Two cases requiring external reinforcement in perceptual learning

Aaron R. Seitz
Department of Psychology, Boston University,
Boston, MA, USA

Jose E. Nanez, Sr.
Department of Social and Behavioral Sciences,
Arizona State University, Phoenix, AZ, USA

Steve Holloway
Department of Social and Behavioral Sciences,
Arizona State University, Phoenix, AZ, USA

Yoshiaki Tsushima
Department of Psychology, Boston University,
Boston, MA, USA

Takeo Watanabe
Department of Psychology, Boston University,
Boston, MA, USA

The role of external reinforcement is an issue of much debate and uncertainty in perceptual learning research. Although it is commonly acknowledged that external reinforcement, such as performance feedback, can aid in perceptual learning (M. H. Herzog & M. Fahle, 1997), there are many examples in which it is not required (K. Ball & R. Sekuler, 1987; M. Fahle, S. Edelman, & T. Poggio, 1995; A. Karni & D. Sagi, 1991; S. P. McKee & G. Westheimer, 1978; L. P. Shiu & H. Pashler, 1992). Additionally, learning without external reinforcement can occur even for stimuli that are irrelevant to the subject's task (A. R. Seitz & T. Watanabe, 2003). It has been thus hypothesized that internal reinforcement can serve a similar role as external reinforcement in learning (M. H. Herzog & M. Fahle, 1998; A. Seitz & T. Watanabe, 2005). This idea suggests that perceptual learning should occur in the absence of external reinforcement provided that easy exemplars are utilized as a basis for the subject to generate internal reinforcement. Here, we report results from two studies that show that this is not always the case. In the first study, subjects participated in two sessions of a motion direction discrimination task with low-contrast dots moving in directions separated by 90°. In the second study, subjects participated in 12 orientation-discrimination sessions using oriented bars (oriented either 70° or 110°) that were masked by spatial noise. Trials of different signal levels (yielding psychometric functions ranging from chance to ceiling) were randomly interleaved. In both studies, subjects experiencing external reinforcement showed significant learning, whereas subjects receiving no external reinforcement failed to show learning. We conclude that while internal reinforcement is an important learning signal, the presence of easy exemplars is not sufficient to generate reinforcement signals.

Keywords: perceptual learning, internal reinforcement, external reinforcement, feedback, motion discrimination, orientation discrimination

Introduction

Perceptual learning refers to changes in sensory abilities that occur through training and is thought to be an important process that helps us adapt to our ever-changing physical environment. A fundamental aspect of perceptual learning is that it is subserved by plasticity in sensory processing stages and is, by and large, a form of implicit learning. Thus, although there is a wide body of evidence showing that external reinforcement, such as response feedback, results in increased performance and learning rates during many types of learning (Goldstein & Rittenhouse, 1954; Goodman & Wood, 2004; Pavlov, Gantt, & Folbort, 1928; Schultz, 2000), the role of external reinforcement is still unclear in perceptual learning. For instance, while some studies indicate a clear role of external reinforcement in perceptual learning (Herzog & Fahle, 1997), many studies find that perceptual learning occurs robustly in the absence of external reinforcement (Ball & Sekuler, 1987; Fahle, Edelman, & Poggio, 1995; Karni & Sagi, 1991; McKee & Westheimer, 1978; Shiu & Pashler, 1992). Furthermore, recent research has shown that perceptual learning of optical motion can occur as a result of mere exposure to a subliminal stimulus, without external reinforcement, without the subject actively attending to the motion stimulus, or with the motion stimulus being a relevant feature of the particular task (Seitz & Watanabe, 2003; Watanabe, Nanez, & Sasaki, 2001). These studies raise the question concerning under what conditions feedback may be required to produce perceptual learning.

doi: 10.1167/6.9.9 Received April 26, 2006; published August 22, 2006 ISSN 1534-7362

This question is not well answered in the literature given that the presence (or absence) of external reinforcement is typically not of primary interest in perceptual learning studies. Also, experiments in which learning fails to manifest often go unreported. Therefore, there exist only a few studies that explicitly use training procedures both with and without external reinforcement. For instance, Herzog and Fahle (1997) found that learning occurred more consistently with feedback, even with partial or blocked feedback, than without feedback, but the results are difficult to interpret given that many subjects in the no-feedback condition showed strong learning and that, on average, no-feedback subjects started off with better performance than those in the feedback conditions. More generally, while external reinforcement has been shown in some cases to be required for within-session learning of texture segmentation (Karni & Sagi, 1991) and contour detection (Shiu & Pashler, 1992), the same studies also found across-session learning (i.e., learning effects evaluated by comparing performance across sessions conducted on distinct days) in the absence of external reinforcement. On the other hand, studies of perceptual learning for hyperacuity found that even within-session learning occurs without feedback (Fahle et al., 1995).

One explanation for the observation that perceptual learning can occur in the absence of external reinforcement is that internal reinforcement can serve as a learning signal (Fahle & Edelman, 1993). For instance, stimuli that are highly discernable can serve as a template that subjects can use to assess stimuli in more difficult conditions. In fact, in some circumstances, allowing subjects to study an easily identifiable stimulus can enable learning in stimulus conditions for which learning does not otherwise occur (Ahissar & Hochstein, 1997). These results would suggest that perceptual learning should occur in training procedures involving a mixture of stimulus difficulties even without external reinforcement.

Here, we report results from two studies that found no learning in the absence of external reinforcement in tasks where easy exemplars were presented consistently throughout training. In both studies, subjects receiving no external reinforcement failed to show any learning, whereas subjects who were given trial-by-trial feedback during training showed significant learning, which was consistent across all signal levels. In the first study, subjects were trained for 1 day to report the direction of motion (of four oblique directions separated by 90°) of a patch of 100% coherent but low-contrast moving dots, and they were tested on a subsequent day in the absence of external reinforcement. In the second study, subjects were trained for 10 days to report the orientation of bars (oriented either 70° or 110°) that were masked by spatial noise and were tested in the absence of external reinforcement before and after training. In both studies, trials of many different signal levels (yielding a psychometric function that ranged from chance to ceiling) were randomly interleaved.

Experiment 1

In this experiment, we investigated the role of feedback in a task where subjects were required to discriminate the direction of low luminance-contrast motion stimuli. In this experiment, subjects in each condition underwent a training session on Day 1, and possible performance improvements were measured on Day 2. This particular task was selected because task-irrelevant perceptual learning has been found in a previous study using the same stimulus set used here (Seitz, Nanez, Holloway, Koyama, & Watanabe, 2005). We hypothesized that if task-irrelevant learning occurs for these stimuli, then learning based on internal reinforcement (i.e., in the absence of external reinforcement) would likely ensue from a training procedure involving the presentation of easy exemplars.

Methods

Twelve participants (age, 19–35 years) were recruited from the Phoenix metropolitan area. Subjects were randomly and evenly assigned between the external reinforcement (ER) and the no-external reinforcement (NoER) conditions. All subjects reported good ocular health and had a best corrected visual acuity (tested on-site) of 20/40 Snellen or better. Additionally, all participants were naive as to the purpose of the study. Stimuli were presented on a 19-in. cathode ray tube monitor with a resolution of 1,152 × 768 pixels and a refresh rate of 75 Hz using custom software written for a Macintosh G4 computer.

In the direction-discrimination sessions, subjects reported the off-cardinal direction (70°, 160°, 250°, or 340°) of the motion display consisting of 200 coherently moving dots that were displayed in an annulus (1° inner diameter, 10° outer diameter) at varying luminance-contrast levels (signal levels). Contrast was varied randomly, beginning below the subject's threshold for detection and discrimination of the stimulus and increasing to suprathreshold levels (0, 0.14, 0.2, 0.28, 0.42, 0.6, 0.9, 1, 1.9, 11.8 cd/m^2 RMS contrast). Subjects viewed a 500-ms stimulus presentation and were asked to report the direction of dot motion by choosing an appropriate directional arrow (for details, see Seitz, Nanez, et al., 2005). Each session lasted approximately 1 hr and consisted of 20 trials at each of 4 directions × 10 signal levels, for a total of 800 trials.

In some conditions, subjects received feedback after each trial regarding whether or not their response was correct: a green "+" symbol coupled with a high-pitched tone for correct responses or a red "x" coupled with a low-pitched tone for incorrect responses.

Results

In the NoER condition, subjects underwent 2 sequential days of training in the absence of any external reinforcement regarding their performance. In the ER condition, subjects underwent the same procedure as in the NoER condition but were given response feedback during the first session. In the second session, neither group received external reinforcement.

Seven subjects in the NoER condition were trained on Day 1 in the absence of external feedback and were tested on Day 2 in the same manner. As can be seen in Figure 1, there was no notable change of performance across sessions, $F(1,119) = 0.2$, $p = .67$, two-way (Signal Level × Day) ANOVA with repeated measures, despite the fact that subjects experienced a wide range of stimulus difficulties, and there was a significant effect of signal level, $F(9,119) = 50.5$, $p < .001$. There was no interaction between session and signal level, $F(9,119) = 0.9$, $p = .53$.

A different group of seven subjects participated in the ER condition. These subjects were trained on Day 1 with external reinforcement consisting of trial-by-trial response feedback. On Day 2, they were tested in the absence of external feedback. As can be seen in Figure 2, there was a significant change in performance across sessions, $F(1,119) = 14.6$, $p < .01$, two-way (Signal Level × Day) ANOVA with repeated measures, a significant effect of signal level, $F(9,119) = 62.3$, $p < .001$, and a small interaction between session and signal level, $F(9,119) = 2.2$, $p < .05$.

Discussion

The results of Experiment 1 indicate that a single training session on a direction-discrimination task in the presence of external reinforcement (i.e., feedback) produces significant learning, whereas conducting the same training in the absence of external reinforcement fails to produce learning.

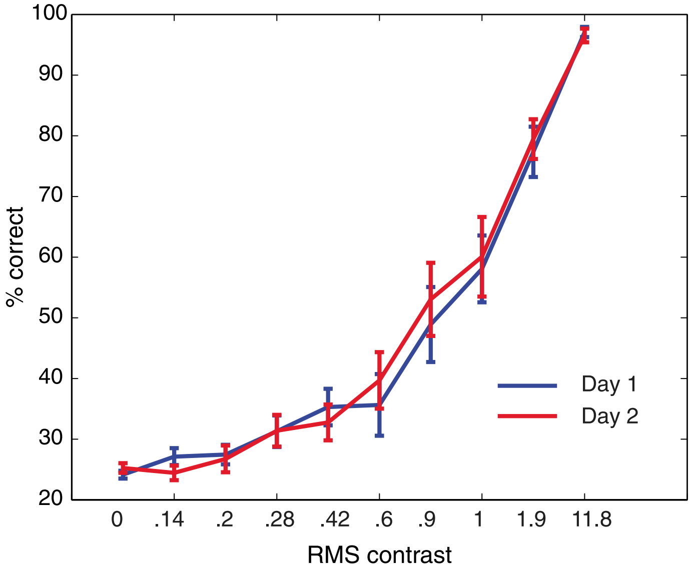

Figure 1. Experiment 1, NoER condition. Subjects receiving no external reinforcement showed no notable performance changes between the first session (blue) and the second session (red). Error bars reflect standard error.

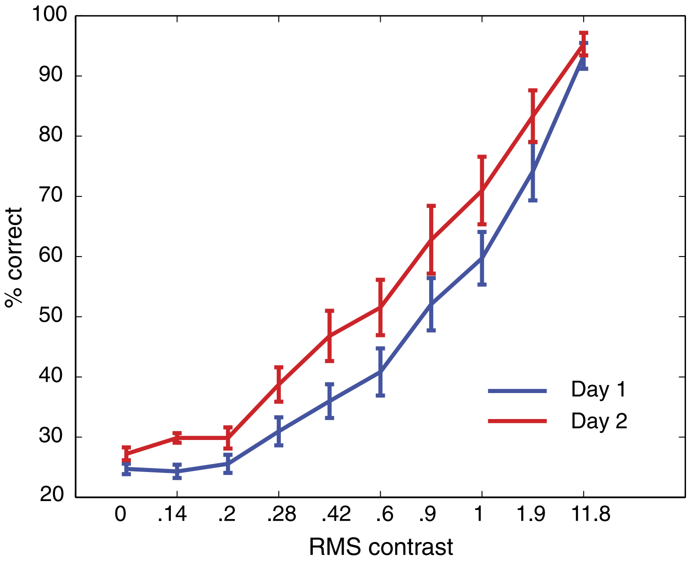

Figure 2. Experiment 1, ER condition. Subjects receiving external reinforcement showed performance improvements between the first session (blue) and the second session (red). Error bars reflect standard error.

Although this demonstrates that internal reinforcement alone is not as efficient for learning as is external reinforcement, there is still a possibility that learning based on internal reinforcement is slow and requires more than a single day of training. For instance, Ball and Sekuler (1987) found that while learning of fine direction discrimination occurred in the absence of external reinforcement, they also found that in the case of oblique directions (such as those used in this study), multiple days of training were required to detect significant learning effects. In general, it has been found that low-level perceptual learning, on tasks such as direction discrimination, often takes many days to manifest even in the presence of external reinforcement (Fine & Jacobs, 2002).

Experiment 2

In this experiment, we investigated the role of feedback in a task where subjects were required to discriminate the orientation of a bar that was masked in spatial noise. Subjects in each condition underwent 10 training sessions and 2 testing sessions so that slowly developing perceptual learning could be measured. While contour discrimination, using similar orientated bars, has been previously found to occur without feedback (Shiu & Pashler, 1992), it has not been studied when the orientation stimuli are masked in noise.

Methods

Eight participants (age, 19–35 years) were recruited from the Boston metropolitan area. Subjects were randomly and

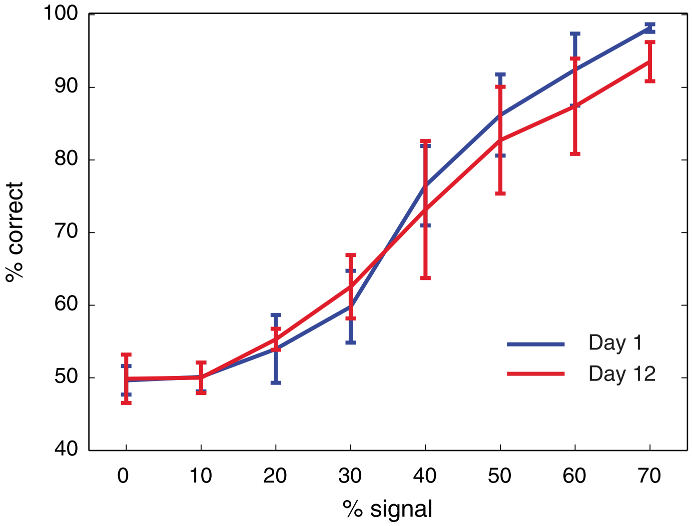

Figure 3. Experiment 2, NoER condition. Subjects without external reinforcement showed no performance improvements between the 1st session (blue) and the 12th session (red). Error bars reflect standard error.

evenly assigned between the ER and the NoER conditions. All subjects reported good ocular health and had a best corrected visual acuity (tested on-site) of 20/40 Snellen or better. Additionally, all participants were naive as to the purpose of the study. Stimuli were presented on a 19-in. cathode ray tube monitor with a resolution of 1,152 × 768 pixels and a refresh rate of 75 Hz using custom software written with use of the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997) for Matlab™ (Natick, MA) on a Macintosh G4 computer.

In the orientation-discrimination sessions, subjects reported the off-cardinal orientation (70° or 110°) of a white (82 cd/m^2) oriented bar (0.75° × 0.15°) presented in a 1° dark-gray (32 cd/m^2) circle in the center of the screen. The bars were spatially masked in noise. Each pixel in the central 1° circle was selected from either a signal stimulus or a noise stimulus, with the percent signal defined as the proportion of pixels in the central 1° circle, including the bar, which were selected from the signal stimulus. For instance, in the case of 60% signal, 60% of the pixels were randomly chosen from stimulus and the other 40% of the pixels were selected from the noise stimulus. The noise stimuli were spatial white-noise patterns, consisting of either white (same color as bar) or dark-gray (same color as background) pixels and were randomly generated for each trial. The presentation of eight different signal levels, ranging from below the subject's discrimination of the stimulus to suprathreshold levels (0%, 10%, 20%, 30%, 40%, 50%, 60%, and 70% signal level), were randomly interleaved. Subjects viewed a 500-ms stimulus presentation and were asked to report the orientation of the stimulus by selecting an appropriate key on the keyboard ("/" or "\"). Each session lasted approximately 45 min and consisted of 40 trials at each of 2 directions × 8 signal levels, for a total of 640 trials.

In the ER condition, subjects received feedback after each trial regarding whether or not their response was correct: a green "Correct" coupled with a high-pitched tone for correct responses or a red "Wrong" coupled with a low-pitched tone for incorrect responses.

In the NoER condition, subjects underwent 10 sequential days of training in the absence of any external reinforcement regarding their performance. In the ER condition, subjects underwent the same procedure as in the NoER condition but received response feedback during the 10 training sessions. Subjects in both groups completed a pretest session before training and a posttest session after training; testing sessions consisted of the same procedure as the training sessions but without feedback in either group.

Results

Four subjects in the NoER condition were trained for 10 days in the absence of external feedback. As can be seen in Figure 3, there was no notable change in performance between the pretest and the posttest sessions, $F(1,63) = 0.30$, $p = .62$, two-way (Signal Level × Testing Day) ANOVA with repeated measures, despite the fact that subjects experienced a wide range of stimulus difficulties. There was a significant effect of signal level, $F(7,63) = 23.1$, $p < .001$, and there was no interaction between session and signal level, $F(7,63) = 0.4$, $p = .89$.

A different set of four subjects participated in the ER condition. These subjects were trained for 10 days with external reinforcement consisting of trial-by-trial response feedback. They were tested before and after the 10 days of training in a session containing no external reinforcement. As can be seen in Figure 4, there was a significant change of performance across sessions, $F(1,63) = 12.9$, $p < .05$,

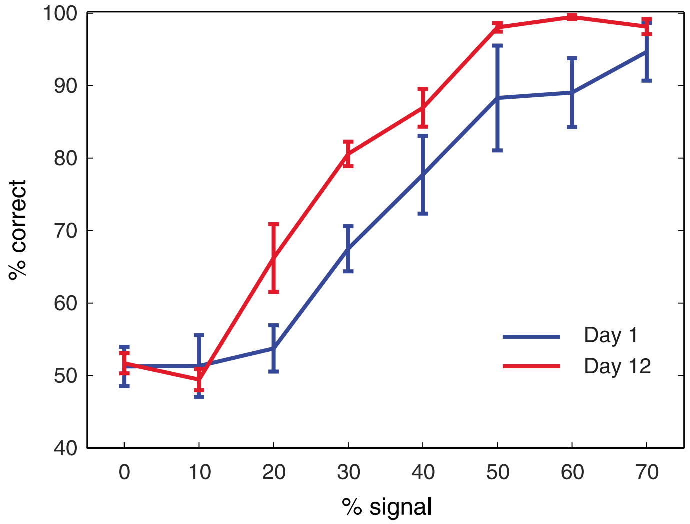

Figure 4. Experiment 2, ER condition. Subjects with external reinforcement showed performance improvements between the 1st session (blue) and the 12th session (red). Error bars reflect standard error.

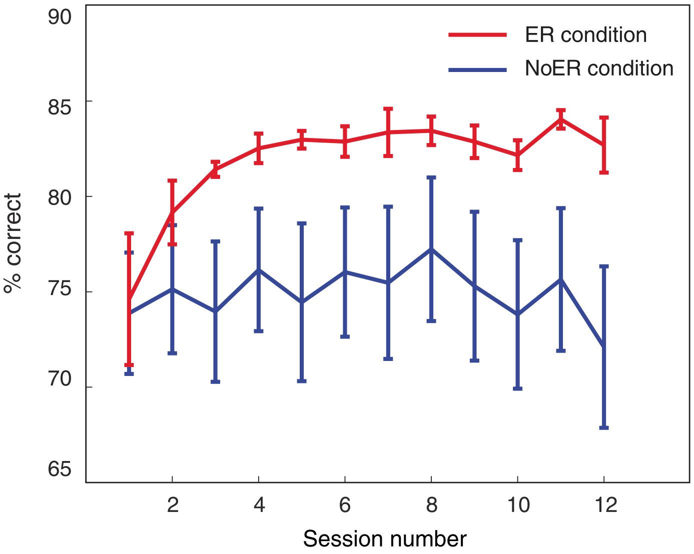

Figure 5. Experiment 2, cross day performance. Subjects in the ER condition showed performance improvements across sessions (red), but subjects in the NoER condition (blue) did not. Error bars reflect standard error.

two-way (Signal Level × Testing Day) ANOVA with repeated measures, a significant effect of signal level, $F(7,63) = 59.4$, $p < .001$, and no interaction, $F(7,63) = 1.8$, $p = .15$.

Figure 5 shows the performance across the 12 sessions for subjects in the ER condition. A significant effect of learning across sessions was found in the ER condition, $F(11,47) = 11.4$, $p < .001$, one-way ANOVA with repeated measures, and there was a significant correlation between training day and performance ($r = .6884$, $p < .01$). On the other hand, no learning was detected for subjects in the NoER condition, $F(11,47) = 1.6$, $p = .14$, one-way ANOVA with repeated measures, and there was no correlation between training day and performance ($r = -.12$, $p = .72$).

Given the negatively accelerating trend in the data from subjects in the ER condition, we fit the data from both groups with power functions ($y = ax^n$). This analysis showed a positive exponent for the ER group ($n = 0.037 \pm 0.015$, 95% confidence; RMSE = 0.013), but the exponent for the NoER group did not significantly deviate from 0 ($n = 0.002 \pm 0.017$, 95% confidence; RMSE = 0.014).

Discussion

The results of Experiment 2 indicate that multiple training sessions on an orientation-discrimination task in the presence of external reinforcement produce significant learning, whereas conducting the same training in the absence of external reinforcement fails to produce learning. Although it is possible that, with further training, perceptual learning would occur in the NoER condition, there was no evidence of this after 12 sessions; in fact, performance was slightly worse in Session 12 than in Session 1.

It is also important to note that subjects in both conditions showed comparable performance in Session 1. Thus, differences in learning cannot be attributed to an initial performance imbalance.

General discussion

The results of Experiments 1 and 2 point toward a failure of internal reinforcement in triggering perceptual learning of motion and orientation discrimination, respectively. An important question is, how are the conditions of these experiments different from others in which learning has been found in the absence of feedback?

One general finding in the literature has been that external reinforcement is often required to find within-session learning but not across-session learning (Karni & Sagi, 1991). An explanation for this is that feedback may help overcome the effects of fatigue, which dissipates across sessions. For instance, Karni and Sagi (1991) found that perceptual learning of texture segmentation occurred in the absence of external reinforcement when measured across sessions, but feedback was required for within-session learning (i.e., “fast learning”) to manifest. Likewise, Shiu and Pashler (1992) found that in an orientation-discrimination task, fast learning only occurred in the presence of external reinforcement, but learning in the absence of external feedback was robust across sessions. However, not all perceptual learning studies support this view; for example, researchers have found learning in the absence of external reinforcement in visual hyperacuity tasks for both fast learning (Fahle et al., 1995) and slow, across-session learning (McKee & Westheimer, 1978). Even so, the results of all these studies would imply that learning should have been found without external reinforcement in the conditions tested here given that both experiments controlled for fatigue effects and that Experiment 2 consisted of 12 sessions—training that should be sufficient enough to detect even slowly developing learning effects.

It should be noted that although the tasks used here resemble those of contour discrimination and motion discrimination, which have been shown to produce perceptual learning in the absence of feedback (Ball & Sekuler, 1987; Shiu & Pashler, 1992), the tasks presented here differ from those used in previous studies in that the angular differences between the motion directions (Experiment 1) and those between the orientations (Experiment 2) were very large. Hence, the primary difficulty encountered in our task is that of detecting the low-signal strength stimuli. In this way, our task resembles those of contrast detection, which has proved to be a case where learning effects are difficult to find even with external reinforcement (Adini, Sagi, & Tsodyks, 2002; Furmanski, Schluppeck, & Engel, 2004; Yu, Klein, & Levi, 2004) and where external reinforcement has been found to be insufficient to produce learning under conditions of signal level uncertainty (Yu et al., 2004).

Studies of task-irrelevant learning show that learning can occur in the absence of external reinforcement even for task-irrelevant, subthreshold stimuli (Seitz & Watanabe, 2003). Seitz and Watanabe (2005) proposed a model to explain both task-irrelevant and task-relevant learning in which task-related signals (either due to external or internal factors) serve to reinforce activity in low-level sensory processing stages in a manner nonspecific to the stimulus. Consistent with this model, robust task-irrelevant learning for low luminance-contrast motion stimuli, which shared the same parameters used in the current study, has been reported (Seitz, Nanez, et al., 2005). In that study, learning occurred in stimulus conditions where task-related internal reinforcement signals, due to accurate detection of targets of an RSVP task, were concurrent with the presentations of the subthreshold motion stimuli. Other research has verified that target processing is integral to this type of learning and that task-irrelevant learning fails to occur when target processing is obstructed by the attentional blink (Seitz, Lefebvre, Watanabe, & Jolicoeur, 2005). This line of research suggests an important role of internal reinforcement signals that are triggered by successful target processing.

While studies of task-irrelevant learning demonstrate that internal reinforcement is sufficient to produce perceptual learning, the extent to which internal reinforcement took place in the present experiments is unclear. The discriminations involved in the present tasks were very difficult, and subjects reported to have low confidence in regard to their performance even at signal levels where they performed quite well. It is thus possible that internal reinforcement signals are weak in these tasks. This possibility is very interesting because it would suggest that internal reinforcement is disassociated with accurate performance in that correct performance may not always generate internal reinforcement signals.

Another factor is that, in both experiments, stimuli of different signal levels were interleaved. Some researchers have reported that such conditions of signal level roving (Yu et al., 2004) produce stimulus uncertainty (Adini, Wilkonsky, Haspel, Tsodyks, & Sagi, 2004) and yield smaller learning effects. On the other hand, these residual learning effects are thought to more likely result from learning at low-level visual stages, which should produce responses that are largely context invariant (Adini et al., 2004; Yu et al., 2004). For instance, Yu et al. (2004) found that learning for contrast detection was absent under conditions of contrast roving and concluded that learning in their contrast detection task was mediated by decision stages. In contrast, Seitz, Yamagishi, et al. (2005) found learning in a hyperacuity task, which is thought to be mediated by low-level plasticity (Poggio, Fahle, & Edelman, 1992), under roving conditions, with feedback, similar to those employed in the experiments presented here.

A recent study (Petrov, Dosher, & Lu, 2006) proposes an augmented Hebbian model that uses selective reweighing to account for "low-level" learning either with or without external reinforcement. In this model, feedback is implemented as an additional input to the decision stage. This model is interesting in that it can account for perceptual learning data without a need for plasticity at the filtering stage of the model. This model exhibits learning without feedback for stimulus and task conditions very similar to those employed in the orientation-discrimination studies presented here. They trained the human subjects (and the model) with and without error feedback to discriminate two distinct Gabor orientations. Gabor stimuli were masked in spatial noise, and three different signal levels (contrast levels) were randomly interleaved using the method of constant stimuli. The subjects (and the model) exhibited similar learning both in the presence and in the absence of error feedback. Why did learning proceed without external reinforcement in their study but not in those presented here? We suggest two potential explanations. First, in their study, the background noise stimulus contained a subtle spatial orientation, and this orientation was changed between blocks of trials. The observed learning effects were largely specific to the orientation of the background stimulus, and thus, a large component of their learning effects could be contextual learning effects (Chun, 2000), which are thought to be a higher level form of learning and accordingly accounted for in the criterion unit of their model. Second, feedback was given for the first few trials in each block. Given reports that block feedback can be as efficient as trial feedback (Herzog & Fahle, 1997), it is possible that this limited feedback was sufficient to account for the reported learning effects.

Although the degree of learning is quite small in our studies (on average, ~10% across subjects and conditions in each study), it must be noted that the tasks used in this study involve simple discriminations of a single, primitive stimulus feature. A number of studies have found that plasticity that occurs in low-level stages of visual processing can take multiple days of training (Fahle, 2004; Sagi & Tanne, 1994; Watanabe et al., 2002) and often demonstrates a lesser degree of learning than that found in studies employing more complex stimuli (Fine & Jacobs, 2002).

It is important to note that although subjects were asked to maintain fixation during stimulus presentation, eye movements were not monitored in the reported experiments. Given the fact that the stimulus durations were 500 ms, it is possible that subjects made eye movements during task performance. However, it is not immediately clear how the presence or absence of external reinforcement would cause subjects to develop different eye-movement patterns. Furthermore, given that the pattern of eye movements would be expected to be very different between the parafoveal motion-discrimination task, for which eye movements could be large, and the foveal orientation-discrimination task, for which eye movements should be minimal, it seems unlikely that motor learning could explain both results.

It should be also noted that many experiments of perceptual learning have used staircase procedures (Adini et al., 2002, 2004; Fahle et al., 1995; Yu et al., 2004) instead of the method of constant stimuli used in our studies. Although we used the method of constant stimuli to assure that easy exemplars would be presented with equal probability at all times during training, it is possible that this procedure contributed to our failure to find learning in the absence of external reinforcement. For instance, the gradual transition between easy and difficult signal levels provided by staircase procedures may allow for better bootstrapping from the easy levels than in conditions where different signal levels are randomly interleaved.

Herzog and Fahle (1997) reported that perceptual learning in a vernier acuity task failed in the absence of external reinforcement. The results presented in this study can be regarded as an extension of that earlier work in that we show that external reinforcement is required, in certain conditions, for perceptual learning of orientation-discrimination and motion-discrimination tasks. It is important to note that in the Herzog and Fahle study, differences in initial performance between the no-feedback group and the feedback groups and intersubject differences in the degree of learning within the no-feedback group led to some ambiguity regarding the necessity of external reinforcement in learning. In a later study, the same authors proposed a model in which external reinforcement modulates the rate of learning but was not regarded as a teaching signal (Herzog & Fahle, 1998). The learning effects in our study are all-or-none depending on the presence of feedback. In addition, our study explicitly examines how the presence of easily detectable signal levels can contribute to learning of difficult signal levels, a topic that has not been directly addressed in previous studies.

Conclusions

The experiments presented in this study demonstrate that interleaving easy and difficult signal levels is not sufficient to produce learning. Still, the generality of our results remains to be clarified. Moving forward, it is important to reach a more complete understanding of the conditions for which external reinforcement plays a role in, and the mechanisms by which reinforcement affects, perceptual learning. To do this, it will be important to explicitly manipulate factors that are known to affect perceptual learning, such as reinforcement, stimulus uncertainty, task difficulty, and so forth, within otherwise unchanged experimental conditions to clarify their respective roles, and how they interact, to produce learning.

Acknowledgments

We would like to thank Andrew Kuljis and Caitlin Donahoe who helped collect data and Nozomi Ito for all her help with the research and manuscript preparation. This research was supported by grants from the NSF (BCS-0549036 and BCS-0345746), the NIH (R01EY015980-01), and Human Frontier Research (RGP18/2004) and by the NSF (CELEST).

Commercial relationships: none.

Corresponding author: Aaron R. Seitz.
Email: aseitz@bu.edu.
Address: 64 Cummington St., Boston, MA 02215, USA.

References

Adini, Y., Sagi, D., & Tsodyks, M. (2002). Context-enabled learning in the human visual system. *Nature, 415,* 790–793. [PubMed]

Adini, Y., Wilkonsky, A., Haspel, R., Tsodyks, M., & Sagi, D. (2004). Perceptual learning in contrast discrimination: The effect of contrast uncertainty. *Journal of Vision, 4*(12), 993–1005, http://journalofvision.org/4/12/2/, doi:10.1167/4.12.2. [PubMed] [Article]

Ahissar, M., & Hochstein, S. (1997). Task difficulty and the specificity of perceptual learning. *Nature, 387,* 401–406. [PubMed]

Ball, K., & Sekuler, R. (1987). Direction-specific improvement in motion discrimination. *Vision Research, 27,* 953–965. [PubMed]

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision, 10,* 433–436. [PubMed]

Chun, M. M. (2000). Contextual cueing of visual attention. *Trends in Cognitive Sciences, 4,* 170–178. [PubMed]

Fahle, M. (2004). Perceptual learning: A case for early selection. *Journal of Vision, 4*(10), 879–890, http://journalofvision.org/4/10/4/, doi:10.1167/4.10.4. [PubMed] [Article]

Fahle, M., & Edelman, S. (1993). Long-term learning in vernier acuity: Effects of stimulus orientation, range and of feedback. *Vision Research, 33,* 397–412. [PubMed]

Fahle, M., Edelman, S., & Poggio, T. (1995). Fast perceptual learning in hyperacuity. *Vision Research, 35,* 3003–3013. [PubMed]

Fine, I., & Jacobs, R. A. (2002). Comparing perceptual learning tasks: A review. *Journal of Vision, 2*(2), 190–203, http://journalofvision.org/2/2/5/, doi:10.1167/2.2.5. [PubMed] [Article]

Furmanski, C. S., Schluppeck, D., & Engel, S. A. (2004). Learning strengthens the response of primary visual cortex to simple patterns. *Current Biology, 14,* 573–578. [PubMed] [Article]

Goldstein, M., & Rittenhouse, C. H. (1954). Knowledge of results in the acquisition and transfer of a gunnery skill. *Journal of Experimental Psychology, 48,* 187–196. [PubMed]

Goodman, J. S., & Wood, R. E. (2004). Feedback specificity, learning opportunities, and learning. *Journal of Applied Psychology, 89,* 809–821. [PubMed]

Herzog, M. H., & Fahle, M. (1997). The role of feedback in learning a vernier discrimination task. *Vision Research, 37,* 2133–2141. [PubMed]

Herzog, M. H., & Fahle, M. (1998). Modeling perceptual learning: Difficulties and how they can be overcome. *Biological Cybernetics, 78,* 107–117. [PubMed]

Karni, A., & Sagi, D. (1991). Where practice makes perfect in texture discrimination: Evidence for primary visual cortex plasticity. *Proceedings of the National Academy of Sciences of the United States of America, 88,* 4966–4970. [PubMed] [Article]

McKee, S. P., & Westheimer, G. (1978). Improvement in vernier acuity with practice. *Perception & Psychophysics, 24,* 258–262. [PubMed]

Pavlov, I. P., Gantt, W. H., & Folbort, G. (1928). *Lectures on conditioned reflexes*. New York: International Publishers.

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision, 10,* 437–442. [PubMed]

Petrov, A. A., Dosher, B. A., & Lu, Z. L. (2006). Perceptual learning without feedback in non-stationary contexts: Data and model. *Vision Research, 46*, 3177–3197. [PubMed]

Poggio, T., Fahle, M., & Edelman, S. (1992). Fast perceptual learning in visual hyperacuity. *Science, 256,* 1018–1021. [PubMed]

Sagi, D., & Tanne, D. (1994). Perceptual learning: Learning to see. *Current Opinion in Neurobiology, 4,* 195–199. [PubMed]

Schultz, W. (2000). Multiple reward signals in the brain. *Nature Reviews: Neuroscience, 1,* 199–207. [PubMed]

Seitz, A., Lefebvre, C., Watanabe, T., & Jolicoeur, P. (2005). Requirement for high-level processing in subliminal learning. *Current Biology, 15,* R753–R755. [PubMed]

Seitz, A., & Watanabe, T. (2005). A unified model for perceptual learning. *Trends in Cognitive Sciences, 9,* 329–334. [PubMed]

Seitz, A. R., Nanez, J. E., Holloway, S. R., Koyama, S., & Watanabe, T. (2005). Seeing what is not there shows the costs of perceptual learning. *Proceedings of the National Academy of Sciences of the United States of America, 102,* 9080–9085. [PubMed] [Article]

Seitz, A. R., & Watanabe, T. (2003). Psychophysics: Is subliminal learning really passive? *Nature, 422,* 36. [PubMed] [Article]

Seitz, A. R., Yamagishi, N., Werner, B., Goda, N., Kawato, M., & Watanabe, T. (2005). Task-specific disruption of perceptual learning. *Proceedings of the National Academy of Sciences of the United States of America, 102,* 14895–14900. [PubMed] [Article]

Shiu, L. P., & Pashler, H. (1992). Improvement in line orientation discrimination is retinally local but dependent on cognitive set. *Perception & Psychophysics, 52,* 582–588. [PubMed]

Watanabe, T., Nanez, J. E., & Sasaki, Y. (2001). Perceptual learning without perception. *Nature, 413,* 844–848. [PubMed]

Watanabe, T., Nanez, J. E., Sr., Koyama, S., Mukai, I., Liederman, J., & Sasaki, Y. (2002). Greater plasticity in lower-level than higher-level visual motion processing in a passive perceptual learning task. *Nature Neuroscience, 5,* 1003–1009. [PubMed] [Article]

Yu, C., Klein, S. A., & Levi, D. M. (2004). Perceptual learning in contrast discrimination and the (minimal) role of context. *Journal of Vision, 4*(3), 169–182, http://journalofvision.org/4/3/4/, doi:10.1167/4.3.4. [PubMed] [Article]