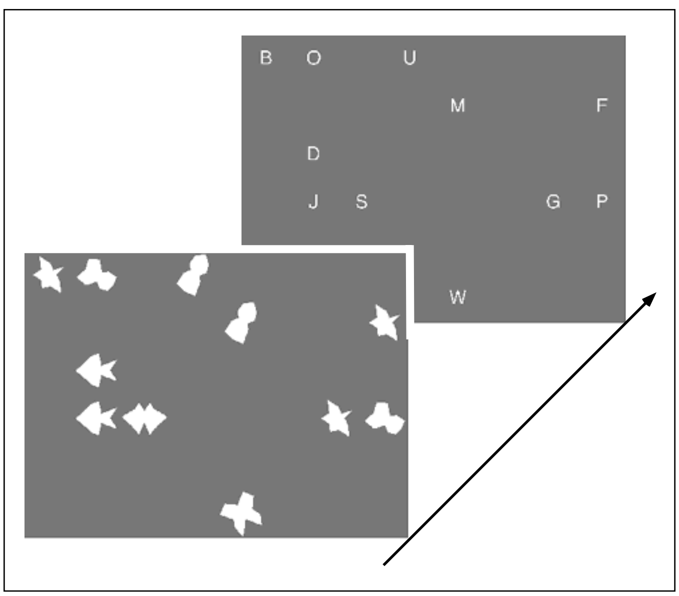

Fig. 1. Schematic display of a search trial in Experiment 1. The target was the single object that was symmetric around the vertical axis. Immediately upon detection, subjects pressed a space bar, reaction time was recorded, and the display was replaced with a probe array of letters. Subjects entered the probe letter that appeared in the target location, providing an accuracy measure of correct target localization.

The task was to search the display and detect the single object that was symmetric around the vertical axis. Subjects were asked to do this as quickly as possible, pressing the space bar immediately upon detection. Once they pressed the space bar, the screen of novel objects was erased and replaced by an array of 11 letters (drawn from a randomized set of 25 alphabet letters, excluding *I*), each occupying the location in which an object had been presented in the previous frame. Subjects were required to type in the identity of the letter that occupied the location that the target object had appeared in. This probe task ensured that subjects correctly localized the target before pressing the space bar. Subjects were told that the probe response was not speeded. Immediate error feedback was provided in the form of auditory tones.

A 6 × 2 design was used. One factor was epoch. The experiment included six epochs of trials, and each epoch comprised four blocks. The epoch factor was crossed with two mapping conditions (consistent vs. variable).[1] In the consistent-mapping (i.e., old) condition, each of the targets was paired with a distractor set, and this pairing was preserved throughout the experiment (repeated across blocks). In the variable-mapping (i.e., new) condition, the pairings between target and distractor sets were randomized from block to block. Each block contained eight consistent-mapping and eight variable-mapping trials. Eight of the 16 target objects were randomly assigned (separately for each subject) to the consistent-mapping condition, and the other 8 were assigned to the variable-mapping condition. In addition, eight distractor sets were assigned to each mapping condition.

Seventeen Yale University subjects with normal or corrected-to-normal visual acuity participated in a 1-hr session for pay or course credit. Subjects were simply instructed to perform the search task, and there was no indication that they should try to encode or learn the displays in any manner.

Results and Discussion

One subject's data were not included in the analysis because of low accuracy (82%). For the other 16 subjects, reaction times (RTs) below 200 ms and above 5,000 ms were removed from the data. Less than 1.3% of the data was omitted as a result of this procedure. Mean accuracy averaged above 98% correct and did not differ between the consistent- and variable-mapping conditions (all Fs < 1).

The mean RT data are shown in Figure 2. The main result was that search performance was faster for the consistent-mapping condition than for the variable-mapping condition, $F(1, 15) = 10.98$, $p < .005$.

1. This manipulation corresponds to the well-established distinction between consistent and variable mapping, introduced by Schneider and Shiffrin (1977; Shiffrin & Schneider, 1977). One difference is that our variable-mapping condition did not involve a mapping reversal between target and distractor shapes. And although consistent mapping is proposed to facilitate learning and performance, we do not claim that this will necessarily lead to complete automaticity in our paradigm.

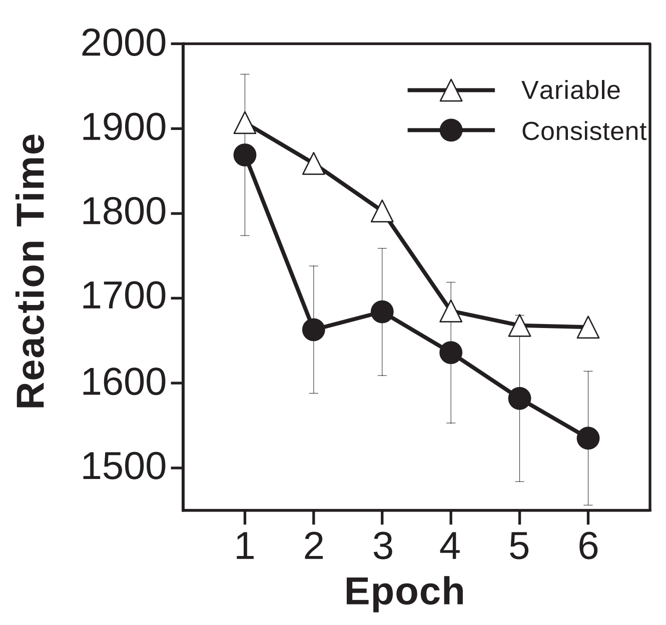

Fig. 2. Mean correct response time for target detection in the variable-mapping and consistent-mapping conditions in Experiment 1, as a function of epoch. The error bars represent the *SEM* ($N = 16$).

This result shows that subjects learned the associations between distractor identities and target identities, so that sensitivity to the distractor context facilitated search performance. There was also a main effect of epoch, $F(5, 75) = 14.79$, $p < .001$. Although the overall interaction between mapping condition and epoch was not significant, the interaction was significant when restricted to the first two epochs, $F(1, 15) = 7.74$, $p < .02$. Thus, the difference between mapping conditions reflects a learning effect that becomes significant after one epoch (four repetitions of each pairing of a target and distractor set).

These results show that subjects encoded the context of targets and used this information to guide object recognition and search. Did conscious recognition of the distractor shapes and awareness of the associative target-distractor pairings drive this learning? To answer this question, we looked at the results of an explicit recognition test administered at the end of the experiment. We informed the subjects that some targets were paired with the same set of distractors throughout the experiment. When asked if they had noticed this, 4 of the 16 subjects reported that they had. None of the subjects reported trying to memorize the object identities or associative pairings, however. In addition, we administered a forced-choice discrimination test. This test consisted of one block of trials that presented a randomly ordered set of eight consistent-mapping displays and eight random-mapping displays, one at a time. Instead of performing the search task, subjects were asked to choose whether each display contained an "old" or "new" identity association. All subjects performed at chance in this task ($M = 48\%$), and hit rates ($M = 17\%$) and false positive rates ($M = 19\%$) did not differ ($p > .49$). The 4 subjects who reported awareness of the repeated pairings also performed at chance ($M = 48\%$).

These findings demonstrate that observers implicitly learn covariation between novel visual objects and are consistent with a previous study that showed learning of co-occurrences between pairs of word stimuli (Logan & Etherton, 1994). In our experiment, subjects were sensitive to the context of target events, and learning of target-context associations occurred when the identities of distractors in the context covaried with the identity of the target. This learning facilitated search performance, a process we call contextual cuing (Chun & Jiang, 1998). In contrast to our 1998 study, search facilitation in this experiment occurred for targets appearing in unpredictable locations, suggesting that context was cuing the identity of the targets.[2] These results point to a covariation learning mechanism that may be fundamental for constructing perceptual schemas that guide top-down expectancies. A multitude of such memory traces may be established through the course of visual experience, and we hypothesize that these memory traces are specific and instance-based (Logan, 1988), allowing for fine-tuned matching (Chun & Jiang, 1998). Top-down knowledge is based on such memory traces for the set of objects a target object typically covaries with.

It is worth noting that visual learning occurred implicitly. Hence, these results demonstrate a new form of implicit learning for encoding covariation between novel visual objects. Existing work has demonstrated implicit covariation learning between visual attributes and other variables such as personality or target location (Lewicki, 1986a, 1986b). Going beyond these interesting findings, our results show robust implicit learning of covariation between novel visual objects. Such covariation learning forms a useful top-down mechanism for visual perception. In the General Discussion, we elaborate on why implicit learning is useful and important.

EXPERIMENT 2

A wide variety of behavioral situations suggest that it is crucial to be able to predict how objects in the environment move and change over time. For instance, safe driving relies on everyone's ability to perceive and predict the movement of other cars. Also, consider what a football quarterback views when trying to find a receiver weaving through a violent sea of other moving players. Predictability of how objects move about, and sensitivity to this predictability, is critical in these and a wide variety of additional contexts. Yet the ability to track multiple moving objects is severely limited in capacity and requires attention (He, Cavanagh, & Intrilligator, 1996; Pylyshyn & Storm, 1988; Yantis, 1992). Top-down knowledge of regularities in the dynamic environment may help overcome the taxing demands of tracking multiple objects.

Experiment 2 examined whether dynamic regularities can be learned and how they may guide search behavior. We employed dynamic visual displays to mimic the dynamic visual world. Subjects performed visual search for a rotated-T target among L-shaped distractors. All of these search items moved independently around the screen, under the constraint that they were not allowed to run into each other or disappear off the screen. The animation sequences were similar to those used in tracking studies (He et al., 1996; Pylyshyn & Storm, 1988; Yantis, 1992).

2. Note that perceptual schemas also specify the locations in which objects can typically appear relative to each other (Biederman et al., 1982). We are currently investigating whether observers can learn both identity and location constraints. However, for present purposes, we focus on identity cuing to generalize contextual cuing beyond configurational cuing of target location (Chun & Jiang, 1998).

Method

Figure 3 shows a schematic trial.[3] Each trial consisted of an animation sequence of eight independently moving objects. One of these items was designated as the target, a T shape pointing left or right, and the subjects were instructed to locate the target and press a button corresponding to whether it was pointing left or right. The orientation of the target remained constant as it moved. The rest of the items were L-shape distractors, each rotated randomly in one of four directions. A target was present on every trial. The white search items (each 1.4° × 1.4° of visual angle in size) moved around a gray field that subtended approximately 40° × 30° of visual angle. The motion sequences were generated using apparent motion of 15 frames presented at a rate of 8 frames per second. All of the items moved at a constant frame-to-frame velocity of about 14° of visual angle per second. The direction of each item's movement from frame to frame was constrained to occur within ±30° of the present trajectory. Trajectories were allowed to change more abruptly at the invisible boundaries of the window or in collision paths with other items.

Each trial started with a small fixation dot appearing in the middle of a computer screen. After a pause of 500 ms, the array of stimuli appeared on the screen. The initial configuration of items was equated for the two conditions to prevent cuing based on this initial array. Also, the targets and distractors were indistinguishable from each other in the first 360 ms of each trial to allow enough time for the motion sequences to evolve in a unique manner. Hence, every item in the first frame appeared as a cross, and over the course of four frames (the first 480 ms), each of the items slowly morphed into either a T target or an L distractor. Subjects monitored each animation sequence for the target and pressed one of two response keys corresponding to the identity of the localized target. Error feedback was given in the form of auditory tones. The target identity (left-facing or right-facing) and corresponding response were randomized for each trial so that they were not correlated with any of the displays. The motion sequence stopped after 2 s, and the final configuration of items remained on the screen until a response was made. The prolonged presence of the last frame did not affect our results because motion contextual cuing was established for RTs under 2 s, while the items were still in motion.

As in Experiment 1, the two main variables were condition (consistent vs. variable) and epoch (1–6). Each epoch contained three blocks. The consistent set of stimuli consisted of randomly generated motion sequences that were repeated throughout the entire experiment, once per block. Embedded target trajectories were fully correlated within each of these invariant global contexts of distractor trajectories. The variable set of stimuli consisted of motion sequences that were newly generated for each block and hence were not correlated with the target trajectories. Because sensitivity to global invariance would cue the target trajectory in the consistent condition, we predicted faster search performance in that condition than in the control baseline condition (variable condition). Each block contained an intermixed set of six different consistent-mapping trials and six different variable-mapping trials. The target trajectories used in the variable-mapping displays were repeated across blocks, controlling for learning of target trajectory per se. Hence, the consistent and variable conditions differed only by their distractor motion trajectories.

3. A dynamic version of a sample trial is available on the World Wide Web at http://pantheon.yale.edu/~yj23/DynamicSearchDemo.html.

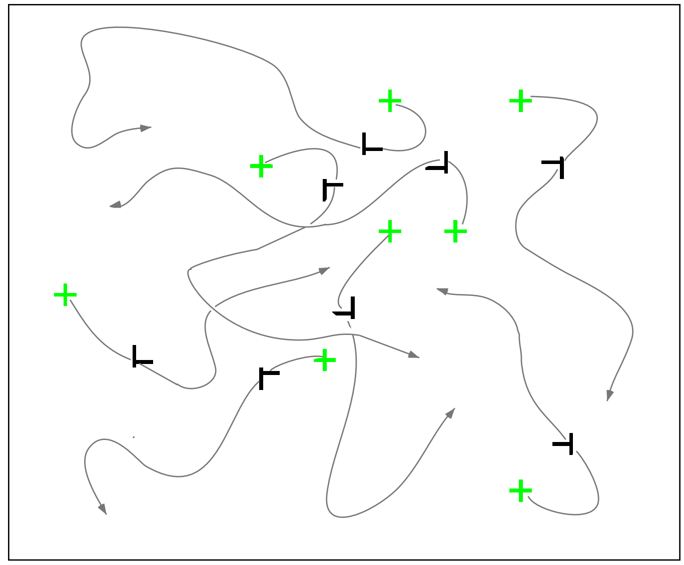

Fig. 3. Sample animation sequence used in Experiment 2. Each trial began with a fixed configuration of crosses (shown in light gray), which evolved into target and distractor items by Frame 4 (shown in black) of the animation sequence. This allowed enough time for the motion sequences to emerge in a unique manner from a fixed starting configuration held constant across conditions. Subjects performed visual search for a rotated T among rotated Ls. All of the objects moved independently around the screen following the schematic motion trajectories illustrated by the arrows. A dynamic demonstration of a sample trial is available on the World Wide Web at http://pantheon.yale.edu/~yj23/DynamicSearchDemo.html.

Twelve subjects participated. As before, subjects were simply instructed to perform the search task as quickly and as accurately as possible.

Results

Accuracy averaged above 95% correct and did not differ between the consistent and variable conditions (all Fs < 1). Mean search performance for targets was significantly faster in consistent displays than in the variable condition, $F(1, 11) = 9.99$, $p < .01$. Figure 4 plots search RTs for targets in the two conditions as a function of epoch. The benefit for consistent displays was significant after just three repetitions (122-ms difference in Epoch 2, $p < .05$, one-tailed t test). Thus, the results indicate that subjects were sensitive to invariant regularities in the complex, dynamic displays, and that these invariants guided attention to embedded targets, facilitating search.

To rule out the possibility that improved search performance was based on explicit recognition of the displays, we queried each subject about the repetition manipulation and administered an explicit recognition test, as in Experiment 1. Only 1 of the 12 subjects reported noticing the repetition manipulation, and no one reported trying to encode the displays explicitly. Moreover, explicit recognition accuracy was at chance levels (overall accuracy = 49%, hit rate = 31%, false alarm rate = 33%). Hence, as in Experiment 1, learning occurred without intent to encode the displays, and the resulting implicit representations of context facilitated search without supporting conscious recognition.

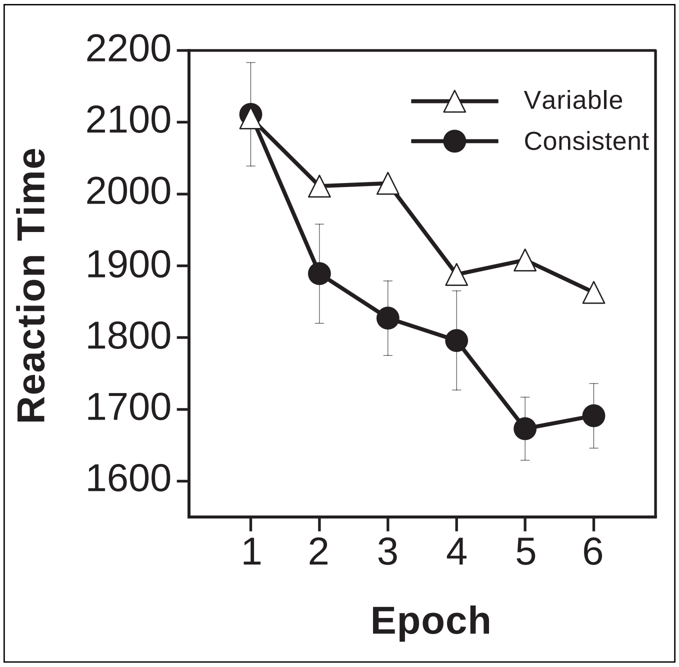

Fig. 4. Mean correct response time for target detection in the variable and consistent conditions of Experiment 2, plotted as a function of epoch. The error bars represent the *SEM* ($N = 12$).

This is a remarkable, new form of implicit learning given that the dynamic contexts were defined by complex motion trajectories of multiple distractor events, and subjects were not informed of any contingencies between targets and distractors. Previous work on implicit procedural learning demonstrated benefits for making motor movements to single target events (lights) that appeared in a fixed sequence of different locations (Lewicki, Hill, & Bizot, 1988; Nissen & Bullemer, 1987) or to single targets that moved around the screen (Frith & Lang, 1979). These tasks involved both perceptual and motor components.[4] In contrast, our task minimizes contributions of motor-response learning, while illustrating a new form of perceptual learning that facilitates search performance on the basis of covariation between multiple perceptual motion trajectories, including those of distractor events. Put more simply, whereas the implicit-procedural-learning studies are relevant for understanding how people become proficient at playing the piano or typing on a keyboard, our demonstration illustrates how people use dynamic visual information to guide behaviors such as driving or selecting targets (e.g., as in determining where to pass the ball amidst a field of moving players in a team sport).

GENERAL DISCUSSION

Top-down knowledge places important constraints on visual processing. This study illustrates how useful top-down knowledge is extracted from the visual input. Specifically, observers are sensitive to regularities and covariation that are invariant (repeated) over time. Perceptual covariation learning is supported by implicit memory traces of perceptual interactions that accumulate during the course of visual experience (Chun & Jiang, 1998; Logan, 1988).

The ability to pick up and encode regularities in the visual environment is of great ecological significance. As Gibson (1969) has long argued, this ability serves to reduce complexity and increase predictability. In Experiment 1, object recognition and search were facilitated by implicit learning of associations between novel distractor sets and novel target shapes, much as an oven may cue the presence of a toaster in everyday life. In Experiment 2, regularities in dynamic environments also cued search, as a global field of moving basketball players constrains how to move and where to pass the ball. In sum, useful top-down perceptual knowledge is routinely derived from structured visual experience. Much of this sophisticated learning appears to occur in an implicit manner.

The implicit nature of learning and memory is highly useful for visual perception (Chun & Jiang, 1998). Implicit operations are durable over time, are robust across interference, and perhaps most important, exhibit high capacity (Reber, 1989; Seger, 1994), releasing limited-capacity cognitive processes so that cognitive resources become available for other processes that may require conscious mediation. Indeed, sophisticated, unconscious mechanisms are proposed to have a broad and profound influence on perception and cognition (Berry & Dienes, 1993; Kihlstrom, 1987; Stadler & Frensch, 1998). For instance, Lewicki and his colleagues (Lewicki, 1986a, 1986b; Lewicki, Hill, & Czyzewska, 1992) have argued that complex covariant information in the environment is typically encoded in an implicit, unconscious manner. Moreover, such implicit mechanisms are potentially available very early in life to guide perceptual development (Gibson, 1969).

We propose that such implicit learning occurs whenever invariant information in the environment is predictive and informative for a behavioral task or situation. In other words, reinforcement occurs only for task-relevant contextual cues that reduce uncertainty and facilitate behavior (Ahissar et al., 1992; Gibson, 1969; Thorndike, 1911/1965). Hence, learning was exhibited in our task because the visual contexts were correlated with target shape or motion trajectory. An alternative hypothesis is that subjects were learning to search through repeated displays more quickly because of low-level repetition priming (Bar & Biederman, 1998; Tulving & Schacter, 1990). This latter account predicts facilitation from repetition (which primes early perceptual mechanisms), independent of whether targets are correlated with invariant contexts. However, note that target shapes and distractor sets were repeated the same number of times in the variable and consistent conditions in Experiment 1, yet contextual cuing was obtained only when targets and contexts were meaningfully correlated. These results indicate a role for associative learning rather than low-level perceptual priming in contextual cuing. Although such a direct comparison is not available in Experiment 2, unpublished data from our lab indicate that no contextual cuing was obtained when target trajectories were decorrelated from their otherwise invariant contexts. Thus, contextual cuing is not produced by repetition per se, but rather is driven by meaningful environmental covariation.[5]

4. Implicit-learning studies have examined separable contributions of perceptual learning versus motor learning in serial RT tasks (Howard, Mutter, & Howard, 1992; Mayr, 1996). However, these effects were still restricted to learning of single-target sequences toward which a response was performed or explicit awareness was engaged.

5. Low-level perceptual priming and procedural learning do contribute to the overall level of performance, of course. This is indicated by the global improvement in the baseline variable-mapping condition and is dissociable from the covariation learning that produces contextual cuing (the additional benefit for the consistent-mapping condition).

In closing, the present results highlight the importance of implicit learning and memory mechanisms in perception and cognition. Complex covariation in the environment can be learned to guide behavior because the informational structure in the environment helps reduce uncertainty. Learning mechanisms bring such useful top-down knowledge of the visual world into the mind, allowing perceivers to benefit from visual experience.

Acknowledgments—This research was supported in part by Grant BCS-9817349 from the National Science Foundation. We thank Ron Rensink and Carol Seger for their helpful comments on an earlier version of this manuscript.

REFERENCES

Ahissar, E., Vaadia, E., Ahissar, M., Bergman, H., Arieli, A., & Abeles, M. (1992). Dependence of cortical plasticity on correlated activity of single neurons and on behavioral context. *Science, 257,* 1412–1415.

Bar, M., & Biederman, I. (1998). Subliminal visual priming. *Psychological Science, 9,* 464–469.

Berry, D.C., & Dienes, Z. (1993). *Implicit learning.* East Sussex, England: Erlbaum.

Biederman, I. (1972). Perceiving real-world scenes. *Science, 177,* 77–80.

Biederman, I., Mezzanotte, R.J., & Rabinowitz, J.C. (1982). Scene perception: Detecting and judging objects undergoing relational violations. *Cognitive Psychology, 14*(2), 143–177.

Chun, M.M., & Jiang, Y. (1998). Contextual cueing: Implicit learning and memory of visual context guides spatial attention. *Cognitive Psychology, 36,* 28–71.

Duncan, J., & Humphreys, G.W. (1989). Visual search and stimulus similarity. *Psychological Review, 96,* 433–458.

Frith, C.D., & Lang, R.J. (1979). Learning and reminiscence as a function of target predictability in a two-dimensional tracking task. *Quarterly Journal of Experimental Psychology, 31,* 103–109.

Garner, W.R. (1974). *The processing of information and structure.* Potomac, MD: Erlbaum.

Gibson, E.J. (1969). *Principles of perceptual learning and development.* New York: Appleton-Century-Crofts.

Gibson, E.J. (1991). Perceptual development and the reduction of uncertainty. In *An odyssey in learning and perception* (pp. 353–363). Cambridge, MA: MIT Press.

He, S., Cavanagh, P., & Intrilligator, J. (1996). Attentional resolution and the locus of visual awareness. *Nature, 383,* 334–337.

Hollingworth, A., & Henderson, J. (1998). Does consistent scene context facilitate object perception? *Journal of Experimental Psychology: General, 127,* 398–415.

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,* 1029–1039.

Kihlstrom, J.F. (1987). The cognitive unconscious. *Science, 237,* 1445–1452.

Lewicki, P. (1986a). *Nonconscious social information processing.* New York: Academic Press.

Lewicki, P. (1986b). Processing information about covariations that cannot be articulated. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 12,* 135–146.

Lewicki, P., Hill, T., & Bizot, E. (1988). Acquisition of procedural knowledge about a pattern of stimuli that cannot be articulated. *Cognitive Psychology, 20,* 24–37.

Lewicki, P., Hill, T., & Czyzewska, M. (1992). Nonconscious acquisition of information. *American Psychologist, 47,* 796–801.

Logan, G.D. (1988). Towards an instance theory of automatization. *Psychological Review, 95,* 492–527.

Logan, G.D., & Etherton, J.L. (1994). What is learned during automatization? The role of attention in constructing an instance. *Journal of Experimental Psychology: Human Perception and Performance, 20,* 1022–1050.

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.

Nissen, M.J., & Bullemer, P. (1987). Attentional requirements of learning: Evidence from performance measures. *Cognitive Psychology, 19,* 1–32.

Palmer, S.E. (1975). The effects of contextual scenes on the identification of objects. *Memory & Cognition,* 3, 519–526.

Pylyshyn, Z.W., & Storm, R.W. (1988). Tracking multiple independent targets: Evidence for a parallel tracking mechanism. *Spatial Vision, 3,* 179–197.

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

Rensink, R.A., O'Regan, J.K., & Clark, J.J. (1997). To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science, 8,* 368–373.

Schneider, W., & Shiffrin, R.M. (1977). Controlled and automatic human information processing: I. Detection, search and attention. *Psychological Review, 84,* 1–66.

Seger, C.A. (1994). Implicit learning. *Psychological Bulletin, 115,* 163–196.

Shiffrin, R.M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending and a general theory. *Psychological Review, 84,* 127–190.

Stadler, M.A., & Frensch, P.A. (Eds.). (1998). *Handbook of implicit learning.* Thousand Oaks, CA: Sage.

Thorndike, E.L. (1965). *Animal intelligence.* New York: Hafner. (Original work published 1911)

Treisman, A.M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology, 12,* 97–136.

Tulving, E., & Schacter, D.L. (1990). Priming and human memory systems. *Science, 247,* 301–306.

Wolfe, J.M., Cave, K.R., & Franzel, S.L. (1989). Guided search: An alternative to the feature integration model for visual search. *Journal of Experimental Psychology: Human Perception and Performance, 15,* 419–433.

Yantis, S. (1992). Multielement visual tracking: Attention and perceptual organization. *Cognitive Psychology, 24,* 295–340.

(RECEIVED 6/1/98; REVISION ACCEPTED 1/5/99)