

Learning to Become an Expert: Reinforcement Learning and the Acquisition of Perceptual Expertise

Olav E. Krigolson^{1,2}, Lara J. Pierce¹, Clay B. Holroyd¹,
and James W. Tanaka¹

Abstract

■ To elucidate the neural mechanisms underlying the development of perceptual expertise, we recorded ERPs while participants performed a categorization task. We found that as participants learned to discriminate computer-generated “blob” stimuli, feedback modulated the amplitude of the error-related negativity (ERN)—an ERP component thought to reflect error evaluation within medial–frontal cortex. As participants improved at the categorization task, we also observed an increase in amplitude of an ERP component associated with

object recognition (the N250). The increase in N250 amplitude preceded an increase in amplitude of an ERN component associated with internal error evaluation (the response ERN). Importantly, these electroencephalographic changes were not observed for participants who failed to improve on the categorization task. Our results suggest that the acquisition of perceptual expertise relies on interactions between the posterior perceptual system and the reinforcement learning system involving medial–frontal cortex. ■

INTRODUCTION

Birdwatchers, dog judges, and other types of perceptual experts acquire their skills through years of systematic training and practice (Ericsson, Krampe, & Tesch-Romer, 1993). Behaviorally, perceptual experts are defined by their ability to discriminate objects of expertise quickly and accurately and furthermore identify them at subordinate levels of representation (Tanaka & Gauthier, 1997; Tanaka & Taylor, 1991). For instance, a bird expert is as fast to recognize a bird at the subordinate level of “sparrow” as a novice is to recognize that same object at the basic level of “bird.” Complementing these behavioral differences, studies using ERPs have shown enhanced electroencephalographic activity at posterior recording sites when experts identify objects from within their domain of expertise (Scott, Tanaka, Sheinberg, & Curran, 2006; Tanaka, Curran, Porterfield, & Collins, 2006; Busey & Vanderkolk, 2005; Gauthier, Curran, Curby, & Collins, 2003; Tanaka & Curran, 2001). However, how experts acquire these preferential abilities and how specific regions of cortex become selectively tuned to objects from domain-specific categories remain unclear.

Insight into the development of perceptual expertise can be found in the category learning literature. From a category learning perspective, one could differentiate the perceptual capabilities of experts and novices in terms of the type of rule set used to identify objects. On one hand, novices could be thought to rely solely upon

an explicit rule set that allows identification of objects at the basic representational level (e.g., if it has beak, wings, and feathers, then it is a bird). On the other hand, in addition to an explicit rule set, experts appear to be able to access an implicit rule set that affords difficult, nonlinear judgments (Ashby & O’Brien, 2005; Ashby, Alfonso-Reese, Turken, & Waldron, 1998). Thus, the development of perceptual expertise can be associated with the development of the set of implicit rules needed to make subordinate level object classifications. Recently, Ashby and O’Brien (2005) and Ashby et al. (1998) have proposed a theoretical account of category learning that suggests that implicit rule sets are developed via a feedback-dependent, procedural, trial-and-error learning process. Indeed, accumulating experimental evidence suggests that the development of an implicit rule set depends on reinforcement learning: If people are not provided with explicit positive and negative feedback, then they do not learn the rules. Furthermore, Ashby’s theory holds that implicit category learning is mediated by structures within the BG, in conjunction with the midbrain dopamine system, ACC, inferotemporal cortex, and other brain areas. Interestingly, some of these same structures are thought to play an important role in reinforcement learning (Holroyd, Yeung, Coles, & Cohen, 2005; Holroyd & Coles, 2002). Specifically, a recent theoretical account suggests that the BG sends learning signals via the midbrain dopamine system to ACC (among other brain areas), where they are used to train the system to select task-appropriate behaviors (Holroyd & Coles, 2002). With all of this in mind, it seems reasonable to suggest that the

¹University of Victoria, ²University of British Columbia

reinforcement learning system involving medial–frontal cortex plays a key role in learning the implicit categorization rules that underlie the development of perceptual expertise.

In the present experiment, we used ERPs to investigate the role of the medial–frontal reinforcement learning system in the development of perceptual expertise. Our approach relied on examining learning related changes in two components of the ERP: the N250 and the error-related negativity (ERN). The N250 is a negative deflection in the ERP that is enhanced in amplitude when experts view objects in their domain of expertise and is associated with the processing of objects at the subordinate representational level (Scott, Tanaka, Sheinberg, & Curran, 2008; Scott et al., 2006; Tanaka et al., 2006); it is typically maximal over posterior areas of the scalp and is thought to be generated in or near the fusiform gyrus (Schweinberger, Kaufmann, Moratti, Keil, & Burton, 2007). The ERN is an ERP component that is differentially sensitive to correct responses and error responses in speeded response time tasks (response ERN) (Gehring, Goss, Coles, & Meyer, 1993; Falkenstein, Hohnsbein, & Blanke, 1991) and to correct feedback and error feedback in trial-and-error learning tasks (feedback ERN) (Holroyd & Krigolson, 2007; Miltner, Braun, & Coles, 1997). Both the response ERN and the feedback ERN are maximal over medial–frontal cortex and are thought to share a common neural generator within ACC (Holroyd et al., 2004; Miltner et al., 1997; Dehaene, Posner, & Tucker, 1994). A recent theory holds that the ERN is elicited by the impact of a dopaminergic reinforcement learning signal sent from the BG on ACC (Holroyd et al., 2005; Holroyd & Coles, 2002). Specifically, it has been proposed that phasic increases and decreases in dopamine activity, indicating whether events are better or worse than expected (Schultz, Dayan, & Montague, 1997; Ljungberg, Apicella, & Schultz, 1991, 1992; Schultz & Romo, 1990), modulate the amplitude of the ERN. Furthermore, according to this position, the response ERN is elicited when performance is evaluated based on the response, and the feedback ERN is elicited when performance is evaluated based on external feedback. Importantly, as the learning progresses and the correctness of a given response can be evaluated internally, the amplitude of the response ERN increases concomitantly. Conversely, as one becomes less reliant upon performance feedback to evaluate the success or the failure of a given action, the amplitude of the feedback ERN decreases.

Participants in the present experiment performed a difficult categorization task in which they learned to discriminate between two different families of computer generated “blob” shapes (*learnable blobs*). Intermixed with the learnable blobs, participants were also exposed to a third family of blobs that could not be classified as belonging to either family because they shared features with both groups (*morph blobs*). After each response, participants received feedback about the accuracy of

their response. We predicted that as participants learned to identify the learnable blobs, two important changes would occur in the ERP. First, we predicted that improved recognition of the learnable blobs would be associated with an increase in the amplitude of the N250, reflecting the development of a representation that afforded subordinate level object classification (Scott et al., 2006, 2008; Tanaka et al., 2006). Second, concomitant with the learning process, we predicted that participants would develop the ability to evaluate the correctness of their responses internally without the benefit of external feedback, as reflected by an increase in amplitude of the response ERN and by a corresponding decrease in the amplitude of the feedback ERN. Importantly, we predicted that these changes in ERP amplitude would be correlated with task performance. Specifically, we predicted that the N250 and the response ERN would increase in amplitude and that the feedback ERN would decrease in amplitude only for participants who demonstrated behavioral improvements in task performance. To ensure that changes in the amplitude of the ERP components were related to learning, we utilized the morph blobs as a control measure. For the morph blobs, we predicted that the amplitudes of the ERP components would remain unchanged for all participants.

METHODS

Participants

Thirty-seven college-aged participants (15 men, 22 women; mean age = 19.7 ± 1.9 years) with no known neurological impairments and with normal or corrected-to-normal vision took part in the experiment. All of the participants were volunteers who received extra credit in a first- or second-year psychology course for their participation. The participants provided informed consent approved by the Office of the Vice President, Research, University of Victoria, and the study was conducted in accordance with the ethical standards prescribed in the 1964 Declaration of Helsinki.

Stimuli

The stimuli in the present experiment were computer-generated, two-dimensional polygons, henceforth termed *blobs* (for examples, see Figure 1). The blobs were generated by a customized Matlab script (Version 7.1, The MathWorks, Inc., Natick, MA) in which two blob prototypes (A and B) were randomly created and family exemplars were produced as distortions of the prototypes (see Curran, Tanaka, & Weiskopf, 2002). The family prototypes were created by dividing a circle into 20 vertices. The possible distance of each vertex from the origin could fall within the specified range of the original circle’s radius (range = 30–70%). Once the distances of the vertices

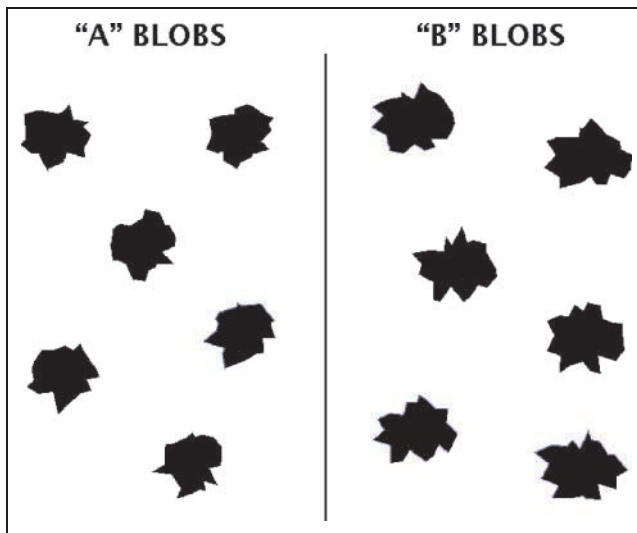


Figure 1. Example exemplar A and B blobs.

were specified, the vertices were interconnected to form a closed polygon. Five hundred exemplars (A and B) were created from each prototype by randomly changing the radius of each vertex (range = $\pm 20\%$). A *morph blob* prototype was created by averaging the vertices of the Family A and the Family B prototypes. Two hundred morph exemplar blobs were generated as distortions of the morph prototype in the same manner as the Family A and the Family B exemplar blobs.

Procedure

Participants were seated comfortably, 60 cm in front of a computer monitor in an electromagnetically shielded booth, and used a standard USB keyboard to perform a category verification task (programmed in EPrime, Version 1.1, Psychology Software Tools, Inc., Pittsburgh, PA) in which participants learned to categorize the A and the B family blobs (henceforth termed *learnable blobs*). Participants completed 10 experimental blocks of a category verification task, each of which consisted of 120 trials. Each trial started with participants viewing a fixation cross for 250 msec (0.72° of visual angle). Subsequent to this, participants viewed a blob selected randomly without replacement from a pool of possible blobs (50 A blobs, 50 B blobs, 20 morph blobs per block). When presented on a computer monitor, each blob subtended 3.8° of visual angle. Concurrent with the presentation of the blob, either the label A or the label B was presented beneath the image. If a participant thought the label matched the blob, they responded by pressing a key that indicated the label and the blob matched. If a participant thought the label did not match the image, then the participant pressed a key indicating that the label and the blob did not match. Feedback stimuli (images of apples and oranges counterbalanced between participants) indicating the accuracy of the response were

presented to participants 500 msec following their guess for 1500 msec. For example, if a participant saw an A blob with a B label and pressed the match key, then they would view the negative feedback stimulus. Note that participants were not aware of the existence of the morph family, and feedback was presented for morph blobs randomly (50% probability for either positive or negative feedback). On each trial, participants were required to respond within a response time deadline that was yoked to performance accuracy. At the start of the experiment, the deadline was 2000 msec; however, if a participant's mean accuracy for an experimental block met or exceeded 90%, then for the next block the response deadline was lowered to the participant's average RT for the previous block plus half of one standard deviation of the RT.

Behavioral Acquisition and Analysis

Response accuracy and RT (msec) were recorded for each trial. These data were analyzed with regard to blob type (learnable and morph) and experimental time (experiment start—the first 50 trials; experiment end—the last 50 trials).

Electroencephalographic Data Acquisition and Analysis

The EEG was recorded and preprocessed with standard techniques and equipment (for more details, see online Supplementary Material). To analyze the ERP components of interest (N250, response ERN, feedback ERN), segments of data from 200 msec before the event of interest (stimulus, response, and feedback) to 600 msec after the event of interest were extracted from the continuous EEG for each trial, channel, and participant. Furthermore, 10 subsets of 50 segments were taken from equally spaced intervals across the experiment to allow examination of ERP component development over time. Subsequent to the creation of the bins, average ERPs were created for the N250, the response ERN, and the feedback ERN for each of the 10 subsets for each participant. The N250 was defined as the mean voltage from 200 to 300 msec following presentation of the stimulus at channel PO7. To analyze the response ERN and the feedback ERN, we constructed difference waves by subtracting the ERP on correct trials from the ERP on error trials for each bin and participant. The response ERN and the feedback ERN were defined as the maximal negative values in the response/feedback difference waves associated with channel FCz. The amplitudes of all three ERP components were directly compared between the first bin (start of the experiment) and the last bin (end of the experiment) via paired samples *t* tests. Note that we compared the first and the last experimental bins to examine changes between the start and the end of the experiment, irrespective of

intermediating variance stemming from individual differences in learning rates. Additionally, trend analyses were also carried out on the ERP component amplitudes across all 10 bins. For all statistical tests, an alpha level of .05 was assumed.

RESULTS

Behavioral Analysis

Accuracy

Mean accuracies for classifying the learnable blobs ranged from 48.0% to 86.8% over the course of the experiment (overall mean accuracy = 67.9%). On the basis of individual accuracy scores, we divided the participants into two groups: participants with mean response accuracies greater than 70.0% (*high learners*) and participants with mean response accuracies below 70.0% (*low learners*). Mean accuracy for the high-learner group improved between the start and the end of the experiment, $t(18) = 4.73, p < .001$ (mean = 80.0%). Mean accuracy for the low-learner group did not change between the start and the end of the experiment, $t(17) = 1.88, p > .05$ (mean = 56.4%). For the morph blobs, the mean accuracy for both groups did not change over the course of the experiment (p 's $> .05$; mean high learners = 50.5%, mean low learners = 50.4%).

Response Time

Mean response times for the learnable blobs for the high-learner group were faster for correct trials (885 msec) than for error trials (950 msec), $t(17) = 4.69, p < .001$. Mean response times did not differ between correct (813 msec) and error (787 msec) trials for the low-learner group for the learnable blobs, $t(16) = 1.50, p > .05$. For the morph blobs, response time did not differ between correct and error trials for the high-learner and the low-learner groups [high learners: 814 msec vs. 803 msec, $t(16) = 1.17, p > .05$; low learners: 935 msec vs. 932 msec, $t(17) = 0.22, p > .05$].

Electroencephalographic Data: Learnable Blobs

We separately analyzed each of the three ERP components of interest (N250, response ERN, and feedback ERN) for the high-learner and the low-learner groups. Furthermore, to evaluate the extent of learning by each of these groups, we compared the amplitude of these ERP components at the start and at the end of the experiment. Where changes in the amplitudes of these ERP components between the start and the end of the experiment were detected, trend analyses were utilized to probe the relationships between these differences.

N250

Examination of the ERPs averaged to the presentation of the learnable blob stimuli revealed a bilateral posterior N250 (maximal at channel PO7; peak latency = 232 msec; Figure 2A) that increased in amplitude between the start and the end of the experiment for the high-learner group, $t(17) = 3.91, p < .001$ (Figure 2B, top panel), but not for the low-learner group, $t(16) = 1.81, p > .05$ (Figure 2C). Additionally, a trend analysis on the amplitude of the N250 across the experiment for the high-learner group revealed a significant cubic relationship, $F(1,18) = 31.12, p < .001$ (Figure 3, top panel).

Response ERN

Examination of the difference waves for the high-learner group at the end of the experiment revealed a negative deflection in the ERP consistent with the response ERN that was greater for error trials than correct trials and that was distributed over frontal-central areas of the scalp (maximal at channel FCz; peak latency = 48 msec; Figure 2D). Importantly, the amplitude of the response ERN was greater at the end of the experiment compared with the start for the high-learner group, $t(17) = 2.80, p < .05$ (Figure 2E), but not for the low-learner group, $t(16) = 0.37, p > .05$ (Figure 2F). A trend analysis on the amplitude of the response ERN indicated that for the high-learner group, the amplitude increased linearly over the course of the experiment, $F(1,18) = 6.93, p < .05$ (Figure 3, middle panel).

Feedback ERN

Analysis of the difference waves locked to the onset of the feedback stimuli revealed an ERP component with a latency and a scalp distribution consistent with the feedback ERN for both the high-learner group and the low-learner group at the start of the experiment (maximal at channel FCz; peak latency = 280 msec; Figure 2G). The amplitude of the feedback ERN decreased between the start and the end of the experiment for the high-learner group, $t(17) = 2.72, p < .05$ (Figure 2H), but not for the low-learner group, $t(16) = 1.91, p > .05$ (Figure 2I). Further, a trend analysis revealed that the amplitude of the feedback ERN for the high-learner group decreased linearly throughout the experiment, $F(1,18) = 8.84, p < .01$ (Figure 3, bottom panel).

Electroencephalographic Data: Morph Blobs

We also analyzed the ERP components associated with the presentation of the morph blobs (N250), with the responses to the morph blobs (response ERN), and with

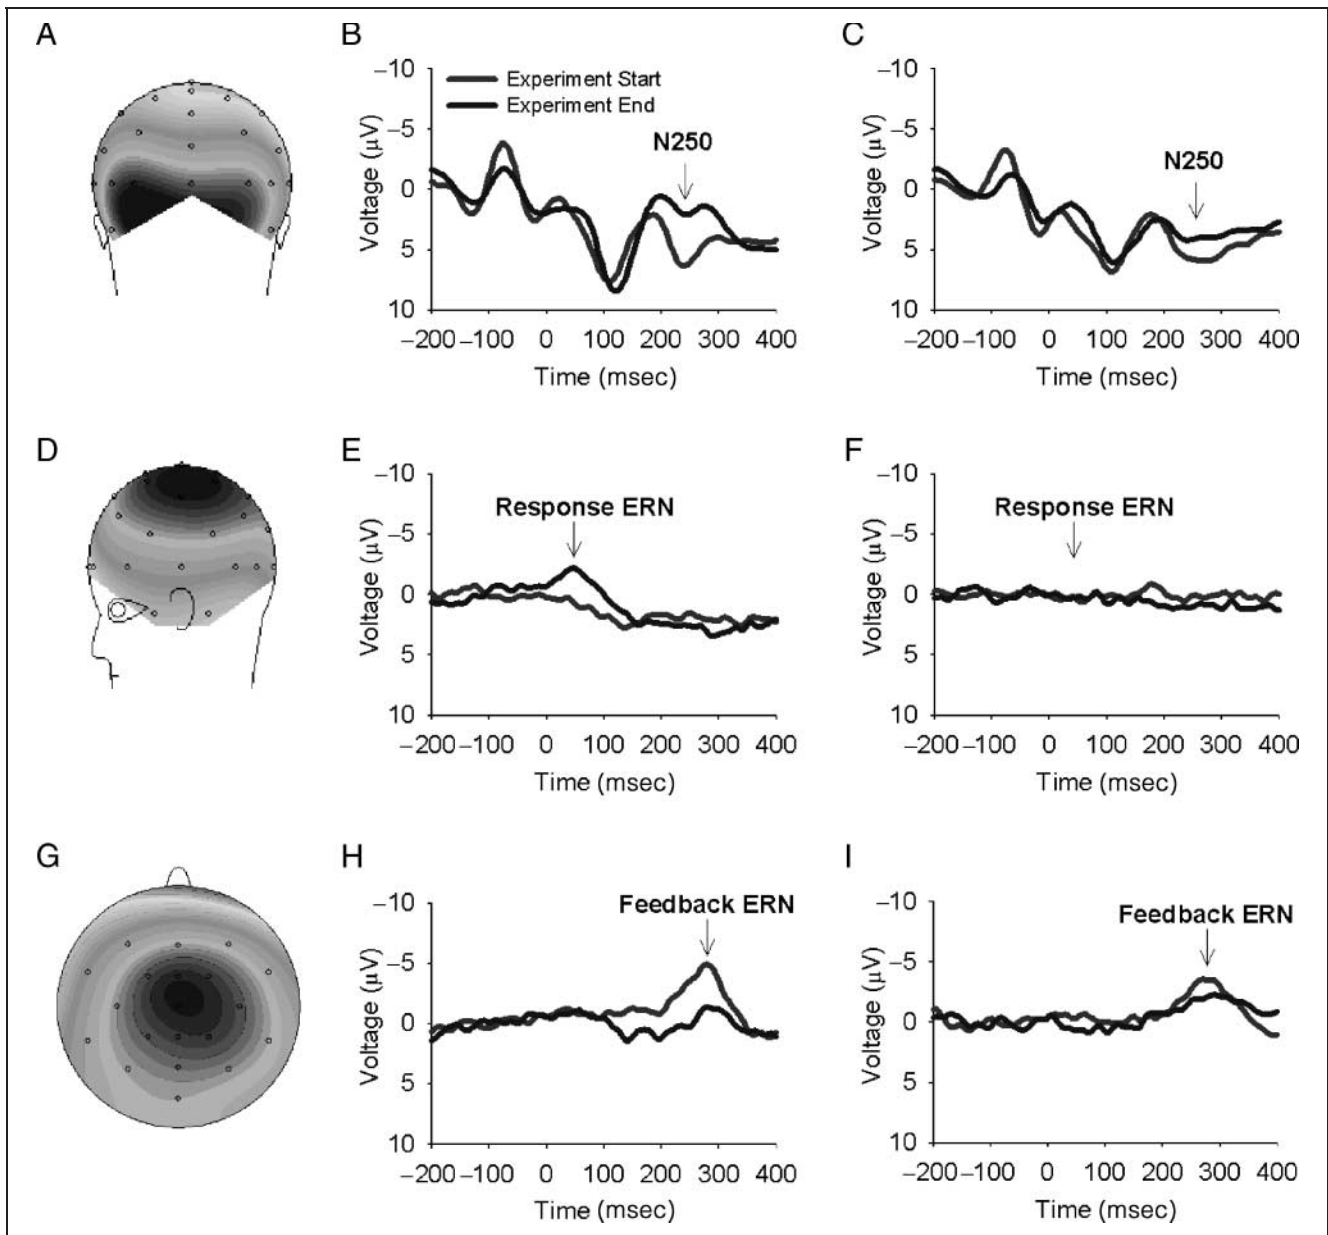


Figure 2. Scalp distributions and ERP waveforms associated with the learnable blob stimuli at stimulus presentation, the participant's response, and the feedback presentation. (A) Scalp distribution of the N250 at the end of the experiment for high-learner participants. ERP waveforms averaged to the time of presentation of the learnable blob stimuli, shown at the start and at the end of the experiment for (B) high-learner participants and (C) low-learner participants. (D) Scalp distribution of the response ERN at the end of the experiment for high-learner participants. ERP difference waves locked to the response following presentation of learnable blob stimuli, shown at the start and at the end of the experiment for (E) high-learner and (F) low-learner participants. (G) Scalp distribution of the feedback ERN at the start of the experiment for high-learner participants. ERP difference waves for feedback presentation on learnable blob stimuli trials, shown at the start and at the end of the experiment for (H) high-learner and (I) low-learner participants. Note that in each ERP panel, zero on the abscissa corresponds to the onset of the event of interest. Negative voltages are plotted up by convention.

the presentation of the feedback stimuli following the responses (feedback ERN) for both the high-learner and the low-learner groups. As only a limited number of morph blobs were shown during the experiment, comparisons were only made between the amplitude of these ERP components at the start and at the end of the task.

N250

The amplitude of the N250 associated with presentation of the morph blobs did not change between the start and the end of the experiment for the high-learner group, $t(17) = 1.18, p > .05$, or for the low-learner group, $t(16) = 0.68, p > .05$.

Response ERN

As with the N250, the amplitude of the response ERN for the high-learner group, $t(17) = 0.05$, $p > .05$, and the low-learner group, $t(16) = 1.49$, $p > .05$, on morph




Figure 3. Changes in ERP component amplitude during the course of the experiment. Blocks of 50 trials were grouped into 10 equivalently spaced bins. (Top) Change in N250 amplitude for the learnable blob stimuli across the experiment for high-learner and low-learner participants. (Middle panel) Change in response ERN amplitude for the learnable blob stimuli across the experiment for high-learner and low-learner participants. (Bottom) Change in feedback ERN amplitude for the learnable blob stimuli across the experiment for high-learner and low-learner participants.

blob trials did not differ between the start and the end of the experiment.

Feedback ERN

The amplitude of the feedback ERN on morph blob trials did not differ between the start and the end of the experiment for either the high-learner group, $t(17) = 0.78$, $p > .05$, or the low-learner group, $t(16) = 1.04$, $p > .05$.

Relationship between Behavioral and Electroencephalographic Changes

To investigate the relationship between the behavioral performance and the ERP components of interest (N250, response ERN, and feedback ERN), we correlated the change in amplitude of these ERP components (ERP component amplitude at the end of the experiment minus ERP component amplitude at the start of the experiment) for the learnable and the morph blobs with each participant's mean accuracy. The results of this analysis revealed that the change in amplitude of the response ERN was correlated with mean accuracy, Pearson $r = -.32$, $p < .05$ (Figure 4). The changes in amplitude of the other ERP components were not correlated with mean accuracy (all p 's $> .05$). We also correlated the change in amplitude of the ERP components (N250, response ERN, and feedback ERN) with each other (e.g., we correlated the change in amplitude of the N250 with the change in amplitude of the response ERN). Although there was a trend suggesting that N250 amplitudes increased with response ERN amplitudes (Pearson $r = .26$), none of these correlations were statistically significant (all p 's $> .05$).

DISCUSSION

The goal of the present study was to examine whether a reinforcement learning system associated with medial-frontal cortex could underlie the acquisition of perceptual expertise. Specifically, we sought to demonstrate that the development of perceptual expertise (as measured by an increase in the amplitude of the N250) was yoked to acquisition of the ability to internally evaluate response errors (as measured by a temporally related increase in the amplitude of the response ERN). Our hypothesis was motivated by a recent theory of category learning (Ashby & O'Brien, 2005; Ashby et al., 1998) that suggests that the development of implicit categorization rules such as those that underlie an expert's ability to make rapid and accurate subordinate level object classifications is dependent upon reinforcement learning principles—if feedback is not provided on a trial-to-trial basis, people do not develop implicit categorization rules. Category learning theory also suggests that the development of implicit rule sets is mediated by regions of




Figure 4. Change in response ERN difference waves amplitude (amplitude at the end of the experiment minus amplitude at the start of the experiment) plotted with respect to participant accuracy. The line represents the best-fit linear regression.

cortex such as the ACC and the BG—a contention supported by studies examining patients with BG lesions. Indeed, experimental results from these paradigms demonstrate that patients with BG lesions exhibit reduced performance in feedback-dependent tasks that require the discrimination of exemplars from two similar perceptual categories (Zaki, Nosofsky, Jessup, & Unverzagt, 2003; Sinha, 1999).

Our results demonstrate that as categorization performance improved for high-learner participants, the amplitude of the N250 elicited by the blob stimuli increased. For these participants, the onset of the enhancement in the amplitude of the N250 preceded an increase in response ERN amplitude and a concomitant decrease in feedback ERN amplitude. In contrast, low-learner participants showed no appreciable change in N250 amplitude, nor was there any evidence of a feedback-to-response shift in the ERN. Further, for both the high- and the low-learner participants, the morph blobs elicited no changes in amplitude of the N250, the feedback ERN, and the response ERN over the course of the experiment. The results for the morph blobs indicate that the ERP changes observed for high-learner participants with regard to the learnable blobs were not artifacts of time-on-task effects such as fatigue but instead were due to category learning. It is worth noting that our results did not indicate any changes in the amplitude of the N170, an ERP component that is also thought to index object exposure (Scott et al., 2006, 2008) for either the high-learner or the low-learner groups. This result may be attributed to the relatively brief amount of practice experienced by participants in the present study relative to previous training studies.

Our results suggest that participants who successfully utilized feedback to develop representations of the blob categories (as indicated by an increase in the amplitude of the N250) were able to use that information to evaluate their behaviors based on the mappings between

those percepts and their internally generated response (as indicated by an increase in the amplitude of the response ERN) rather than on externally provided feedback (as indicated by a decrease in the amplitude of the feedback ERN). However, although visual inspection of Figure 3 suggests a temporal relationship between the increase in the N250 and the response ERN amplitudes, we did not find a significant correlation between N250 and response ERN amplitudes. It is worth noting that even the high-learner participants continued to make multiple errors throughout the experiment due to increasing speed pressure, which may have hindered the development of the learnable blob percepts and may explain the positive but nonsignificant N250–response ERN correlation. By contrast, participants who were not able to utilize feedback (for whatever reason) failed to acquire new representations for the learnable blobs as indicated by the lack of an enhanced N250 component and the absence of the response ERN amplitude. Additionally, for these participants, the amplitude of the feedback ERN did not significantly decrease throughout the task, suggesting that they were still actively engaged in trying to learn perceptual categories.

How do reinforcement learning signals modify perceptual representations? In principle, reinforcement learning systems that utilize prediction errors can be used to train a variety of neural systems (Sutton & Barto, 1998). In the context of the present experiment, we suggest that a prediction error is initially elicited by the feedback stimuli (observed as the feedback ERN), and that on a trial-by-trial basis, this prediction error signal was used to develop and to refine perceptual representations (i.e., the implicit categorization rules) for the learnable blob stimuli. As learning progressed, participants who learned to correctly identify the learnable blobs (the high learners) developed a representation that afforded the ability to internally evaluate the consequences of their behavioral responses. As such, for the high-learner participants, a prediction error was elicited when a blob was incorrectly classified at the time of the response (observed as a response ERN) in the later stages of the experiment. In principle, the prediction error associated with the response ERN could be utilized to further improve the implicit category representations of the blob stimuli, but the data from the present study do not address this issue. In terms of the low-learner participants, for whatever reason (fatigue, boredom, inability to relate the feedback to the perceptual process), the prediction error elicited by feedback did not always result in the development of an enhanced perceptual representation, and as such there was no resulting changes in the ERP or their behavioral performance.

In conclusion, the results of the present experiment demonstrate that within the context of a perceptual learning task, individual differences in behavioral task performance (e.g., accuracy) can be attributed to specific

changes in neural processes. We found that the development of the N250 component and the shift from the feedback-to-response ERN are reliable predictors of the acquisition of perceptual expertise. Furthermore, in line with recent category learning theories, our results suggest that a reinforcement learning system associated with medial–frontal cortex may be responsible for the development of the implicit categorization rule sets used by experts to make subordinate level judgments.

Acknowledgments

The first author would like to thank the Michael Smith Foundation for Health Research (Research Trainee Grant) and the National Sciences and Engineering Research Council of Canada (PDF-343168-2007) for supporting this research. Additionally, grants to J. W. T. by the National Sciences and Engineering Research Council of Canada, the James S. McDonnell Foundation, and the Temporal Dynamics of Learning Center (NSF grant SBE-0542013) and to C. B. H. by the National Sciences and Research Council of Canada (Discovery grant RGPIN 312409-05) helped support this research.

Reprint requests should be sent to Olav E. Krigolson, Department of Psychology, University of British Columbia, 2136 West Mall, Vancouver, BC, Canada V6T 1Z4, or via e-mail: krigolson@psych.ubc.ca.

REFERENCES

- Ashby, F. G., Alfonso-Reese, L. A., Turken, U., & Waldron, E. M. (1998). A neuropsychological theory of multiple systems in category learning. *Psychological Review*, *105*, 442–481.
- Ashby, F. G., & O'Brien, J. B. (2005). Category learning and multiple memory systems. *Trends in Cognitive Science*, *9*, 83–89.
- Busey, T. A., & Vanderkolk, J. R. (2005). Behavioral and electrophysiological evidence for configural processing in fingerprint experts. *Vision Research*, *45*, 431–448.
- Curran, T., Tanaka, J. W., & Weiskopf, D. M. (2002). An electrophysiological comparison of visual categorization and recognition memory. *Cognitive, Affective & Behavioral Neuroscience*, *2*, 1–18.
- Dehaene, S., Posner, M. I., & Tucker, D. M. (1994). Localization of a neural system for error detection and compensation. *Psychological Science*, *5*, 303–305.
- Ericsson, K. A., Krampe, R. T., & Tesch-Romer, C. (1993). The role of deliberate practice in the acquisition of expert performance. *Psychological Review*, *100*, 363–406.
- Falkenstein, M., Hohnsbein, J. H. J., & Blanke, L. (1991). Effects of crossmodal divided attention on late ERP components. II. Error processing in choice reaction tasks. *Electroencephalography and Clinical Neurophysiology*, *78*, 447–455.
- Gauthier, I., Curran, T., Curby, K. M., & Collins, D. (2003). Perceptual interference supports a non-modular account of face processing. *Nature Neuroscience*, *6*, 428–432.
- Gehring, W. J., Goss, B., Coles, M. G. H., & Meyer, D. E. (1993). A neural system for error detection and compensation. *Psychological Science*, *4*, 385–390.
- Holroyd, C. B., & Coles, M. G. H. (2002). The neural basis of human error processing: Reinforcement learning, dopamine, and the error-related negativity. *Psychological Review*, *109*, 679–709.
- Holroyd, C. B., & Krigolson, O. E. (2007). Reward prediction error signals associated with a modified time estimation task. *Psychophysiology*, *44*, 913–917.
- Holroyd, C. B., Nieuwenhuis, S., Yeung, N., Nystrom, L., Mars, R., Coles, M. G. H., et al. (2004). Dorsal anterior cingulate cortex shows fMRI response to internal and external error signals. *Nature Neuroscience*, *7*, 497–498.
- Holroyd, C. B., Yeung, N., Coles, M. G. H., & Cohen, J. D. (2005). A mechanism for error detection in speeded response time tasks. *Journal of Experimental Psychology: General*, *134*, 163–191.
- Ljungberg, T., Apicella, P., & Schultz, W. (1991). Responses of monkey dopamine neurons to external stimuli: Changes with learning. In G. Bernardi, M. B. Carpenter, G. Di Chiara, M. Morelli, & P. Stanzione (Eds.), *The basal ganglia III* (pp. 487–494). New York: Plenum Press.
- Ljungberg, T., Apicella, P., & Schultz, W. (1992). Responses of monkey dopamine neurons during learning of behavioral reactions. *Journal of Neurophysiology*, *67*, 145–163.
- Miltner, W. H. R., Braun, C. H., & Coles, M. G. H. (1997). Event-related brain potentials following incorrect feedback in a time-estimation task: Evidence for a “generic” neural system for error detection. *Journal of Cognitive Neuroscience*, *9*, 788–798.
- Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, *275*, 1593–1599.
- Schultz, W., & Romo, R. (1990). Dopamine neurons of the monkey midbrain: Contingencies of responses to stimuli eliciting immediate behavioral reactions. *Journal of Neurophysiology*, *63*, 607–624.
- Schweinberger, S. R., Kaufmann, J. M., Moratti, S., Keil, A., & Burton, A. M. (2007). Brain responses to repetitions of human and animal faces, inverted faces, and objects—A MEG study. *Brain Research*, *1184*, 226–233.
- Scott, L., Tanaka, J. W., Sheinberg, D., & Curran, T. (2006). A reevaluation of the electrophysiological correlates of expert object processing. *Journal of Cognitive Neuroscience*, *18*, 1453–1465.
- Scott, L. S., Tanaka, J. W., Sheinberg, D. L., & Curran, T. (2008). The role of category learning in the acquisition and retention of perceptual expertise: A behavioral and neurophysiological study. *Brain Research*, *1210*, 204–215.
- Sinha, R. R. (1999). Neuropsychological substrates of category learning. *Dissertation Abstracts International. B. The Sciences and Engineering*, *60(5-B)*, 2381.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning* (1st ed.). Cambridge, MA: MIT Press.
- Tanaka, J. W., & Curran, T. (2001). A neural basis for expert object recognition. *Psychological Science*, *12*, 43–47.
- Tanaka, J. W., Curran, T., Porterfield, A., & Collins, D. (2006). The activation of pre-existing and acquired face representations: The N250 ERP as an index of face familiarity. *Journal of Cognitive Neuroscience*, *18*, 1488–1497.
- Tanaka, J. W., & Gauthier, I. (1997). Expertise in object and face recognition. In R. L. Goldstone, P. G. Medin, & D. L. Schyns (Eds.), *Psychology of learning and motivation series, special volume: Perceptual mechanisms of learning* (Vol. 36, pp. 83–125). San Diego, CA: Academic Press.
- Tanaka, J. W., & Taylor, M. (1991). Object categories and expertise: Is the basic level in the eye of the beholder? *Cognitive Psychology*, *23*, 457–482.
- Zaki, S. R., Nosofsky, R. M., Jessup, N. M., & Unverzagt, F. W. (2003). Categorization and recognition performance of a memory-impaired group: Evidence for single-system models. *Journal of International Neuropsychological Society*, *9*, 394–406.