

Comparing the effects of positive and negative feedback in information-integration category learning

Michael Freedberg¹ · Brian Glass² · J. Vincent Filoteo³ · Eliot Hazeltine¹ · W. Todd Maddox⁴

Published online: 25 July 2016
© Psychonomic Society, Inc. 2016

Abstract Categorical learning is dependent on feedback. Here, we compare how positive and negative feedback affect information-integration (II) category learning. Ashby and O'Brien (2007) demonstrated that both positive and negative feedback are required to solve II category problems when feedback was not guaranteed on each trial, and reported no differences between positive-only and negative-only feedback in terms of their effectiveness. We followed up on these findings and conducted 3 experiments in which participants completed 2,400 II categorization trials across three days under 1 of 3 conditions: positive feedback only (PFB), negative feedback only (NFB), or both types of feedback (CP; control partial). An adaptive algorithm controlled the amount of feedback given to each group so that feedback was nearly equated. Using different feedback control procedures, Experiments 1 and 2 demonstrated that participants in the NFB and CP group were able to engage II learning strategies, whereas the PFB group was not. Additionally, the NFB group was able to achieve significantly higher accuracy than the PFB group by Day 3. Experiment 3 revealed that these differences remained

even when we equated the information received on feedback trials. Thus, negative feedback appears significantly more effective for learning II category structures. This suggests that the human implicit learning system may be capable of learning in the absence of positive feedback.

Keywords Categorization · Implicit learning · Positive feedback · Negative feedback

Feedback plays a critical role in many forms of learning, so it is not surprising that feedback optimization has been the subject of intense investigation (Abe et al., 2011; Ashby & O'Brien, 2007; Brackbill & O'Hara, 1957; Dunn, Newell, & Kalish, 2012; Edmunds, Milton, & Wills, 2015; Galea, Mallia, Rothwell, & Diedrichsen, 2015; Meyer & Offenbach, 1962; Maddox, Love, Glass, & Filoteo, 2008; Wächter, Lungu, Liu, Willingham, & Ashe, 2009). For example, delaying feedback (Dunn et al., 2012; Maddox, Ashby, & Bohil, 2003) and adding reward to feedback (Abe et al., 2011; Freedberg, Schacherer, & Hazeltine, 2016; Nikooyan & Ahmed, 2015) impact learning. Of particular interest is the comparative effectiveness of positive and negative feedback to learning (Abe et al., 2011; Brackbill & O'Hara, 1957; Frank, Seeberger, & O'Reilly, 2004; Galea et al., 2015; Meyer & Offenbach, 1962; Wächter et al., 2009). Here, we define positive feedback as a signal that a task has been performed correctly and negative feedback as a signal that a task has been performed incorrectly.

In terms of category learning, several studies (Brackbill & O'Hara, 1957; Meyer & Offenbach, 1962) indicate a stronger influence of negative feedback (such as punishments) over positive feedback (such as rewards) when solving rule-based (RB) category problems, which can be solved by applying a verbal strategy (Ashby & O'Brien, 2005). However, it is less clear how effective positive and negative feedback are when solving

Electronic supplementary material The online version of this article (doi:10.3758/s13421-016-0638-3) contains supplementary material, which is available to authorized users.

✉ Michael Freedberg
Michael-Freedberg@uiowa.edu

¹ Department of Psychological and Brain Sciences, University of Iowa, E11 Seashore Hall, Iowa City, IA 52242-1407, USA

² Department of Computer Science, University College, London, London, UK

³ Department of Psychiatry, University of California, San Diego, La Jolla, CA, USA

⁴ Department of Psychology, Austin, TX, USA

information-integration (II) category problems (Ashby & O'Brien, 2007). Information-integration category learning involves the predecisional (nonverbalizable) synthesis of two or more pieces of information (Ashby & O'Brien, 2005). Consider the RB and II category structure examples in Fig. 1. The “discs” differ in terms of their bar frequency (x -axis) and bar orientation (y -axis). The left panel is an example of an RB category structure. The optimal linear bound (the black line) represents the best possible method for dividing the stimuli into categories and involves paying attention to the bar frequency and ignoring the bar orientation. The right panel is an example of an II category structure. It uses the same stimuli, but the optimal linear bound is now a diagonal. Here, both dimensions must be used on each trial to make an accurate category judgment. This category structure is difficult for participants to describe even when they perform with high accuracy.¹

Previously, Ashby and O'Brien (2007) examined the effectiveness of positive and negative feedback to II learning under four different feedback conditions: (1) partial positive feedback only (PFB), (2) partial negative feedback only (NFB), (3) partial positive and negative feedback (CP; control partial), and (4) full negative and positive feedback (CF; control full). To control the rate of feedback, the researchers employed an adaptive algorithm that adjusted feedback based on each participant's error rate. Therefore, the PFB, NFB, and CP groups received roughly equivalent feedback frequencies during the course of the experiment. Moreover, the researchers told participants that on trials where they did not receive feedback, they should not assume that they were right or wrong. Thus, participants were instructed to use only the feedback trials to guide their decisions. Note that when there are just two categories, as in Ashby and O'Brien (2007) (and this study), positive and negative feedback provide the same amount of information in the context of a single trial; both indicate what the correct answer should have been. However, it is possible that this information is more or less useful on correct than incorrect trials, or that positive and negative feedback engage different learning systems that are differentially suited for encoding II categories. The primary finding from Ashby and O'Brien's (2007) study was that II learning was only observed in the groups that received *both* types of feedback. Overall, the researchers did not observe a significant difference between the

PFB and the NFB groups, nor did they observe significant II learning in either the PFB or NFB groups.

Comparing the effectiveness of positive and negative feedback

We return to this issue and evaluate two possibilities regarding the utility of positive and negative feedback in shaping behavior, as previously discussed by Kubanek, Snyder, and Abrams (2015). The first possibility is that positive and negative feedback are equal reinforcers in terms of magnitude, but differ in the sign of their effect of behavioral frequency (Thorndike, 1911). This hypothesis predicts that we should find an equal benefit for positive and negative feedback in supporting II learning. A second possibility is that positive and negative reinforcement represent distinct influences on behavior (Yechiam & Hochman, 2013). In contrast, this hypothesis predicts an asymmetrical influence of positive and negative feedback (as in Abe et al., 2011; Galea et al., 2015; Wchter et al., 2009); one type of feedback may be more useful than the other. Note, that when feedback is guaranteed during categorical learning, one may expect a mutual benefit of positive and negative feedback. However, when feedback is ambiguous on some trials (when information is limited), it is possible that one type of feedback may be more useful than the other, or that they may be mutually beneficial, as in the case of Ashby and O'Brien's (2007) experiment.

Brackbill and O'Hara (1957) and Meyer and Offenbach (1962) demonstrated a clear advantage for negative feedback over positive feedback in solving RB category problems. Thus, as a starting point, we expect that participants who receive only negative feedback will demonstrate significantly stronger II learning than participants only receiving positive feedback, consistent with the asymmetry hypothesis.

However, this hypothesis runs counter to the conclusions of Ashby and O'Brien (2007). It is important to note that Ashby and O'Brien used an II category structure where each category was defined as a bivariate normal distribution and the categories partially overlapped (see Fig. 2, left panel). This category structure has three consequences. The first is that the optimal accuracy that could be achieved by obeying the optimal-linear bound was 86 %. Second, the category structure included items that were distant from the category boundary (to illustrate this we have imposed a line orthogonal to the optimal linear bound in each panel of Fig. 2). The consequence of including these items is that these trials can be solved using RB strategies because they are sufficiently far from the optimal bound. Thus, these trials may act as “lures” to initiate an RB strategy. Third, the bivariate overlapping distribution of the categories may have hindered the ability of the PFB and NFB groups to achieve II learning. The optimal accuracy that could be achieved by the best II strategy was 86 %, but the best RB strategy was satisfactory enough to yield an accuracy of 77.8 %. This difference (8.2 %) may not have been sufficiently

¹ We note that there are mixed findings regarding whether RB and II category problems are solved by one or more systems (Ashby et al., 2002; Ashby & Maddox, 2011; Dunn et al., 2012; Edmunds et al., 2015; Filoteo, Maddox, Salmon, & Song, 2005; Maddox et al., 2003; Maddox & Ashby, 2004; Maddox & Ing, 2005; Newell, Dunn & Kalish, 2011; Stanton & Nosofsky, 2007; Tharp & Pickering, 2009; Waldron & Ashby, 2001), but here we do not wish to enter the debate regarding whether one or more systems are engaged to support RB or II learning. Rather, the main objective of this study is to characterize II learning, whether positive or negative feedback is more effective for solving II category problems.

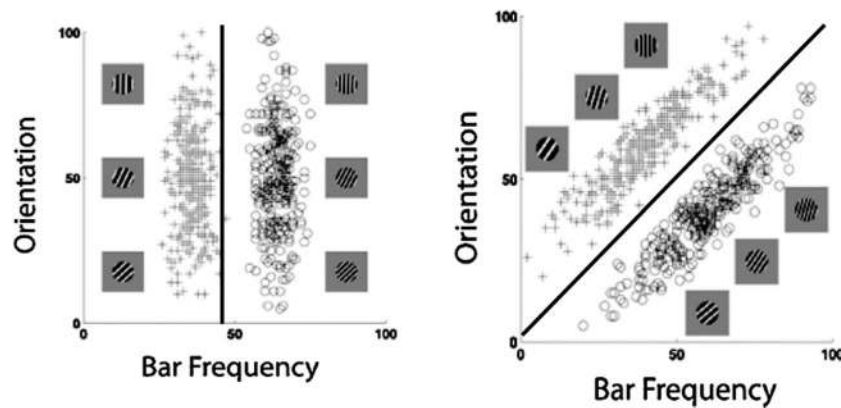


Fig. 1 Examples of rule-based (left panel) and information-integration (right panel) category structures. The optimal linear bound for each category structure is denoted by the black line

compelling to promote the abandonment of the default RB strategy. Thus, it is possible that the pattern of results found by Ashby and O'Brien (2007) were shaped by the choice of category structure.

To resolve these issues, we employed a modified version of Ashby and O'Brien's (2007) II category learning paradigm. First, we modified the category structure so that the two categories were nonoverlapping. In this way, the optimal II strategy would produce the correct response on 100 % of trials. Thus, the difference between the optimal RB strategy (81.6 %) and the II strategy (100 %) is relatively large, to maximize the incentive to abandon the RB strategy. Second, we eliminated trials that were further from the optimal linear bound. The left panel of Fig. 2 shows the category structure used by Ashby and O'Brien (2007). Although most trials are concentrated around the optimal linear bound (denoted by the solid diagonal line), many items are distant from the bound and therefore relatively easy. Rule-based strategies generally provide the correct answer for these stimuli and fail for the more difficult ones. Therefore, we

opted to exclude the easier items (see Fig. 2, middle and right panel) to promote II learning while discouraging RB learning.

Experiment 1

Experiment 1 tested the hypothesis that negative feedback would benefit II category learning more than positive feedback when the category structure minimized the effectiveness of rule-based strategies. As a starting point, we used the same adaptive algorithm used by Ashby and O'Brien (2007), with the exception that the category structure was nonoverlapping without stimuli that were 0.3 diagonal units beyond the optimal linear bound (see Fig. 2, middle panel).

Method

Participants Fifteen participants were recruited from the University of Texas at Austin community, in accordance with

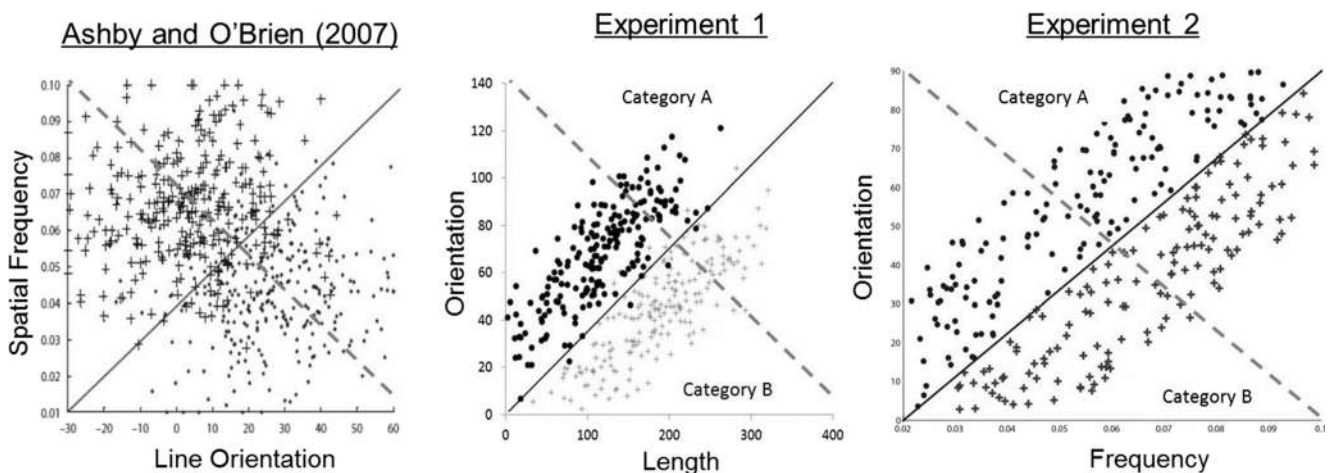


Fig. 2 Left panel. Category structure used in Ashby and O'Brien (2007). Plus signs represent Category A stimuli and dots represent Category B stimuli. The optimal linear bound is denoted by the solid diagonal line. The dashed gray line is imposed on the category structure to represent the distinction between easy and hard trials; harder trials are located closer to

the optimal linear bound and easier trials are located farther from the bound. The remaining panels represent the category structures used in Experiment 1 (middle panel) and Experiments 2 and 3 (right panel), based on 400 randomly sampled Category A (dots) and Category B stimuli (crosses)

the university’s institutional review board. Participants were randomly assigned to one of three conditions: PFB, NFB, or CP. All participants had normal or corrected-to-normal vision and were paid \$7 per session.

Stimuli On each trial, participants were shown a line that varied along two dimensions: length and orientation. Stimuli sets were pregenerated by drawing 10 sets of 80 random values of arbitrary units from two distributions (see Table 1). Each set represented one block of trials, and the presentation order of the blocks were randomized between participants. To generate a line stimulus, the orientation value was converted to radians by applying a scaling factor of $\pi/500$ (see Ashby, Maddox, & Bohil, 2002). The length value represented the length of the generated line in screen pixels. Unlike Ashby and O’Brien (2007), the large positive covariance between the two dimensions ensured that stimuli represented values close to the optimal linear decision bound of $y = x$. The rationale for this strategy is that trials that exist on the extreme ends of each category are easier to categorize because one dimension becomes increasingly more important than the other. For instance, if a trial stimulus has an orientation of 90 degrees, then there is a significantly greater chance that it belongs to Category A than a stimulus that has an orientation of 45 degrees. Likewise, a stimulus with a length of 350 has a significantly greater chance of belonging to Category B than a stimulus with a frequency of 175. Therefore, we excluded these trials.

Procedure Participants in the positive feedback (PFB) condition only received positive feedback. Participants in the negative feedback (NFB) condition only received negative feedback. Finally, participants in the control condition (CP) received both types of feedback (positive and negative) on ~26 % of trials. For all groups, the overall proportion of feedback was approximately 27 %. The PFB condition never received feedback after an incorrect response, and the NFB condition never received feedback after a correct response.

To control the proportion of feedback trials across sessions and conditions, an adaptive algorithm was used (Equation 1). The algorithm, developed by Ashby and O’Brien (2007) was designed to roughly equate feedback between all groups. Whereas the NFB group was given feedback on 80 % of

incorrect trials, the PFB group was given feedback on trials according to the following algorithm:

$$P(\text{Feedback} \mid \text{Correct Trial}) = .8 \frac{Q(\text{errors on last 50 trials})}{Q(\text{correct on last 50 trials})}, \tag{1}$$

where P is probability and Q is proportion. The main function of this algorithm is to decrease the probability of feedback for the PFB group as performance improves (see the **Supplementary Method** section for a graphical representation).² In the CP condition, feedback was given at a rate which would provide the same amount of feedback as the NFB and PFB conditions. The probability of feedback on each trial was $P(\text{Feedback}) = 0.8Q(\text{errors on last 50 trials})$. During the first 50 trials, the error rate was fixed to 0.5.

Participants completed three sessions, each with 800 trials. Participants were instructed to classify the line stimuli into two categories. The PFB group was instructed that on a portion of trials they would receive positive feedback that would be helpful toward making their judgments. The NFB group was instructed that on a portion of trials they would receive negative feedback that would indicate that their selection was wrong. The CP group was told that they would receive both types of feedback. On no feedback trials, participants in all conditions were instructed not to assume that they were right or wrong, but to use their partial feedback to guide their decisions.

Each trial began with the presentation of a single line that remained on the screen until the participant made their judgment. Participants responded on a standard keyboard by pressing the z key if they believed the stimulus belonged to Category A and the / key if they believed the stimulus belonged to Category B. The stimulus and feedback were presented in white on a black background. Positive feedback took the form of the phrase “Correct, that was an A” and negative feedback took the form of the phrase “Error, that was a B.” Each trial began with a 500 ms fixation cross, followed by a response terminated stimulus presentation, followed by 1,000 ms of feedback (present or absent), and a 500 ms intertrial interval (ITI). Blocks included 100 trials that were separated by participant-controlled rest screens. Participants completed eight blocks in each session. Sessions were usually completed on consecutive days, with no more than 3 days between consecutive sessions.

Table 1 Category distribution characteristics. Dimension x (length) is represented in pixels. Dimension y (orientation) was converted to radians with a scaling factor of $\pi/500$. μ = Mean, δ = Standard Deviation, $\text{Cov}_{x,y}$ = Covariance between dimensions x and y

	μ_x	μ_y	σ_x	σ_y	$\text{COV}_{x,y}$
Category A	185	115	65	65	339
Category B	115	185	65	65	339

² Another important consequence of this feedback mechanism is that improvements in performance result in decreases in positive feedback for the PFB group. Thus, the PFB group is penalized with less feedback for improving their performance, and rewarded with more feedback for performing worse. Consequently, this feature of the algorithm may have caused an elimination of group differences, which may also account for Ashby and O’Brien’s (2007) pattern of results.

Model-based analyses We fit seven classes of decision bound models to each participant's data. Four of the models assume a rule-based strategy: (1) Conjunctive A, (2) Conjunctive B, (3) unidimensional length, and (4) unidimensional orientation. The two conjunctive models assume that the participant sets a criterion along the length dimension that divides the stimuli into short and long bars and a criterion along orientation that divided the stimuli into shallow and steep bars. Conjunctive A assumes that the participant classifies stimuli into Category A if they were short and shallow, and into Category B otherwise. Conversely, Conjunctive B assumes the participant classifies stimuli into Category B if they were long and steep and into Category A otherwise. The unidimensional strategies assume that the participant ignores one dimension when making their judgments. Unidimensional length assumes that the participant sets a criterion on bar width and categorized based on that value. Similarly, unidimensional orient assumes that the participant sets a criterion on orientation and categorized based on that value (see Fig. 3 for categorization strategy examples). A fifth model assumes that the participant responds randomly (the random responder model). Because our primary interest was in how learning differed between groups, we excluded participants (a) if the best fitting model was the random responder and (b) if accuracy scores on Day 3 did not exceed 50 %.

The final two models assumed that the participant uses an II strategy when making their judgments. The optimum general linear classifier (OPT-GLC) represents the most accurate strategy for dividing the stimuli and is denoted by the gray line in the right panel in Fig. 3. The suboptimal GLC represents a slightly inferior, but still nonverbal, strategy for dividing the stimuli based on the angle of the diagonal line. Thus, whereas the optimal GLC strategy assumes a 45° angle in the diagonal linear bound, the suboptimal GLC assumes a diagonal that deviates slightly from 45 degrees.³ The best model fit for each participant was determined by estimating model parameters relevant to each strategy and using the method of maximum likelihood. Maximum likelihood was defined as the smallest Bayesian information criterion (BIC; Schwarz, 1978) reached for each model fit. BIC was calculated by the following equation:

$$BIC = r \ln N - 2 \ln L, \quad (2)$$

where N equals sample size, r is the number of free parameters, and L is the likelihood of the model given the data.

Data archiving Trial-level raw data for Experiments 1–3 are available at

<http://psychology.uiowa.edu/hazelab/archived-data>.

³ Note that although Ashby and O'Brien (2007) do not mention the use of a suboptimal GLC, they do include this model in their analyses.

Results

Proportion of feedback trials To evaluate the algorithm's feedback rate, we submitted the proportion of feedback trials to a 3 (condition: PFB, NFB, CP) \times 3 (day) repeated measures ANOVA. There was a significant main effect of day, $F(2, 24) = 14.6, p < .001, \eta_p^2 = 0.55$, as well as a marginally significant Day \times Condition interaction, $F(4, 24) = 2.61, p = .06, \eta_p^2 = 0.30$. This interaction indicates that the NFB and CP groups experienced a significant reduction in feedback across days (NFB Day 1: 31 %; Day 2: 23 %, Day 3: 21 %; CP Day 1: 29 %, Day 2: 28 %, Day 3: 21 %), whereas the PFB group did not (PFB Day 1: 30 %, Day 2: 29 %, Day 3: 28 %). There was no main effect of condition ($F < 1$). Post hoc comparisons revealed no significant pairwise differences between the groups ($|t| < 1$). These results indicate that all groups received feedback on a roughly similar proportion of trials, but that the NFB and CP groups received slightly less feedback than the PFB group on Days 2 and 3. Although the NFB group received slightly less feedback than the PFB group, the NFB group demonstrated greater improvements in accuracy.⁴

Accuracy-based analysis We performed a pairwise Wilcoxon sign test on the twenty-four 100-trial blocks of the experiment between all groups, similar to Ashby and O'Brien's (2007) analysis (see Fig. 4, left panel). There was a significant difference between CP and PFB (sign test: $S = 19$ of 24 blocks, $p < .01$), but not between NFB and PFB (sign test: $S = 15$ of 24 blocks, $p = .31$), nor between NFB and CP (sign test: $S = 12$ of 24 blocks, $p = 1.00$). Overall, these results indicate that there was only a significant pairwise difference between the CP and PFB groups, indicating that the CP group outperformed the PFB group.

Model-based analysis The left panel of Fig. 5 reveals the results of the modeling analysis for Experiment 1. The category boundaries for Day 3 were modeled for each participant separately. The best fit model for each participant indicated that an II strategy was used by three participants in the CP group, three participants in the NFB group, and one participant in the PFB group.

Discussion

Although the NFB group did not significantly outperform the PFB group, there was a trend suggesting that negative feedback led to more learning than positive feedback, and participants receiving both types of feedback performed no better than participants receiving only negative feedback. Moreover, the two groups appeared to prefer different strategies. Three out of five

⁴ This is also the case for Experiments 2 and 3.




Fig. 3 Examples of unidimensional (left panel), conjunctive (middle panel), and optimal-II (right panel) categorization strategies from three participants in Experiment 1. Gray lines denote the linear bound(s) for each categorization strategy

participants receiving only negative feedback engaged an II strategy, whereas only one of the participants receiving only positive feedback did so. Thus, although we did not find differences in accuracy between the PFB and NFB groups, our results suggest that negative feedback may be more helpful toward promoting II learning than positive feedback

Although not predicted, the bulk of the NFB group’s learning improvements were observed *between* days (offline changes; changes in performance between consecutive sessions) and not *within* each day (online changes; changes in performance during task engagement; see Fig. 6, top panels). To confirm this impression, we conducted two additional analyses. First, we submitted within-day learning scores (defined as accuracy on Block 8 minus accuracy on Block 1 for each day) to a group (PFB, NFB, CP) by day repeated-measures ANOVA. This revealed a marginally significant main effect of day, $F(2, 42) = 3.170, p = .06, \eta_p^2 = 0.209$, but no main effect of group ($F < 1$). The interaction between group and day, however, was marginally significant, $F(4, 42) = 2.490, p = .07, \eta_p^2 = 0.293$. These results suggest that within-day accuracy changes were lower for the NFB group on Days 2 and 3, and that accuracy improvements decreased across days (see Fig. 6, top middle panel).

Second, we submitted between-day scores (defined as accuracy on the first block of the following day minus accuracy on

the last block of the previous day) to a group by day (Day 2 minus Day 1, Day 3 minus Day 2) repeated-measures ANOVA. The results revealed no main effect of day, and no interaction (F s < 1). The main effect of group, however, was marginally significant, $F(2, 12) = 3.545, p = .06, \eta_p^2 = 0.371$. Post hoc comparisons revealed that NFB differed marginally from PFB ($p = .06$), CP did not differ from PFB ($p = .741$), nor did NFB differ from CP ($p = .199$). This analysis showed that the NFB group experienced larger between-day improvements than the PFB group. Thus, it is possible that negative feedback may engage offline processes not engaged by positive feedback. However, because we did not identify strong behavioral differences between groups, the findings are only suggestive. To resolve this ambiguity, Experiment 2 used the same category structure from Experiment 1, but substituted a more precise method for equating feedback, and included three additional participants to each group to increase statistical power.

Experiment 2

Despite the fact that the NFB group was given less feedback than the NFB group, Experiment 1 suggested a trend toward




Fig. 4 Accuracy plotted across day and group for Experiments 1 (left panel) and 2 (right panel). PFB = positive feedback only, NFB = negative feedback only, CP = control partial (both feedback types)




Fig. 5 Number of participants in each group that engaged a unidimensional, conjunctive, or II categorization strategy on Day 3 for Experiment 1 (left panel) and Experiment 2 (right panel)

more successful engagement of II strategies and higher accuracies for the NFB group. Therefore, in Experiment 2 we used an alternative algorithm to more precisely equate feedback between groups. Additionally, we included eight participants in each group (a total of 24 participants) to increase our power to detect a potential difference between groups. We hypothesized that we would detect stronger learning for the NFB group over the PFB group. Additionally, since Experiment 1 suggested that negative feedback leads to greater offline changes in category learning than positive feedback, we predicted that offline changes in accuracy would be greater for the NFB group over the PFB group.

Method

Participants Thirty participants were recruited from the University of Iowa community in accordance with the

universities institutional review board. Six participants were excluded for poor performance or if they were classified as a random responder by our modeling analysis (1, for PFB, 3 for NFB, and 2 for CP). Participants were randomly assigned to one of three conditions: PFB, NFB, or CP. Eight participants were assigned to each group and balanced based on age and sex (PFB: average age = 24.89 ± 4.49 years, four females; NFB: average age = 24.13 ± 4.48 years, four females; CP: average age = 22.85 ± 4.69 years, five females). All participants had normal or corrected-to-normal vision and were paid \$10 per session.

Stimuli Participants were shown a Gabor patch that varied along two dimensions: frequency and orientation. On each trial a random value for each dimension was generated and combined to form a Gabor patch. The orientation was free to rotate between 0° (completely vertical lines) and 90° (completely horizontal lines). Frequency was free to vary

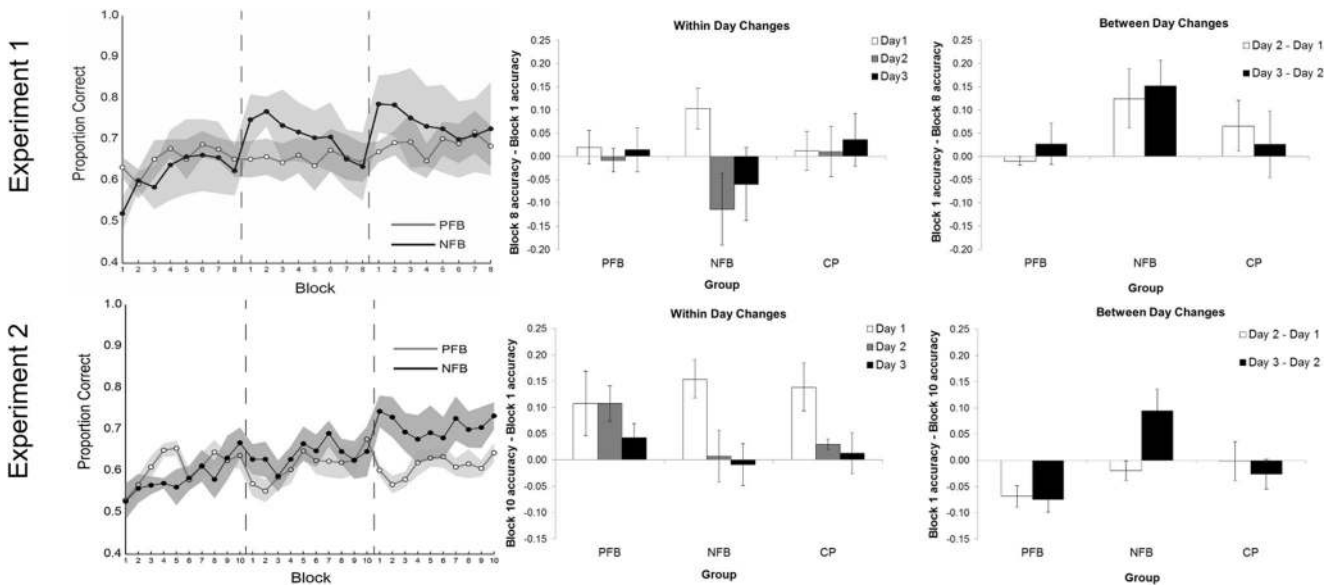


Fig. 6 Comparison of within and between-day changes for Experiment 1 (top panels) and Experiment 2 (bottom panels). The left panels plot accuracy for the PFB and NFB groups across days. The shaded region represents the SEM, and the dashed lines represent breaks between days. The middle panels plot within-day changes (last block accuracy minus

first block accuracy) for all groups across each day. The right panel plots between-day changes (first block accuracy on the following day minus last block accuracy on the previous day) for both between-day periods. Error bars represent the SEM

between 0.02 and 0.10 cycles per degree. Table 2 details the characteristics for the category structure used in Experiments 2 and 3. The right panel of Fig. 2 represents 400 randomly drawn trials from each category distribution. As in Experiment 1, items that extended beyond a distance of 0.3 diagonal units perpendicular to the optimal linear bound were not presented to participants (see Fig. 2 for a comparison between category structures).⁵

Procedure Participants in the positive feedback (PFB) condition only received positive feedback. Participants in the negative feedback (NFB) condition only received negative feedback, and participants in the control condition (CP) received both types of feedback (positive and negative) on ~20 % of trials. For all groups, the overall proportion of feedback was approximately 20 %. Furthermore, for the CP condition, there was the constraint that equal amounts of positive and negative feedback be given. The PFB condition received no feedback after an incorrect response, and the NFB condition received no feedback after a correct response.

To control for the proportion of feedback trials across session and condition, we used an adaptive algorithm (Equation 3). Although Ashby and O’Brien (2007) were able to roughly equate feedback between the PFB and NFB groups in their experiment, the CP group received less feedback than the NFB and PFB groups (although this was only significant for the PFB group) and participants received different amounts of feedback on each day. This was because the error rate determined how much feedback participants in the PFB group received.

For Experiment 2, feedback was given on trials eligible for feedback (i.e., incorrect trials for the NFB group and correct trials for the PFB group) if the following expression was true:

$$abs\left(\left(\frac{total\ feedback\ trials}{Total\ Trials}\right)-0.20\right) > abs\left(\left(\frac{total\ feedback\ trials + 1}{Total\ Trials}\right)-0.20\right). \tag{3}$$

This mechanism adjusted the trial-by-trial feedback so that the total amount of feedback given on each day was as close to 20 % of all trials for all groups. After each trial in which a response was made making feedback possible, we calculated the overall percentage of feedback (1) if feedback was to be given on the current trial (right side of Equation 3), and (2) if feedback was not to be given (left side of Equation 3). The option that brought the total percentage feedback closer to 20 % was chosen (see [Supplementary Method](#) section for a detailed example). In sum, the feedback algorithm favors the distribution of feedback when the percentage of feedback distributed falls below 20 % of all previous trials, and favors the

⁵ Note that while Ashby and O’Brien (2007) employed a category structure with a bivariate normal distribution, we used an evenly distributed category structure.

Table 2 Category distribution characteristics. Dimension y (orientation) is represented in degrees. Dimension x (frequency) was converted to degrees by norming dimension x and multiplying by 90. μ = Mean, σ = Standard Deviation, $Cov_{x,y}$ = Covariance between dimensions x and y

	μ_x	μ_y	σ_x	σ_y	$COV_{x,y}$
Category A	56	36	21	20	304
Category B	36	56	22	22	304

withholding of feedback when feedback exceeds 20 %. Thus, there is constant adjustment after each trial response to keep the amount of total feedback anchored towards 20 %.

For the CP condition, feedback type (positive or negative) was dependent on the percentage of positive and negative feedback trials as calculated throughout the experiment. If the response was correct, then the proportion of positive feedback trials was calculated and the circumstance (presenting or withholding feedback) that promoted positive feedback closer to 20 % of correct trials was chosen. Likewise, if the response was incorrect, then the proportion of negative feedback trials was calculated and the circumstance (presenting or withholding feedback) that brought the total proportion of feedback closer to 20 % was chosen. Thus, participants in the CP condition received no feedback on ~80 % of all trials, positive feedback on ~20 % of correct trials, and negative feedback on ~20 % of incorrect trials. Feedback instructions for Experiment 2 were identical to Experiment 1.

Each trial began with the presentation of a single Gabor patch and remained on the screen until the participant made his or her judgment. Participants responded on a standard keyboard by pressing the z key if they believed the stimulus belonged to Category A and the m key if they believed the stimulus belonged to Category B. The stimulus and feedback were presented on a gray background. Positive feedback took the form of the word “Correct,” presented in green font and negative feedback took the form of the word “Incorrect,” presented in red font. All feedback remained on the screen with the stimulus for 1,500 ms. Trials with no feedback showed only the stimulus on a gray background for 1,500 ms. Ten blocks of 80 trials were completed and were separated by participant-controlled rest periods. Sessions were usually completed on consecutive days, with no more than 3 days between consecutive sessions. One participant in the NFB group only completed nine of the 10 blocks on Day 1, but completed all blocks on Day 2 and Day 3.

Model-based analysis The modeling analysis for Experiment 2 was identical to Experiment 1.

Results

Proportion of feedback trials To determine how our feedback algorithm controlled the rate of feedback, we submitted the

proportion of feedback trials to a two-factor ANOVA using condition (PFB, NFB, and CP) and day as factors. There was a significant effect of day, $F(2, 42) = 3.60, p < .05, \eta_p^2 = 0.146$, condition, $F(2, 21) = 7.783, p < .005, \eta_p^2 = 0.426$, and a significant Day \times Condition interaction, $F(4, 42) = 3.876, p < .01, \eta_p^2 = 0.270$. Post hoc comparisons revealed a significant difference in the proportion of feedback received between the NFB and PFB group, $t(7) = 3.50, p < .05$, and between the NFB group and the CP group, $t(7) = 4, p < .01$, but not between the PFB and CP groups ($|t| < 1$). Note that these effects are the product of the low variance caused by the precision of our feedback mechanism. This is supported by the fact that the CP and PFB groups both received 20 % feedback on each day, while the NFB group experienced 19 %, 19 %, and 18 % feedback across the three days. We do note that, similar to Experiment 1, the NFB group received significantly less feedback than the PFB groups.

Accuracy We performed a pairwise Wilcoxon sign test on the 30 blocks between all groups, similar to Ashby and O'Brien (2007) (see Fig. 4, right panel). There was a significant difference between CP and PFB (sign test: $S = 24$ of 30 blocks, $p < .005$), NFB and PFB (sign test: $S = 23$ of 30 blocks, $p < .01$), but not between NFB and CP (sign test: $S = 17$ of 30 blocks, $p = .585$). This analysis indicates a strong advantage in learning for the NFB and CP groups over the PFB group.

Although the PFB group experienced a gain in accuracy on Day 1 from Block 1 (53 %) to Block 10 (64 %), no further learning was observed for the rest of the experiment. In contrast, the NFB and CP group continued to show strong evidence of learning throughout the experiment, reaching accuracies of 73 % and 70 %, respectively by Block 10 of Day 3 (64 % for PFB). Despite the fact that the NFB group received significantly less feedback than the PFB and CP groups, we observed a significant advantage for the NFB group over the PFB group.

Within- and between-day accuracy changes As in Experiment 1, we submitted the within-day accuracy scores (defined as accuracy on Block 10 minus accuracy on Block 1 for each day) to a group (PFB, NFB, CP) by day repeated measures ANOVA. This revealed a significant main effect of day, $F(2, 42) = 6.726, p < .005, \eta_p^2 = 0.243$, but no main effect of group, and no interaction ($F_s < 1$). These results resembled the results of Experiment 1 (within-day improvements decreased across days), with the exception that Day 1 improvements were more similar across groups. Thus, within-day changes in accuracy were statistically similar between groups (see Fig. 6, bottom middle panel).

Furthermore, we submitted between-day accuracy scores (defined as accuracy on the first block of the following day minus accuracy on the last block of the previous day) to a group by day (Day 2 minus Day 1, Day 3 minus Day 2) repeated-measures ANOVA. The results revealed no main effect of day, $F(1, 21) = 1.417, p = .247$, and a marginally

significant interaction, $F(2, 21) = 3.349, p = .06$. The main effect of group, however, was significant, $F(2, 21) = 9.966, p < .005, \eta_p^2 = 0.479$. Post hoc comparisons revealed that NFB differed significantly from PFB ($p < .005$), CP differed marginally from PFB ($p = .072$), but NFB did not differ significantly from CP ($p = .124$). These results are similar to those of Experiment 1 where learning differences were only identified between days and provide a clearer picture regarding the benefit of negative feedback over positive feedback; it appears that negative feedback affords the engagement of offline processes that cannot be engaged by positive feedback alone.

Model-based analyses As in Experiment 1, the model-based analyses of the patterns of responses in Experiment 2 suggested that participants in the NFB group were more likely to use an II strategy than participants in the PFB group. The right panel of Fig. 5 reveals the number of participants whose strategy were best modeled by either a unidimensional, conjunctive, or II strategy. The PFB group almost unanimously favored a unidimensional strategy; seven of eight participants used a unidimensional categorization strategy. For the NFB group, four participants used a GLC strategy while the other four participants chose to use either a unidimensional or conjunctive strategy. Finally, the CP group mostly engaged a unidimensional or GLC strategy (unidimensional: 4, conjunctive: 1, GLC: 3).⁶

Although our modeling process selects the best-fitting model based on the lowest BIC, it does not indicate the probability that the best-fitting model is adequately superior to the other models (Edmunds et al., 2015). To determine how likely the best-fitting

⁶ To test whether our modeling analysis favored the selection of RB models, we conducted three model recovery simulations, similar to Donkin, Newell, Kalish, Dunn, and Nosofsky (2015). We performed this analysis for the suboptimal GLC, the unidimensional-length, and Conjunctive B models. To conduct this analysis, we extracted the best fitting parameters for each participant from our original modeling analysis. Next, we used the parameters to form new optimal linear bounds for each of our 24 participants and simulated their Day 3 responses as if they were using that bound. Finally, we reran our modeling analyses for all participants with these new responses. If our modeling analysis biases the selection of RB models over II models, then we should identify several participants who cannot be fit with the II models, even when we assume that they are using a GLC strategy. Similarly, if our modeling analysis biases the selection of II strategies over RB strategies, then we should identify several participants who cannot be fit with RB strategies, even when we assume they are using either a conjunctive or unidimensional strategy.

When we modeled each participant's responses using their suboptimal II parameters, all participants were best fit by an II strategy. Thus, our modeling analysis did not favor the selection of an RB strategy over an II strategy; all participants in our original modeling analysis had the chance to be modeled using an II strategy, and either were, or another model fit better. Finally, we reran our models assuming all participants adhered to a unidimensional and a conjunctive strategy, in separate analyses. For both of these model recovery simulations, we only identified one participant (of 24) who was best fit by an II strategy. Thus, although our modeling analysis represented a very slight bias toward II strategies, we still did not identify any participants in the PFB group who used an II strategy.

model derived from our analysis is actually the most appropriate model over the alternative models, we computed model probabilities based on Bayesian weights (Wagenmakers & Farrell, 2004; See [Supplementary Method](#) section). These probabilities for each model and each participant are plotted in Fig. 7 as a heat map (including Experiments 1 and 3). The darker the corresponding box, the higher the model probability is for that model. For the NFB group, the probability of the winning model being the GLC, and not the unidimensional, is 97 %; the probability of the winning model being the GLC, and not conjunctive, is 78 %. Thus, there is a high probability that the correct model is the winning model derived from our model analysis. Similarly, we can infer with strong confidence that the PFB group was correctly modeled by an RB strategy; the probability of the winning model being an RB model (unidimensional or conjunctive), rather than either II model, is 68 %. This analysis confirms that the PFB group was best modeled by an RB strategy, and the NFB group was best modeled by an II strategy.

Summary of Experiments 1 and 2

Experiments 1 and 2 reveal an advantage for negative feedback over positive feedback in promoting II learning. Both experiments used similar category structures (e.g., nonoverlapping), but different mechanisms for controlling the rate of feedback; whereas Experiment 1 used an error-based method of controlling feedback, similar to Ashby and O’Brien (2007), feedback in Experiment 2 did not depend on the error rate. This pattern of results is likely the product of the type of feedback received, but the distribution of trials that received

feedback may also play a critical role. Given that our mechanisms for controlling the feedback rate did not control which stimuli yielded feedback across perceptual space, it could be the case that the PFB and NFB group may have received qualitatively different information on their feedback trials.

For instance, consider a situation where two participants, one in the PFB group and the other in the NFB group, are both achieving 75 % accuracy. The PFB participant can receive feedback on a range of trials that span the 75 % of accurate trials, and the distribution of these trials in the stimulus space should be biased toward stimuli far from the category boundary. However, the NFB participant can only receive feedback on the 25 % of trials that were incorrect. These trials are more likely to involve stimuli that are closer to the category boundary. Thus, the feedback received by the NFB participant may be more useful because it focuses on the harder trials (trials closer to the optimal linear bound). This is supported by research showing that II category learning is facilitated by initial training on “harder” trials over “easier” trials (Spiering & Ashby, 2008). Figure 8 plots the percentage of feedback given to the PFB and NFB groups for three levels of difficulty (i.e., distances from the boundary; top) and across perceptual space (bottom). Whereas the NFB group appears to have a strong concentration of feedback trials close to the optimal linear bound (denoted by the white line), the PFB group appears to have received more distributed feedback.

This presents a possible explanation for our results: perhaps the NFB group performed significantly better than the PFB group because the NFB group received more useful information from their feedback (i.e., more feedback concentrated




Fig. 7 Heat plot of model probabilities for each model and participant. Each participant is plotted across rows, and each model type is plotted across columns. Black represents a model probability of 1 and white represents a model probability of 0. SUB-OPT GLC = suboptimal

GLC, OPT-GLC = optimal GLC, UNI-L = unidimensional length, UNI-O = unidimensional orientation, CONJ A = Conjunctive A, CONJ B = Conjunctive B, Flat = flat model

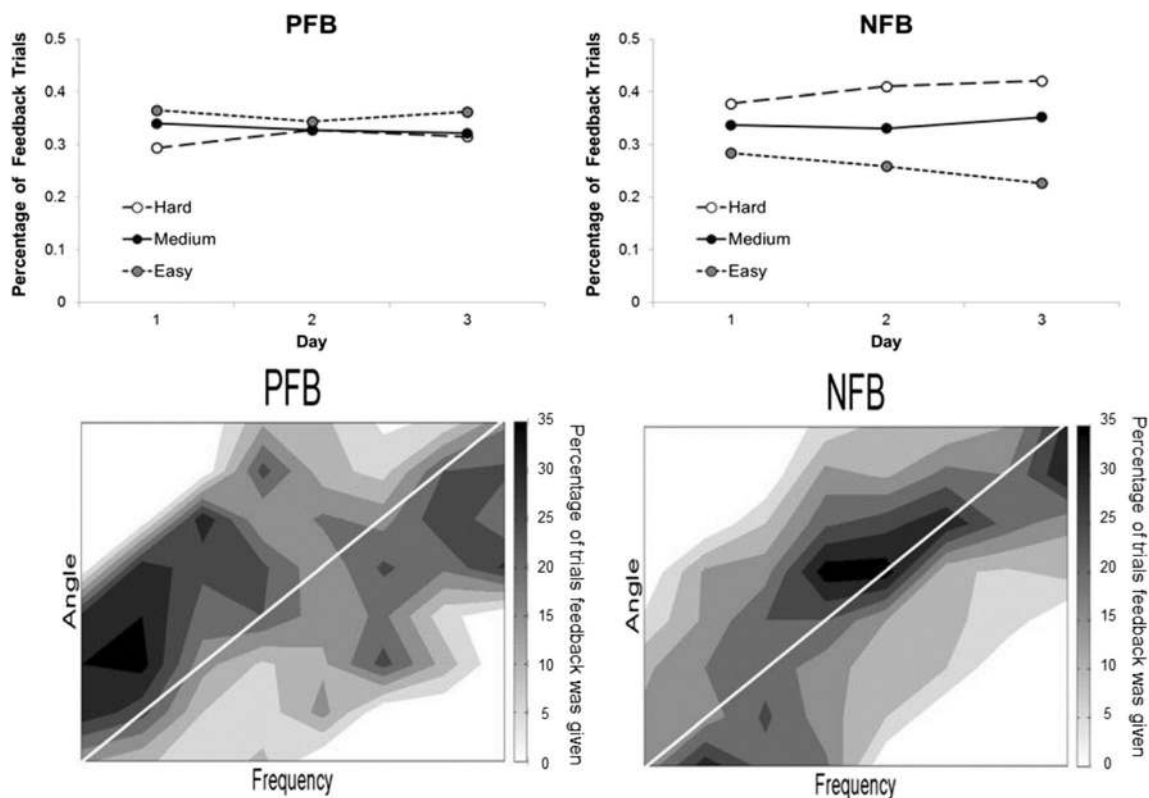


Fig. 8 (Top Panels) Percentage of feedback trials plotted across days (x-axis) for Hard, Medium, and Easy trials (separate lines) for Experiment 2. (Bottom Panels) Percentage of feedback plotted across perceptual space

for PFB (left panel) and NFB (right panel) on Day 3. Darker shades reflect higher percentages of feedback given for those trial types. The white lines represent the optimal linear bound

toward the optimal linear bound). To claim that negative feedback is more effective for teaching II categories than positive feedback, the type of stimuli that receive feedback must be equated across the PFB and NFB groups. Thus, in Experiment 3 we adjusted trial feedback so that participants in the PFB group received feedback mostly on difficult trials in a corresponding fashion to the NFB group run in Experiment 2.

Experiment 3

For Experiment 3, we ran a group that received positive feedback with an adaptive algorithm designed to match the biased distribution of feedback towards harder trials as in the NFB group. We refer to the new PFB group as PFB-HF (harder feedback).

Method

Participants Eight participants were recruited from the University of Iowa community in accordance with the universities institutional review board. Participants were balanced with the NFB group from Experiment 2 based on age and sex (PFB-HF: average age = 20.88 ± 1.36 years, five females; NFB: average age = 24.13 ± 4.48 years, four females; CP: average age = 22.85 ± 4.69 years, five females). All

participants had normal or corrected-to-normal vision and attended all three sessions. All participants were paid \$10 per session.

Stimuli All stimuli were identical to Experiment 2.

Procedures All procedures were the same as the PFB group in Experiment 2, except that the probability of feedback was adjusted so that the PFB-HF group was more likely to receive feedback on harder trials (see [Supplementary Method](#) section).

Model-based analysis The modeling analysis for Experiment 3 was identical to Experiments 1 and 2.

Results

Proportion of feedback trials Figure 9 illustrates the distribution of feedback for the PFB-HF group across trial difficulty (left panel) and across perceptual space (right panel). To determine whether we equated the information received by participants between groups, we compared the proportion of feedback trials for the PFB-HF group in Experiment 3 and the NFB group in Experiment 2. Thus, we performed a three-way ANOVA using day, condition (NFB vs. PFB-HF), and trial difficulty (easy, medium, or hard) as factors, and percentage of feedback trials as our

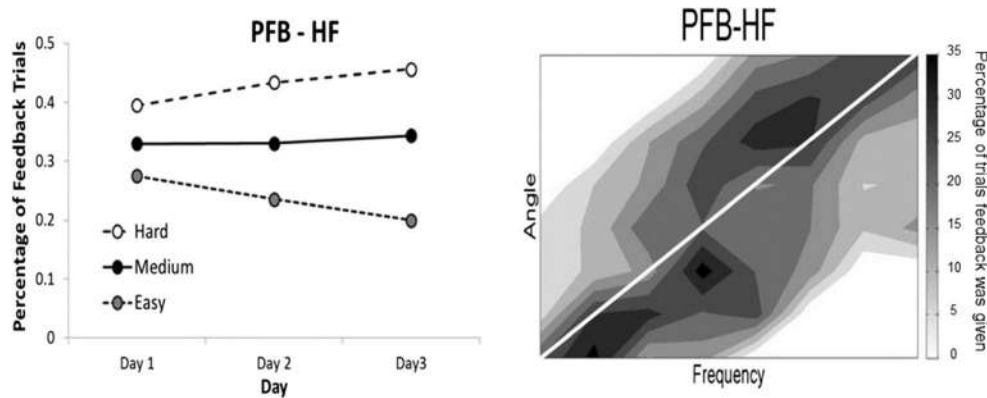


Fig. 9 Left panel. Percentage of trials plotted across days (x-axis) for hard, medium, and easy trials (separate lines). Right panel. Percentage of feedback plotted across perceptual space for PFB-HF on Day 3. Darker

shades reflect a higher percentage of feedback given for those trial types. The white line represents the optimal linear bound

dependent variable. The ANOVA revealed a significant main effect of day, $F(2, 26) = 20.495, p < .001, \eta_p^2 = 0.612$, and trial difficulty, $F(2, 26) = 184.909, p < .001, \eta_p^2 = 0.934$, and a significant interaction between trial difficulty and day, $F(4, 52) = 13.081, p < .001, \eta_p^2 = 0.502$. No other main effects or interactions were significant. Critically, no significant interaction between condition and trial difficulty was revealed ($F < 1$), so we conclude that our algorithm successfully equated feedback across perceptual space for all groups (for a comparison of the distribution of feedback between the NFB and PFB-HF groups, see the upper right panel of Fig. 8 and the left panel of Fig. 9). Participants received feedback on approximately 20 % of all trials across all days.

Accuracy-based analysis We performed a pairwise Wilcoxon sign test on the thirty 80-trial blocks of the experiment between the PFB-HF group and all of the groups in Experiment 2. The pairwise analysis revealed a significant difference between CP and PFB-HF (sign test: $S = 28$ of 30 blocks, $p < .001$), NFB and PFB-HF (sign test: $S = 23$ of 30 blocks, $p < .01$), but not between PFB-HF and PFB (sign test: $S = 16$ of 30 blocks, $p = .856$). Similar to the original PFB group, the PFB-HF group experienced a gain in accuracy on Day 1 from Block 1 (47 %) to Block 10 (63 %), but this growth in accuracy did not increase by the end of the experiment on Block 10 of Day 3 (63 %).

Within- and between-day learning Within and between-day learning was contrasted between all four groups, similar to Experiment 2. The within-day analysis revealed a significant effect of day, $F(2, 56) = 10.883, p < .001, \eta_p^2 = 0.280$, but no effect of group and no interaction (F s < 1). The between-day analysis revealed no significant effect of day, $F(1, 28) = 1.107, ns$, and a marginally significant interaction, $F(3, 28) = 2.658, p = .07, \eta_p^2 = 0.222$. We also identified a significant effect of group, $F(3, 28) = 5.079, p < .01, \eta_p^2 = 0.352$. Post hoc tests revealed group differences between NFB and PFB ($p < .01$), NFB and PFB-HF ($p < .05$), but no other comparisons were significant (p s

$> .24$). These results reveal that although both groups experienced similar online changes, the NFB group experienced stronger offline changes compared to the PFB and PFB-HF groups.⁷

Model-based analysis The data for the PFB-HF group were modeled similarly to the previous groups. This analysis revealed that no participant’s data in the PFB-HF group was best modeled by an II strategy: five used a unidimensional strategy and three used a conjunctive strategy. Thus, although equating information between groups increased the number of participants using a conjunctive strategy from one (in Experiment 2) to three, we still did not see more than one participant engage an II strategy among this group. Note also that the model probabilities are high for the PFB-HF group (see Fig. 7).

Discussion

Experiment 3 revealed that equating information was not sufficient to eliminate performance differences between groups. After roughly equating the regions of perceptual space that received feedback, we still observed a difference between the PFB-HF and NFB group in terms of accuracy achieved across days. In addition, we only identified a single PFB participant who was able to engage an II strategy across 21 total PFB participants compared to 7 out of 13 NFB participants who adopted an II strategy. These results provide further evidence that negative feedback is significantly more effective for teaching II categories than positive feedback.

General Discussion

This study demonstrates a clear advantage for negative feedback over positive feedback for II category learning. This conclusion is supported by higher accuracy for the NFB over

⁷ Note that this pattern of results is the same when directly comparing the PFB-HF and NFB groups.

the PFB group and greater use of II strategies in groups that were given negative feedback. In addition, we observed that the advantage for the NFB group was driven by between-session changes in accuracy, rather than within-session changes. These findings remained robust when the information that each group received was equated.

These results contrast with Ashby and O'Brien's (2007) finding that there was no difference in the effectiveness between positive and negative feedback. One possible reason for the disparity relates to the category structure. Unlike Ashby and O'Brien, we used categories that were nonoverlapping and excluded trials further than 0.3 diagonal units perpendicular to the optimal linear bound. Overlapping category structures lead to feedback that is inconsistent with the optimal bound, and this may be particularly detrimental to performance when the overall rate of feedback is low. Feedback on trials that are far from the bound may provide little information about the location of the bound, thereby diluting the proportion of trials on which useful information was given. In the case of limited feedback (such as in the PFB and NFB groups), these conditions are likely to promote the continued use of an RB strategy. By increasing the gap between the optimal accuracy using an II strategy versus using an RB strategy from 82 % to 100 %, we promoted conditions necessary to see a difference between these groups. These changes were sufficient to reveal a distinct advantage for negative feedback over positive feedback in promoting II learning.

Possible Mechanisms

One might propose that negative feedback is necessary to engage an II strategy because negative feedback signals the need to update the current categorization strategy, whereas positive feedback does not signal the need to change strategy. In other words, negative feedback may present a global signal that the current strategy being used is incorrect on top of the signal that the trial was performed incorrectly. For example, participants in the NFB group who may have used a unidimensional or conjunctive strategy early on may have realized that their current strategy was inadequate, leading to a strategic shift. Thus, negative feedback may signal the need to break out of an inadequate rule-based strategy. In contrast, in the absence of negative feedback, the PFB group may have assumed that their strategy was adequate, resulting in acquiescence to inferior performance. This may explain why participants in the PFB group failed to engage an II strategy.

Another possible explanation for our pattern of results is that negative feedback is required to unlearn incorrect associations between a stimulus and a response that were formed early on during training. In other words, it is possible that when an incorrect response is produced, an association is formed between a stimulus and a response. Evidence for this comes from Wasserman, Brooks, and McMurray (2015), who demonstrated

that pigeons can learn to categorize stimuli into multiple categories over the course of thousands of trials. On each trial, the pigeons were shown a target stimulus and a distractor stimulus and were cued to categorize the stimulus into one of 16 categories. Interestingly, the researchers noted that pigeons were less accurate when the current trial display included a distractor that had been rewarded as a target on the previous trial. This suggests that the pigeons had difficulty suppressing responses to stimuli that had just been rewarded. They posited that associative learning benefited from "pruning" incorrect associations, as well as the formation of correct associations. Based on our experiments, it is reasonable to suggest that the PFB group formed many incorrect associations early on in training, but that those associations were never corrected in the absence of negative feedback. If this was the case, it may imply that the positive feedback group had difficulty pruning these incorrect or irrelevant associations, which may explain the pattern of results we observed in our experiments.

A final possibility is that positive feedback on correct trials may not have been as informative as negative feedback on incorrect trials in our experiments. Typically, in a two-choice categorization task, positive and negative feedback are equally informative; positive feedback indicates that the response was correct whereas negative feedback indicates that the alternate option was correct. Nonetheless, it is possible that as training proceeded, positive feedback became less useful than negative feedback. This is because positive feedback may predominantly include information about trials the participant has already mastered. In contrast, negative feedback always indicates information about trials that participants are unsure about, or at least erred on. This may explain the differences in the pattern of information the PFB and NFB group received during Experiment 2. However, a problem with this explanation is that even when biasing feedback towards harder trials (the PFB-HF group in Experiment 3), which is presumably where one would find feedback most informative, we still observed a significant advantage for the NFB group over the PFB-HF group. Future research will be needed to disambiguate which possibility best explains our pattern of results.

Our analyses pointed to a potential mechanism to explain the advantage for the NFB group: differences in offline changes between groups rather than online changes. Although we did identify that this was the key difference between groups, we did not specifically isolate the locus of this effect (time away from the task, sleep-dependent consolidation, etc.). Thus, we are unable to give a precise reason why offline changes were greater for the NFB group over the PFB groups. Generally, however, it is assumed that offline processes do not involve any intentional shifts in categorization strategy. Thus, the between-days advantage suggests that negative feedback affords the engagement of incidental learning (e.g., learning not guided by intention) processes between days, that positive feedback cannot. Future work will be needed to determine if this is the case.

Conclusion

Our objective was to investigate the effectiveness of positive and negative feedback toward promoting II learning. Contrary to previous findings (e.g., Ashby & O'Brien, 2007), we demonstrated a stronger advantage for negative feedback over positive feedback. We observed higher accuracies as well as the successful engagement of II strategies for the negative feedback group whereas only one participant in the PFB group was able to engage an II strategy. These results were observed even after equating the information that was received between groups. In addition, although online changes were similar between groups, stronger offline changes were observed for participants that received negative feedback compared to those that received positive feedback. These results suggest that negative feedback may act as a more effective signal for teaching II categories.

Acknowledgements The authors would like to thank Darrell Worthy for his help with the modeling analyses. Additionally, we would like to thank Bob McMurray and Gavin Jenkins for their insight with respect to the interpretation of our findings. This research was supported in part by a National Institute on Drug Abuse under Award Number DA032457 to W. T. M. and 5R03DA031583-02 to E. H.

References

- Abe, M., Schambra, H., Wassermann, E. M., Luckenbaugh, D., Schweighofer, N., & Cohen, L. G. (2011). Reward improves long-term retention of a motor memory through induction of offline memory gains. *Current Biology*, *21*(7), 557–562.
- Ashby, F. G., & Maddox, W. T. (2011). Human category learning 2.0. *Annals of the New York Academy of Sciences*, *1224*(1), 147–161.
- Ashby, F. G., Maddox, W. T., & Bohil, C. J. (2002). Observational versus feedback training in rule-based and information-integration category learning. *Memory & Cognition*, *30*(5), 666–677.
- Ashby, F. G., & O'Brien, J. B. (2005). Category learning and multiple memory systems. *Trends in Cognitive Sciences*, *9*(2), 83–89.
- Ashby, F. G., & O'Brien, J. R. B. (2007). The effects of positive versus negative feedback on information-integration category learning. *Perception & Psychophysics*, *69*(6), 865–878.
- Brackbill, Y., & O'Hara, J. (1957). *Discrimination learning in children as a function of reward and punishment*. Eugene: Western Psychol. Ass.
- Donkin, C., Newell, B. R., Kalish, M., Dunn, J. C., & Nosofsky, R. M. (2015). Identifying strategy use in category learning tasks: A case for more diagnostic data and models. *Journal of Experimental Psychology: Learning, and Cognition*, *41*(4), 933–949.
- Dunn, J. C., Newell, B. R., Kalish, M. L., & 4. (2012). The effect of feedback delay and feedback type on perceptual category learning: The limits of multiple systems. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *38*, 840.
- Edmunds, C. E. R., Milton, F., & Wills, A. J. (2015). Feedback can be superior to observational training for both rule-based and information-integration category structures. *The Quarterly Journal of Experimental Psychology*, *68*(6), 1203–1222.
- Filoteo, J. V., Maddox, W. T., Salmon, D. P., & Song, D. D. (2005). Information-integration category learning in patients with striatal dysfunction. *Neuropsychology*, *19*(2), 212.
- Frank, M. J., Seeberger, L. C., & O'Reilly, R. C. (2004). By carrot or by stick: Cognitive reinforcement learning in parkinsonism. *Science*, *306*(5703), 1940–1943.
- Freedberg, M., Schacherer, J., & Hazeltine, E. (2016). Incidental learning of rewarded associations bolsters learning on an associative task. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *42*(5), 786–803.
- Galea, J. M., Mallia, E., Rothwell, J., & Diedrichsen, J. (2015). The dissociable effects of punishment and reward on motor learning. *Nature Neuroscience*, *18*(4), 597–602.
- Kubaneck, J., Snyder, L. H., & Abrams, R. A. (2015). Reward and punishment act as distinct factors in guiding behavior. *Cognition*, *139*, 154–167.
- Maddox, W. T., & Ashby, F. G. (2004). Dissociating explicit and procedural-learning based systems of perceptual category learning. *Behavioural Processes*, *66*(3), 309–332.
- Maddox, W. T., Ashby, F. G., & Bohil, C. J. (2003). Delayed feedback effects on rule-based and information-integration category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *29*, 650.
- Maddox, W. T., & Ing, A. D. (2005). Delayed feedback disrupts the procedural-learning system but not the hypothesis-testing system in perceptual category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *31*, 100.
- Maddox, W. T., Love, B. C., Glass, B. D., & Filoteo, J. V. (2008). When more is less: Feedback effects in perceptual category learning. *Cognition*, *108*(2), 578–589.
- Meyer, W. J., & Offenbach, S. I. (1962). Effectiveness of reward and punishment as a function of task complexity. *Journal of Comparative and Physiological Psychology*, *55*(4), 532.
- Newell, B. R., Dunn, J. C., & Kalish, M. (2011). 6 Systems of Category Learning: Fact or Fantasy?. *Psychology of Learning and Motivation-Advances in Research and Theory*, *54*, 167.
- Nikooyan, A. A., & Ahmed, A. A. (2015). Reward feedback accelerates motor learning. *Journal of Neurophysiology*, *113*(2), 633–646.
- Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, *6*(2), 461–464.
- Spiering, B. J., & Ashby, F. G. (2008). Initial training with difficult items facilitates information integration, but not rule-based category learning. *Psychological Science*, *19*(11), 1169–1177.
- Stanton, R. D., & Nosofsky, R. M. (2007). Feedback interference and dissociations of classification: Evidence against the multiple-learning-systems hypothesis. *Memory & Cognition*, *35*(7), 1747–1758.
- Tharp, I. J., & Pickering, A. D. (2009). A note on DeCaro, Thomas, and Beilock (2008): Further data demonstrate complexities in the assessment of information-integration category learning. *Cognition*, *111*(3), 410–414.
- Thorndike, E. L. (1911). *Animal intelligence: Experimental studies*. New York: Macmillan.
- Wächter, T., Lungu, O. V., Liu, T., Willingham, D. T., & Ashe, J. (2009). Differential effect of reward and punishment on procedural learning. *The Journal of Neuroscience*, *29*(2), 436–443.
- Wagenmakers, E. J., & Farrell, S. (2004). AIC model selection using Akaike weights. *Psychonomic Bulletin & Review*, *11*(1), 192–196.
- Waldron, E. M., & Ashby, F. G. (2001). The effects of concurrent task interference on category learning: Evidence for multiple category learning systems. *Psychonomic Bulletin & Review*, *8*(1), 168–176.
- Wasserman, E. A., Brooks, D. I., & McMurray, B. (2015). Pigeons acquire multiple categories in parallel via associative learning: A parallel to human word learning? *Cognition*, *136*, 99–122.
- Yechiam, E., & Hochman, G. (2013). Losses as modulators of attention: Review and analysis of the unique effects of losses over gains. *Psychological Bulletin*, *139*(2), 497.