Semisupervised category learning: The impact of feedback in learning the information-integration task

KATLEEN VANDIST, MAARTEN DE SCHRYVER, AND YVES ROSSEEL
Ghent University, Ghent, Belgium

In a standard supervised classification paradigm, stimuli are presented sequentially, participants make a classification, and feedback follows immediately. In this article, we use a semisupervised classification paradigm, in which feedback is given after a prespecified percentage of trials only. In Experiment 1, feedback was given in 100%, 0%, 25%, and 50% of the trials. Previous research reported by Ashby, Queller, and Berretty (1999) indicated that in an information-integration task, perfect accuracy was obtained supervised (100%) but not unsupervised (0%). Our results show that in both the 100% and 50% conditions, participants were able to achieve maximum accuracy. However, in the 0% and the 25% conditions, participants failed to learn. To discover the influence of the no-feedback trials on the learning process, the 50% condition was replicated in Experiment 2, substituting unrelated filler trials for the no-feedback trials. The results indicated that accuracy rates were similar, suggesting no impact of the no-feedback trials on the learning process. The possibility of ever learning in a 25% setting was also researched in Experiment 2. Using twice as many trials, the results showed that all but 2 participants succeeded, suggesting that only the total number of feedback trials is important. The impact of the semisupervised learning results for ALCOVE, COVIS, and SPEED models is discussed.

In category learning, there are two supervised learning paradigms: the supervised classification learning paradigm, and the supervised observational learning paradigm. Current research is dominated by the supervised *classification* learning paradigm. The typical ingredients of this paradigm are that (1) participants know in advance the number of contrasting categories; (2) stimuli are presented one at a time; (3) participants classify each stimulus; and (4) feedback (showing the correct category label of the stimulus) follows immediately. (See Shepard, Hovland, & Jenkins, 1961, for a description of the basic experiment.) Numerous studies have demonstrated that, using the supervised classification learning paradigm, participants can learn most kinds of complex category tasks (e.g., Ashby & Gott, 1988; Ashby & Maddox, 1990, 1992; Ashby, Queller, & Berretty, 1999; Medin & Schwanenflugel, 1981). In the supervised *observational* learning paradigm, participants are shown the category label of each stimulus prior to its presentation, and simply confirm the category label with an appropriate buttonpress (see Ashby, Maddox, & Bohil, 2002, for a description of the basic experiment). Learning performance is tested in test trials, in which neither feedback nor the category label is given. Results indicate successful learning in simple and in complex category tasks (Cincotta & Seger, 2007). However, using complex category tasks, higher accuracy rates are obtained in supervised feedback learning when training progresses (Ashby et al., 2002).

Few studies have focused on *unsupervised category* learning. In this setting, neither feedback (as in the supervised classification learning paradigm) nor category labels (as in the supervised observational learning paradigm) are given. Unsupervised learning can be intentional or incidental. In incidental learning, the main task does not require any category responses; instead, participants are asked to rate the pleasantness of the stimuli, for example, or to judge the relative positions of the stimuli. Yet these interactions with the stimuli can result in category formation, even though participants may have never received category information (for more details, see Love, 2002, 2003). By contrast, the main task in intentional category learning is the categorization of the stimuli. Several paradigms have been developed to study intentional unsupervised category learning. Perhaps the most used paradigm is the free sorting task, in which stimuli are presented simultaneously. The task of the participants is to order the stimuli in a way that seems natural to them (see, e.g., Billman & Davies, 2005; Lassaline & Murphy, 1996; Medin & Wattenmaker, 1987; Spalding & Murphy, 1996). A second paradigm that we adopted in our article is the *unsupervised classification learning paradigm*. The procedure of this paradigm is similar to the supervised classification learning paradigm: Participants are informed of the number of categories, stimuli are shown sequentially, and participants have to classify each stimulus. However, feedback is never given. The main

---

K. Vandist, katleen.vandist@ugent.be

---

results for both unsupervised learning paradigms are that performance is dominated by the use of unidimensional rules, regardless of the underlying category task (Ashby et al., 1999; Lassaline & Murphy, 1996; Medin & Wattemaker, 1987; Spalding & Murphy, 1996). These unidimensional rules are easy to verbalize and to apply (e.g., the small stimuli in Category A and the long stimuli in Category B).

Learning in the supervised and in the unsupervised classification learning paradigm is, among other category tasks, investigated in the information-integration task (Ashby et al., 1999). This category task is frequently used in research (e.g., Ashby & Ell, 2001; Ashby & Gott, 1988; Ashby & Maddox, 1990, 1992, 2005; Ashby & O'Brien, 2007; Ashby & Waldron, 1999; Ell & Ashby, 2006). The defining characteristic of the information-integration task is that participants have to combine the perceptual information of both underlying stimulus dimensions simultaneously to obtain good performance. This is difficult to accomplish, because the optimal decision bound has no verbal analogue (e.g., Ashby, Alfonso-Reese, Turken, & Waldron, 1998; Ashby et al., 2002; Maddox, Ashby, Ing, & Pickering, 2004). Using an information-integration task, Ashby et al. (1999) found that in the supervised condition, participants were able to master the task. In the unsupervised condition, however, learning failed, due to the use of unidimensional rules.

In our view, neither supervised nor intentional unsupervised learning seems ecologically plausible. Translated into daily life, supervised learning implies that our environment acts as an always available, completely reliable, stimulus-by-stimulus teacher. Consequently, for every object we observe, there is somebody or something telling us what it is. In addition, this information will always be unambiguously correct. Even in the most scholastic forms of learning, total feedback (as in supervised feedback learning) or total information (as in supervised observational learning) is improbable. On the other hand, unsupervised learning implies that we never receive any information about the categories. So during our entire lives, nobody would ever tell us the name of an object, or which objects belong together: again, a very unlikely situation. In our opinion, the supervised and the unsupervised classification learning paradigms are two extremes of a continuum, and we believe that the daily reality falls somewhere in between. When confronted with (new) objects, sometimes category information is received, and sometimes it is not. Therefore, we propose a semisupervised category learning paradigm¹ in which the prespecified percentage of trials with feedback can range from 0% (resulting in an unsupervised classification learning paradigm) to 100% (resulting in a supervised classification learning paradigm). Category responses following feedback will be referred to as "feedback trials," whereas in "no-feedback trials," feedback will not be provided. For instance, a 25% semisupervised classification learning paradigm consists of 25% feedback trials and 75% no-feedback trials.

This investigation fits in with the current renewed interest in the impact of feedback in category learning. In the Maddox, Ashby, and Bohil (2003) study, for example, the effect of delayed feedback was investigated.

The results showed impaired learning in the information-integration task, whereas learning in the simple category task was unaffected. Furthermore, Ashby and O'Brien (2007) examined the efficacy of positive and negative feedback using the information-integration task. Positive feedback signals that the current categorization response was correct, whereas negative feedback indicates the opposite. In the full condition, the supervised classification learning paradigm was applied. In the negative feedback condition, feedback followed after 80% of the incorrect responses. In the positive feedback condition and in the partial condition (with positive and negative feedback), the number of feedback trials was manipulated to be equal to the negative feedback condition; this resulted in approximately 25% of the 3,000 trials. To deal with the higher need of feedback at the start of learning, more feedback trials were given in the first blocks. The results suggest that both positive and negative feedback are needed for optimal learning in the information-integration task. Interestingly, comparing learning performance after 750 feedback trials in the supervised and the partial conditions reveals that higher accuracy was obtained in the partial condition. Ashby and O'Brien see two possible explanations for this result: First, in the partial condition, participants benefit from the no-feedback trials and learn from them; second, more between-session consolidation could occur, due to the distribution of the feedback trials over 5 different days.

Although Ashby and O'Brien (2007) show in the partial condition that semisupervised category learning can be successful, some questions remain unanswered. For example, it remains unclear whether more feedback trials are indeed needed in the beginning of learning (as in the Ashby & O'Brien [2007] article) or not (as investigated in this article). Second, we address what contributes to successful semisupervised category learning: Is the percentage of feedback trials decisive? Or is there a need for a minimum number of feedback trials? Finally, a third goal deals with the role of the no-feedback trials in the learning process. Are no-feedback trials helpful for learning, as suggested by Ashby and O'Brien? Do they interfere with the learning process? Or are the no-feedback trials just irrelevant and can they simply be ignored?

EXPERIMENT 1

The objective of Experiment 1 was twofold. First, we aimed to replicate the findings of Ashby et al. (1999) in a 100% condition and in a 0% condition. Second, the semisupervised classification learning paradigm was tested in 25% and 50% conditions to find out whether these particular amounts of feedback are sufficient for successful learning.

Method

Participants. In total, 34 participants (26 women; average age, 22 years) took part in the experiment in return for a small payment.

Design. Experiment 1 consisted of four between-subjects conditions (9 participants in the 100% and 25% conditions; 8 in the 0% and 50% conditions). Two conditions, the 100% and the 0%, were replications of Ashby et al. (1999). In the 100% condition, the super-

Figure 1. Two examples of Gabor patches.

vised classification learning paradigm was used; in the 0% condition, the unsupervised classification learning paradigm was used. In the other two conditions, the semisupervised classification learning paradigm was used. Note that the name of the condition indicates the percentage of feedback trials given. This percentage of feedback trials was the same in all the blocks. The order of feedback trials and no-feedback trials was randomized within blocks.

Stimuli and Apparatus. In all conditions, gray 300×300 square-pixel Gabor patches were shown on a black screen (two examples can be seen in Figure 1). In this study, the “gratings” varied continuously on two dimensions: the spatial frequency and the spatial orientation. Frequency and orientation are perceptually separable dimensions. The arbitrary stimulus coordinates were converted to physical units using the following transformation: spatial frequency (in cycles/pixel) was converted using $f(y) = 0.01 + (y/1,500)$, with y varying between 0 and 100. Spatial orientation was referred to in x° , varying between 0° and 100° . These coordinates originated in the information-integration task, as displayed, for example, in Figure 2. The optimal decision bound, which classified the stimuli perfectly in two categories, was diagonal. Participants viewed the stimuli on a 17-in. LCD monitor with an 800×600 resolution, at a distance of approximately one arm’s length.

For every participant, 800 stimuli were generated by randomly sampling from the two bivariate normal distributions, leading to 400 X stimuli and 400 N stimuli. Category X had a different mean than Category N did, but the variance and covariance remained the same. Due to the random sampling, the optimal decision bound varied slightly from block to block, although the mean optimal decision bound was $y = x$. The exact parameter values are shown in Table 1. As in the Ashby et al. (1999) study, the mean, the variance, and the covariance values were chosen in such a way that a linear decision bound based on one dimension would account for a maximum accuracy of 80%.

Procedure. All participants were tested individually in a dimly lit room. For every participant, 800 trials were divided into 10 blocks, so that an equal number of stimuli of each category was guaranteed in every block. Depending on the condition, the procedure differed. In the 100% condition, participants were told that they would see stimuli one by one, originating from two categories, X and N. They were asked to respond by pressing “X” on the keyboard if they believed that the stimulus was an X, and to press “N” when they believed that the stimulus was an N. Participants were informed

that they would receive feedback (i.e., the true category label) after each category response. At the end of each block, the percentage of the correct responses was shown; consequently, participants were encouraged to do better in the next block. At the beginning of the experiment, they were informed that it was possible to achieve maximum accuracy. A trial started when the stimulus was projected in the middle of the screen. After the participant responded, the feedback became visible under the stimulus for 1,500 msec, followed by an intertrial interval (ITI) of 1,000 msec.

The 0% condition is based on Experiment 1A of Ashby et al. (1999). Two types of blocks were adopted: observation (odd) blocks and response (even) blocks. In the observation block, participants had to watch the stimuli one by one, whereas the procedure with the response blocks was similar to that of the supervised condition, except that feedback was given neither after a category decision nor at the end of the block. The participants were requested to choose the category labels arbitrarily at the start of each response block, but, once a decision was made, participants had to be consistent throughout the entire block.

In the semisupervised conditions with 25% or 50% feedback, the participants got the same instructions as in the supervised condition, except that feedback might or might not follow the response. At the end of each block, the accuracy of the feedback trials was displayed in percentages. These feedback trials were identical to the 100% condition. In the no-feedback trials, the stimulus remained visible after the response for 1,500 msec, which was then followed by the ITI.

Results

Two types of analyses were performed. First, the mean accuracy over blocks was calculated for each condition. Second, model-based analyses were run to determine the category decision strategy used by the participant.

Accuracy analysis. Recall that all participants can achieve perfect accuracy when they use the optimal decision bound—that is, by integrating the information from the two stimulus dimensions at some predecisional stage. On the other hand, using a unidimensional decision rule, the accuracy amounts to a maximum of 80%. This implies that participants with a result of more than 80% probably adopted a (suboptimal) information-integration decision rule; hence, in this article, category learning is considered successful when a score higher than 80% is obtained. The accuracy analyses are based solely on the trials with feedback. The procedure for calculating the accuracy in the 0% condition was the same as in Ashby et al. (1999). In this condition, category assignment of the labels X and N was arbitrary; it was assumed that participants assigned the labels to obtain the highest percentage accuracy within each block. Consequently, for each participant, the classification resulting in the highest percentage accuracy was used; it was, therefore, impossible for participants to have an outcome of less than 50%.

Figure 3 shows for every condition the average percent correct of each (response) block. In Tables 2 and 3, the accuracy for every participant in the last two blocks is displayed. In the 100% and 50% conditions, learning was successful: The mean accuracy increased from an average of 70% ($SD = 13.86$) for the 100% condition and an average of 74% ($SD = 11.26$) for the 50% condition to almost perfect accuracy in the last block [92% ($SD = 6.64$) and 92% ($SD = 8.44$), respectively]. The individual accuracy rates show that all participants obtained a score higher than 80%,

Table 1
Parameter Values That Define the Categories

Category	Frequency		Orientation	
	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
X*	40	11.88	60	11.88
N*	60	11.88	40	11.88

Note—Average optimal slope, 1; average optimal intercept, 0. *Within-category correlation = .99.

Figure 2. Participant 1: In this figure, an example of the underlying category structure is shown. Stimuli of Category X are printed in gray squares, stimuli of Category N are printed in filled black circles. Figure 2 as a total: Category responses during the last two blocks for each participant in the 100% condition of Experiment 1. Also shown are the decision bounds of the best-fitting models. The y-axis is the orientation dimension; the x-axis, the frequency dimension.

except Participant 9 in the 100% condition, who failed to reach the learning criterion. Nevertheless, there is no indication that successful learning occurred in the 0% and 25% conditions, because the average percent correct stayed below the 80% accuracy level throughout all blocks. For the 0% condition, the average percent correct for Response Block 1 was 73% ($SD = 4.71$), and in the last response block, 72% ($SD = 7.09$). All participants scored less than 80%. In the 25% condition, the average percent correct for Block 1 was 61% ($SD = 16.48$), and it was 74% ($SD = 8.46$) for the last block. Participants 1 and 3, however, reached 80% and 88% accuracy, respectively, in the last two blocks.

Model-based analysis. Model-based analyses were run to determine the category decision strategy used by the participant. The category decision strategy for each participant was calculated on the last two response blocks. Four different decision-bound models were fit to these data (see the Appendix for details). Decision-bound models assume that participants divide the perceptual space into response regions by constructing a decision bound. For each stimulus, participants determine the region in the perceptual space, and this elicits the appropriate response (see, e.g., Ashby & Gott, 1988; Maddox & Ashby, 1993). Three of the tested models—the horizontal one-

Figure 3. Mean percent correct by block in each (semi)-supervised condition of Experiment 1.

dimensional classifier, the vertical one-dimensional classifier, and the general conjunctive classifier—were rule based (see, e.g., Maddox & Ashby, 2004). If participants adopted one of these category decision strategies, category learning failed. The fourth model, the general linear classifier, was information-integration based, and could approximate the optimal (diagonal) bound. Only with this category decision strategy could perfect accuracy, indicating successful learning, be obtained.

The participants' responses for the last two blocks, along with the decision bound of the best-fitting model, are shown in Figures 2, 4, 5, and 6 for the 100%, 0%, 25%, and 50% conditions, respectively. The BIC scores corresponding to the four decision bound models are depicted in Tables 2 and 3.

In the 100% condition, all participants but one used a strategy based on information-integration, indicating that they had learned the category task. In the 0% condition, all participants failed to learn the category task. Participants 2, 3, 6, and 7 used a one-dimensional decision bound based on the frequency, whereas Participant 5 used a one-dimensional decision bound based on the orientation. The responses of Participants 1 and 8 were best fitted by the general conjunctive classifier. In the 25% condition, several category decision strategies were used. Participants 2, 4, 8, and 9 applied a one-dimensional model based on the frequency, whereas the data of Participants 6 and 7 were best fitted with a one-dimensional model based on the orientation. Participants 1 and 5 used a general conjunctive strategy. These results suggest that they were not able to learn the category task. However, the responses of Participant 3 were best fitted by the general linear classifier. It seems that this participant learned the category task. Finally, in the 50% condition, all participants (except Participant 1) succeeded in learning the category task, because their responses were best fitted by the general linear classifier. These results suggest that it is possible to learn the

information-integration task when feedback is randomly given after only 50% of the responses.

Discussion

The objective of Experiment 1 was to replicate the findings of Ashby et al. (1999) using an information-integration structure in the 100% (supervised) and the 0% (unsupervised) conditions. The results show that this repli-

Table 2
Bayesian Information Criterion (BIC) Scores for All Participants in the 100% and the 0% Conditions of Experiment 1

Participant	GLC	DIM-O	DIM-F	GCC	% Correct*
100%					
1	<u>25.83</u>	166.91	110.63	60.61	99
2	<u>24.13</u>	147.84	155.14	58.97	99
3	<u>15.23</u>	150.51	138.83	43.45	100
4	128.15	197.47	175.16	<u>95.00</u>	88
5	<u>94.57</u>	197.33	170.35	119.04	93
6	<u>90.32</u>	210.50	114.38	97.70	91
7	<u>97.90</u>	190.43	156.52	100.30	93
8	<u>122.91</u>	205.07	166.58	134.20	89
9	<u>132.82</u>	133.88	241.41	137.77	79
0%					
1	115.19	223.72	115.25	<u>115.11</u>	69
2	71.48	231.53	<u>66.47</u>	71.54	78
3	106.54	234.78	<u>102.88</u>	106.62	72
4	76.41	220.86	82.23	<u>76.11</u>	73
5	37.75	<u>33.21</u>	237.39	38.29	79
6	108.45	226.52	<u>106.54</u>	107.64	71
7	106.13	227.69	<u>102.34</u>	106.15	72
8	146.73	168.30	195.94	<u>144.48</u>	56

Note—The BIC scores corresponding to the best-fitting words are underlined. GLC, general linear classifier; DIM-O, one-dimensional classifier, orientation; DIM-F, one-dimensional classifier, frequency; GCC, general conjunctive classifier. *Last two blocks.

Table 3
Bayesian Information Criterion (BIC) Scores for All Participants in the 25% and the 50% Conditions of Experiment 1

Participant	GLC	DIM-O	DIM-F	GCC	% Correct*
25%					
1	122.34	138.94	226.73	<u>120.79</u>	80
2	56.32	229.80	<u>51.38</u>	56.46	70
3	<u>91.58</u>	212.23	133.84	112.49	88
4	147.40	226.49	<u>145.03</u>	146.96	68
5	146.94	220.93	153.35	<u>145.02</u>	75
6	136.12	<u>131.16</u>	240.03	135.51	75
7	125.15	<u>120.10</u>	239.06	125.01	73
8	175.60	237.27	<u>171.88</u>	175.62	65
9	173.71	240.48	<u>168.71</u>	172.62	75
50%					
1	105.40	190.13	167.51	<u>102.69</u>	93
2	<u>29.81</u>	144.62	155.96	72.78	99
3	<u>128.79</u>	210.66	163.37	146.11	89
4	<u>50.35</u>	157.71	167.73	85.89	99
5	<u>100.86</u>	186.73	170.69	113.59	93
6	<u>93.03</u>	188.27	160.30	104.52	91
7	<u>86.15</u>	174.72	173.10	112.70	98
8	<u>57.76</u>	168.41	156.77	97.16	96

Note—The BIC scores corresponding to the best-fitting model are underlined. GLC, general linear classifier; DIM-O, one-dimensional classifier, orientation; DIM-F, one-dimensional classifier, frequency; GCC, general conjunctive classifier. *Last two blocks.

Figure 4. Category responses during the last two blocks for each participant in the 0% condition of Experiment 1. Also shown are the decision bounds of the best-fitting models. The y-axis is the orientation dimension; the x-axis, the frequency dimension.

cation is successful in both conditions. The mean accuracy of the 100% condition in the last block is close to perfect; moreover, all participants but one used an information-integration decision bound close to the optimal bound. By contrast, in the 0% condition, the mean accuracy stays far below 80%, and the model-based analyses show the ubiquitous use of rule-based decision strategies. These results indicate that participants succeeded in learning the category structure in the 100% condition, whereas in the 0% condition, participants failed to learn the category task.

In the semisupervised condition with 50% feedback, the high accuracy and the information-integration-based decision strategies of 8 participants indicate successful learn-

ing. This result suggests that feedback is not required after each response to learn the information-integration task, but participants apparently fail to learn when the feedback is limited to 25%. In this condition, the responses of 8 participants were best fitted by a rule-based decision bound.

EXPERIMENT 2

In Experiment 2, two follow-up conditions were run. First, we extended the 50% condition of Experiment 1 to study the effect of the no-feedback trials. Second, the possibility of ever learning in a 25% semisupervised classification learning paradigm was studied.

Figure 5. Category responses during the last two blocks for each participant in the 25% condition of Experiment 1. Also shown are the decision bounds of the best-fitting models. The y-axis is the orientation dimension; the x-axis, the frequency dimension.

Impact of the no-feedback trials. In the 50% condition of Experiment 1, participants could master the category task, although feedback was given only after half of the trials. Even so, the role of the no-feedback trials in the learning process remains unclear. Therefore, the influence, or lack thereof, of the no-feedback trials was studied by comparing the 50% condition of Experiment 1 with a new condition in which the no-feedback trials were replaced by unrelated filler trials. We envisioned three possible ways in which the no-feedback trials might influence the learning process. A first possibility is that the no-feedback trials benefited the learning

process, perhaps compensating for the lack of feedback on half of the trials. If this was indeed true, we expected learning to be slower in Experiment 2. A second possibility is that the no-feedback trials harmed the learning process. In this case, learning would be faster in Experiment 2. A final alternative was that the no-feedback trials did not interfere with the learning process at all. In this case, we would expect similar performance in both conditions.

25% feedback learning possibilities. When the 25% condition of Experiment 1 ended at Block 10, most participants seemed unable to learn the category structure.

Figure 6. Category responses during the last two blocks for each participant in the 50% condition of Experiment 1. Also shown are the decision bounds of the best-fitting models. The y-axis is the orientation dimension; the x-axis, the frequency dimension.

The reason for this failure could be that to learn this type of category task one needs more than 25% feedback trials. This hypothesis will be referred to as the *percentage-feedback hypothesis*. However, the responses of 1 participant were best fitted by an information-integration decision bound model (i.e., the general linear classifier). Participants might have been able to learn the category task but the experiment might have stopped too soon, leaving open the possibility that, with more blocks in total, the category task could be learned after all. This *absolute-number hypothesis* suggested that the absolute number of feedback trials is indicative for learning, and that learning would occur if participants received

more feedback trials in total. The objective of this condition was to find out which was correct: the percentage-feedback hypothesis or the absolute-number hypothesis. Therefore, the 25% condition of Experiment 1 was replicated, but this time the participants were subjected to two sessions of 800 trials.

Method

Participants. Seventeen participants (12 women; average age, 21.1 years) participated in the experiment in return for a small payment.

Procedure, Stimuli, and Apparatus. For the 50% feedback learning trials, the procedure, the apparatus, the stimuli, and the category task were identical to those of the 50% condition, as de-

scribed in Experiment 1. The difference between the 50% condition of Experiment 1 and the present condition was that Gabor patches were shown on half of the trials, on which, after response, subjects were provided with feedback (i.e., feedback trials). These trials were randomly alternated with unrelated filler trials in which the stimuli were the characters "X" (odd trials) and "N" (even trials). These fillers had the same size as the Gabor patches but did not disclose anything about the underlying category task or cue the outcome of the following feedback trial.

For the 25% feedback learning trials, the procedure, the stimuli, and the apparatus were exactly the same as those described in the 25% condition of Experiment 1, except that participants received two sessions of 800 trials. The two sessions took place on 2 consecutive days.

Results

Impact of the no-feedback trials. For this condition, only the feedback trials were analyzed. The accuracy analysis indicated successful learning: The average percent correct in the last block was 95% ($SD = 4.52$) (see Table 4 for the individual accuracy rates for the last two blocks). Figure 7 shows the average percent correct over the blocks. Note that the learning process observed was similar to that in the 50% condition of Experiment 1. Next, the best-fitting decision bounds models were calculated on all responses of the last two blocks (see Table 4 and Figure 8). The results show that all participants used an information-integration decision strategy.

25% feedback learning possibilities. The accuracy analysis suggested successful learning: The average percent correct in the last block was 96% ($SD = 6.82$) (see Table 4 for the individual accuracy rates for the last two blocks; see Figure 9 for the average percent correct over the blocks). Similarly, the model-based analyses revealed the use of information-integration decision strategies, except for Participants 2 and 9. The results are shown in Figure 10 and Table 4.

Discussion

In Experiment 2, two additional conditions were run for a better comprehension of semisupervised feedback learning. First, the effect of the no-feedback trials on the learning process was addressed. The results demonstrated no loss or gain in performance in this unrelated filler condition, as indicated by a high average percent correct, similar to the 50% condition of Experiment 1. Additionally, 8 participants appeared to have used an information-integration-based decision strategy. In conclusion, these results suggest that, when no feedback is given, observing and responding to relevant stimuli does not influence the learning process in this task. Second, the ambiguous results of the 25% condition of Experiment 1 were clarified. Two feasible hypotheses were tested: the percentage-feedback hypothesis and the absolute-number hypothesis. The results reveal that by giving two sessions instead of one, all but 2 participants mastered the category task. This suggests that the total number of feedback trials is indicative of learning or failure to learn. Interestingly, after one session, the accuracy rates were much higher than in the 25% condition. We believe that this result was due to the greater com-

mitment of the participants in Experiment 2. Perhaps because the participants knew in advance that they would return and perform the same task, they tried harder to learn the categories.

GENERAL DISCUSSION

In the category learning literature, the focus has been mainly on strictly supervised learning or on strictly unsupervised learning. In the supervised category learning paradigm, feedback followed immediately after each classification response, whereas in the unsupervised category learning paradigm, feedback was never provided. In this article, we argue that both these paradigms are ecologically implausible. In daily life, we sometimes get feedback about the category membership of an object, and sometimes we do not. It is important that category learning researchers take this into account. Therefore, in this article, a semisupervised category learning paradigm was used in which the percentage of feedback trials was manipulated. In order to facilitate comparison with earlier work, we used the information-integration task. A first finding is that participants were able to learn the information-integration task when feedback followed after half of the trials (Experiment 1). This finding is in line with the partial condition of Ashby and O'Brien (2007), but it is surprising, too, since Maddox et al. (2003) found that supervised category learning in the information-integration task was impaired when feedback was delayed. In the 50% condition, feedback was even omitted after half the trials, and still no interference in the learning process was observed. Second, we found that learning occurred even when feedback was provided after only a quarter of the trials, but only if a sufficient number of feedback trials were presented (Experiment 2). The major conclusion is that participants do not need feedback after every single response

Figure 7. Mean percent correct by block in the conditions with 50% feedback.

Figure 8. Category responses during the last two blocks for each participant in the impact of the no-feedback trials condition of Experiment 2. Also shown are the decision bounds of the best-fitting models. The y-axis is the orientation dimension; the x-axis, the frequency dimension.

to master the information-integration task, as long as the absolute number of feedback trials is sufficiently high. Additionally, since the percentage of feedback trials was equal across blocks, it seems that there is no need for a higher percentage of feedback trials in the first blocks in order to achieve successful learning, as suggested by Ashby and O'Brien. Further research is needed to find out whether the same conclusions hold for other complex category tasks and other (semisupervised) paradigms.

This article is the first in category learning literature to directly investigate the nature of the no-feedback trials during the learning process (Experiment 2). Clearly,

adopting filler trials instead of no-feedback trials did not influence learning. Although participants had to respond to the stimuli (and presumably this decision involved a lot of brain processing), they seemed to simply ignore these trials when no feedback followed. This finding rebutted the suggestion made by Ashby and O'Brien (2007) that no-feedback trials might benefit the learning process.

To the best of our knowledge, no category learning models make explicit predictions about semisupervised category learning (using continuous stimulus dimensions). In order to relate our results to the broader literature, three (supervised) category learning models are discussed: the

Table 4
Bayesian Information Criterion (BIC) Scores
for All Participants in Experiment 2

Participant	GLC	DIM-O	DIM-F	GCC	% Correct*
50% (Filler)					
1	<u>13.15</u>	69.64	79.80	29.36	100
2	<u>23.37</u>	82.58	60.02	34.33	99
3	<u>27.50</u>	67.37	95.18	40.52	98
4	<u>28.49</u>	81.01	73.60	41.92	98
5	<u>52.33</u>	139.38	81.48	79.64	94
6	<u>15.23</u>	141.02	75.68	27.95	100
7	<u>71.42</u>	136.12	79.12	78.63	83
8	<u>69.65</u>	164.88	96.67	84.76	89
25% (Two Sessions)					
1	<u>24.76</u>	128.15	173.95	68.71	98
2	<u>122.03</u>	<u>118.55</u>	247.00	122.64	63
3	<u>29.63</u>	180.74	143.99	50.49	95
4	<u>15.23</u>	144.26	140.52	58.08	100
5	<u>50.51</u>	171.07	154.19	91.82	95
6	<u>42.96</u>	167.00	149.23	68.80	95
7	<u>15.23</u>	178.55	123.10	70.69	100
8	<u>136.70</u>	183.97	201.43	146.91	88
9	84.51	232.20	<u>79.99</u>	85.01	63

Note—The BIC scores corresponding to the best-fitting model are underlined. GLC, general linear classifier; DIM-O, one-dimensional classifier, orientation; DIM-F, one-dimensional classifier, frequency; GCC, general conjunctive classifier. *Last two blocks.

Figure 9. Mean percent correct by block in the conditions with 25% feedback.

ALCOVE model, the COVIS model and the SPEED model. The ALCOVE (attention learning covering map) model, according to which learning is error driven, was developed by Kruschke (1992). Kruschke explicitly stipulated that “When there is no error, nothing changes” (p. 25). Due to the absence of any feedback in the no-feedback trials, it was structurally impossible to obtain an error. Consequently, and in line with our results, the ALCOVE model predicts no influence of the no-feedback trials on the learning process.

On the other hand, the competition between verbal and implicit systems (COVIS) model predicts interference of

the no-feedback trials. COVIS was developed by Ashby et al. (1998) and postulated two systems that compete during learning. Learning the information-integration task relies on the implicit system (e.g., Ashby et al., 1998; Ashby & Waldron, 1999), and only this system will be explained. Critical to learning are the synapses between the visual cortical cells and the medium spinal cells in the caudate. To strengthen these synapses and enhance learning, three factors are required: (1) a strong presynaptic activation, (2) strong postsynaptic activation, and (3) dopamine release. Factor 3, dopamine, serves as a reward-mediated signal and will, therefore, be released only when positive feedback is given. Although there is large empirical support for COVIS (for recent reviews, see Ashby & Maddox, 2005; Ashby & Valentin, 2005; Ell & Ashby, 2006; Maddox & Ashby, 2004; Maddox & Filoteo, 2005), the semisupervised conditions seem incompatible. According to COVIS, the feedback trials in which a correct classification was made would strengthen the synapses and boost learning. On the contrary, the absence of Factor 3 in the no-feedback trials and the negative feedback trials would reduce the strength of the active synapses, and learning would decrease (Ashby et al., 1998). To our knowledge, there is no specific prediction of the rate at which active synapses are enhanced (by positive feedback) or diminished (by absence of feedback or negative feedback); but at least a disturbance in learning may be expected in the semisupervised conditions. Especially in the 25% condition, where a theoretical maximum of 25% of the category responses can lead to strengthening and a minimum of 75% of the responses will result in a weakening of the synapses, the interference effects should be detectable; but successful learning did occur in the 25% condition of Experiment 2. Apparently, the weakening of the synapses in more than 75% of the trials did not impede learning. Moreover, when we compare the learning processes of the 50% and the 100% conditions of Experiment 1, no difference can be found, in spite of the fact that, in the 50% condition, more than half of the trials should have disturbed learning, according to COVIS. Finally, Experiment 2 demonstrated no advantage in the learning process when unrelated (hence, undisturbing) fillers were used.

Recently, the implicit learning system of COVIS was extended by the subcortical pathways enable expertise development (SPEED) model (Ashby, Ennis, & Spiering, 2007). Besides the cortical–striatal pathway that relies on the same three-factor learning as does the COVIS model, a cortical–cortical pathway is added. This pathway links the sensory association cortex directly to the premotor area, resulting in faster reaction times. In this pathway, two factors are needed to strengthen the synapses. First, there must be presynaptic activation; second, there must be postsynaptic activation. When the postsynaptic activation passes a certain threshold, the synapses are strengthened. On the other hand, when the postsynaptic activation fails to reach the threshold, the link between the active synapses is diminished. Consequently, when this cortical–cortical pathway is used, the lack of (positive) feedback in the no-feedback trials should not disturb learning. Nevertheless, there

Figure 10. Category responses during the last two blocks for each participant in the 25% feedback trials learning possibilities condition of Experiment 2. Also shown are the decision bounds of the best-fitting models. The y-axis is the orientation dimension; the x-axis, the frequency dimension.

might be one problem for SPEED to explain our results, probably because SPEED did not consider semisupervised learning. In the case when learning depends on the cortical–striatal pathway (three-factor learning), the no-feedback trial corresponds to the condition in which striatal activation could pass the threshold but dopamine stays below the baseline. This condition is described as “long-term depression” in SPEED (Ashby et al., 2007, p. 640) and results in weaker synapse strength. Since the main goal of SPEED is to simulate categorization automaticity, it is postulated that early in category learning the cortical–striatal pathway is dominant. Hence, early in

learning, feedback is crucial; when the learning process progresses, the role of the cortical–cortical pathway increases, and the importance of feedback diminishes (see Ashby et al., 2007, for more details). This implies that in the first blocks, the strengthening of the synapses (and, therefore, learning) is entirely dependent on the positive feedback trials. As with the COVIS model, the absence of positive feedback trials in the first blocks will reduce the active synapses and interfere with learning. Although the rate at which the synapses are enhanced or weakened is unknown, these competing mechanisms will complicate the enhancing of the cortical–cortical pathway. Con-

sequently, disturbance in learning is expected, at least in the first blocks. Again, no such interference effects could be observed, since the learning processes in the 50% condition of Experiment 1 and the 100% condition were similar. Besides, no advantage of the unrelated (and hence undisturbing) fillers could be found in the first blocks of Experiment 2.

However, it cannot be ruled out that the irrelevance of the no-feedback trials was specific to the information-integration task used in this article. Note that the stimulus dimensions were highly correlated, and that both the feedback trials and no-feedback trials were randomly chosen from this highly correlated bivariate normal distribution. Consequently, the no-feedback trials did not provide any new information about the category task. This constellation could be masking possible influences of the no-feedback trials on the learning process; in other words, if the no-feedback trials give extra information about the category task, they may have an impact on the learning process. To discover the nature of the no-feedback trials entirely, more research is needed.

AUTHOR NOTE

Correspondence concerning this article should be addressed to K. Vandist, Department of Data Analysis, Ghent University, H Dunantlaan 2, B-9000 Ghent, Belgium (e-mail: katleen.vandist@ugent.be).

REFERENCES

- ASHBY, F. G. (1992). Multidimensional models of categorization. In F. G. Ashby (Ed.), *Multidimensional models of perception and cognition* (pp. 449-483). Hillsdale, NJ: Erlbaum.
- ASHBY, F. G., ALFONSO-REESE, L. A., TURKEN, A., & WALDRON, E. (1998). A neuropsychological theory of multiple systems in category learning. *Psychological Review*, **105**, 442-481.
- ASHBY, F. G., & ELL, S. (2001). The neurobiology of human category learning. *Trends in Cognitive Sciences*, **5**, 204-210. doi:10.1016/S1364-6613(00)01624-7
- ASHBY, F. G., ENNIS, J., & SPIERING, B. (2007). A neurobiological theory of automaticity in perceptual categorization. *Psychological Review*, **114**, 632-656. doi:10.1037/0033-295x.114.3.632
- ASHBY, F. G., & GOTT, R. E. (1988). Decision rules in the perception and categorization of multidimensional stimuli. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **14**, 33-53.
- ASHBY, F. G., & MADDUX, W. T. (1990). Integrating information from separable psychological dimensions. *Journal of Experimental Psychology: Human Perception & Performance*, **16**, 598-612.
- ASHBY, F. G., & MADDUX, W. T. (1992). Complex decision rules in categorization: Contrasting novice and experienced performance. *Journal of Experimental Psychology: Human Perception & Performance*, **18**, 50-71.
- ASHBY, F. G., & MADDUX, W. T. (2005). Human category learning. *Annual Review of Psychology*, **56**, 149-178. doi:10.1146/annurev.psych.56.091103.070217
- ASHBY, F. G., MADDUX, W. T., & BOHIL, C. (2002). Observational versus feedback training in rule-based and information-integration category learning. *Memory & Cognition*, **30**, 666-677.
- ASHBY, F. G., & O'BRIEN, J. (2007). The effects of positive versus negative feedback on information-integration category learning. *Perception & Psychophysics*, **69**, 865-878.
- ASHBY, F. G., QUELLER, S., & BERRETTY, P. M. (1999). On the dominance of unidimensional rules in unsupervised categorization. *Perception & Psychophysics*, **61**, 1178-1199.
- ASHBY, F. G., & VALENTIN, V. (2005). Multiple systems of perceptual category learning: Theory and cognitive tests. In I. H. Cohen & C. Lefebvre (Eds.), *Handbook of categorization in cognitive science* (pp. 547-572). New York: Elsevier.
- ASHBY, F. G., & WALDRON, E. M. (1999). On the nature of implicit categorization. *Psychonomic Bulletin & Review*, **6**, 363-378.
- BILLMAN, D., & DAVIES, J. (2005). Consistent contrast and correlation in free sorting. *American Journal of Psychology*, **118**, 353-383.
- CINCOTTA, C., & SEGER, C. (2007). Dissociation between striatal regions while learning to categorize via feedback and via observation. *Journal of Cognitive Neuroscience*, **19**, 249-265.
- ELL, S. W., & ASHBY, F. G. (2006). The effects of category overlap on information-integration and rule-based category learning. *Perception & Psychophysics*, **68**, 1013-1026.
- KRUSCHKE, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review*, **99**, 22-44.
- LASSALINE, M., & MURPHY, G. (1996). Induction and category coherence. *Psychonomic Bulletin & Review*, **3**, 95-99.
- LOVE, B. C. (2002). Comparing supervised and unsupervised category learning. *Psychonomic Bulletin & Review*, **9**, 829-835.
- LOVE, B. C. (2003). The multifaceted nature of unsupervised category learning. *Psychonomic Bulletin & Review*, **10**, 190-197.
- MADDUX, W. T., & ASHBY, F. G. (1993). Comparing decision bound and exemplar models of categorization. *Perception & Psychophysics*, **53**, 49-70.
- MADDUX, W. T., & ASHBY, F. G. (2004). Dissociating explicit and procedural-learning based systems of perceptual category learning. *Behavioural Processes*, **66**, 309-332. doi:10.1016/j.beproc.2004.03.011
- MADDUX, W. T., ASHBY, F. G., & BOHIL, C. (2003). Delayed feedback effects on rule-based and information-integration category learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **29**, 650-662.
- MADDUX, W. T., ASHBY, F. G., ING, A., & PICKERING, A. (2004). Disrupting feedback processing interferes with rule-based but not information-integration category learning. *Memory & Cognition*, **32**, 582-591.
- MADDUX, W. T., & FILOTEO, J. (2005). The neuropsychology of perceptual category learning. In I. H. Cohen & C. Lefebvre (Eds.), *Handbook of categorization in cognitive science* (pp. 573-599). New York: Elsevier.
- MEDIN, D., & SCHWANENFLUGEL, P. (1981). Linear separability in classification learning. *Journal of Experimental Psychology: Human Learning & Memory*, **7**, 355-368.
- MEDIN, D., & WATTENMAKER, W. (1987). Family resemblance, conceptual cohesiveness and category construction. *Cognitive Psychology*, **19**, 242-279.
- SCHWARZ, G. (1978). Estimating dimension of a model. *Annals of Statistics*, **6**, 461-464.
- SHEPARD, R., HOVLAND, C., & JENKINS, H. (1961). Learning and memorization of classifications. *Psychological Monographs*, **75**, 1-42.
- SPALDING, T., & MURPHY, G. (1996). Effects of background knowledge on category construction. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **22**, 525-538.

NOTE

1. Unknown during the data collection of this study, the semisupervised classification learning paradigm was simultaneously used in the partial condition by Ashby and O'Brien (2007).

APPENDIX

Four models were fit to each participant's data: three rule-based decision bound models (i.e., the vertical and the horizontal unidimensional models, the general conjunctive classifier), and one information-integration model (i.e., the general linear classifier). More details of the models can be found in Ashby (1992), Ashby and Gott (1988), and Maddox and Ashby (1993).

Rule-Based Models

The one-dimensional classifier (DIM) assumes that participants set a decision criterion on one of the stimulus dimensions. Because our stimuli have two stimulus dimensions, orientation and frequency, two one-dimensional classifiers were fit, one for the frequency dimension (DIM-F), and one for the orientation dimension (DIM-O). An example of a one-dimensional rule used for categorization is: "Respond A if the orientation is steep; otherwise, respond B." These models have two parameters: a decision criterion along the chosen classification dimension, and a perceptual noise variance parameter.

The general conjunctive classifier (GCC) assumes that the rule used by participants is a conjunction, for example: "Respond A if the orientation is steep and the frequency small; otherwise, respond B." The GCC has three parameters: a criterion on each dimension (i.e., orientation and frequency) and one parameter for the perceptual noise variance.

Information-Integration Model

The general linear classifier (GLC) assumes that a linear decision bound divides the stimulus space into response regions. Confronted with a stimulus, the perceived location in the stimulus space is determined and the corresponding categorization response elicited. These decision bounds require linear integration of both stimulus dimensions, resulting in an information-integration decision strategy. The GLC has three parameters: the slope and the intercept of the linear decision bound, and a perceptual noise variance.

Selecting the Best-Fitting Model

The model parameters were estimated using the method of maximum likelihood. To select the best-fitting model for every participant's set of responses, the model with the smallest Bayesian information criterion (BIC) is chosen. The BIC penalizes for the number of free parameters and is defined as $BIC = r \cdot \ln N - 2 \cdot \ln L$, where r is the number of free parameters, N is the sample size, and L is the likelihood of the model, given the data (Schwarz, 1978). These BIC scores are shown in Tables 2, 3, and 4. The BIC scores corresponding to the best-fitting model are underlined.

(Manuscript received July 15, 2007;
revision accepted for publication August 20, 2008.)