Perceptual dimensions influence auditory category learning

Casey L. Roark & Lori L. Holt

Springer

Springer

PERCEPTUAL/COGNITIVE CONSTRAINTS ON THE STRUCTURE OF SPEECH COMMUNICATION: IN HONOR OF RANDY DIEHL

CrossMark

Perceptual dimensions influence auditory category learning

Casey L. Roark[1] · Lori L. Holt[1]

Abstract

Human category learning appears to be supported by dual learning systems. Previous research indicates the engagement of distinct neural systems in learning categories that require selective attention to dimensions versus those that require integration across dimensions. This evidence has largely come from studies of learning across perceptually separable visual dimensions, but recent research has applied dual system models to understanding auditory and speech categorization. Since differential engagement of the dual learning systems is closely related to selective attention to input dimensions, it may be important that acoustic dimensions are quite often perceptually integral and difficult to attend to selectively. We investigated this issue across artificial auditory categories defined by center frequency and modulation frequency acoustic dimensions. Learners demonstrated a bias to integrate across the dimensions, rather than to selectively attend, and the bias specifically reflected a positive correlation between the dimensions. Further, we found that the acoustic dimensions did not equivalently contribute to categorization decisions. These results demonstrate the need to reconsider the assumption that the orthogonal input dimensions used in designing an experiment are indeed orthogonal in perceptual space as there are important implications for category learning.

Keywords Categorization · Perceptual categorization and identification · Audition

Introduction

Learning to treat distinct perceptual experiences as functionally equivalent is vital for perception, action, language, and thought. In the auditory domain, we can interpret a child's squeal as thrilled or terrified, judge a kettle to be boiling from its gurgle, or understand diverse acoustic signals from different talkers to each be the word "thanks." Auditory categories, whether speech or nonspeech, are most often complex and defined across multiple acoustic dimensions. Often times, these acoustic dimensions can be quite difficult to describe verbally or to attend to selectively (Francis, Baldwin, & Nusbaum, 2000; Grau & Kemler Nelson, 1988; Hillenbrand, Getty, Clark, & Wheeler, 1995).

Recently, an influential theory of category learning that was originally developed to explain visual category learning

has been expanded into the auditory domain (Chandrasekaran, Koslov, & Maddox, 2014; Chandrasekaran, Yi, & Maddox, 2014; Maddox, Molis, & Diehl, 2002). Many behavioral, neuropsychological, and neuroimaging studies have provided considerable evidence for the involvement of at least two distinct systems in visual category learning (Ashby & Maddox, 2005, 2011; Morrison, Reber, Bharani, & Paller, 2015; Smith & Grossman, 2008; but see Newell, Dunn, & Kalish, 2011). The Competition between Verbal and Implicit Systems (COVIS) model specifically posits the involvement of an explicit, hypothesis-testing system and an implicit, procedural-learning system (Ashby, Alfonso-Reese, Turken, & Waldron, 1998). The explicit, hypothesis-testing system is optimal for learning so-called "rule-based" (RB) categories that can be described with verbalizable rules. The kind of rules that are most often used in the literature vary across a single input dimension. However, rules can be based on multiple dimensions or require a more complex rule than can be easily verbalized. In the current study, we investigate RB categories based on one dimension, but acknowledge that rules can be more complex. This explicit system is thought to involve top-down processes and involves the prefrontal cortex as well as the head of the caudate nucleus in the striatum (Ashby &

✉ Lori L. Holt
loriholt@cmu.edu

1 Department of Psychology, Carnegie Mellon University, and the Center for the Neural Basis of Cognition, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA

Maddox, 2005). The implicit system learns via slower procedural learning mechanisms and is optimal for learning "information-integration" (II) categories that require integration across at least two input dimensions. The implicit system is thought to implement these processes by involving the body and tail of the caudate nucleus in the striatum as well as the putamen (Ashby & Maddox, 2005).

The expansion of this model into the auditory modality reveals some of the challenges in applying visual theories to audition (Roark & Holt, 2018). One issue concerns input dimensions. Whereas most visual category-learning studies have examined learning across simple input dimensions that are easily described verbally, acoustic dimensions like modulation frequency, or amplitude envelope, or the formant frequencies of speech, may be difficult for untrained listeners to describe (Francis et al., 2000; Hillenbrand et al., 1995). Additionally, the visual input dimensions that have been typically used in research tend to be perceptually separable, in that they are processed independently and are easy to attend to selectively (Garner, 1974). In contrast, acoustic dimensions are often integral; they are difficult to attend to selectively (Garner, 1974). Pitch and loudness, for example, are perceived integrally such that they are processed in a unitary fashion (Grau & Kemler Nelson, 1988; Melara & Marks, 1990). Although both auditory and visual dimensions can be separable or integral, a challenge in translating categorization models developed in the visual modality to auditory and speech category learning is that integral, interacting, and difficult-to-verbalize acoustic dimensions may differ from easily-verbalized categorization rules across separable dimensions in the visual modality, such as the frequently used spatial frequency and line orientation dimensions in a Gabor patch (Maddox, Ashby, & Bohil, 2003), or line length and orientation (Maddox, Filoteo, Lauritzen, Connally, & Hejl, 2005).

Across both visual and auditory tasks, integral or interacting dimensions have received much less attention than dimensions that are separable and easy to verbalize. Research on integral and separable dimensions demonstrates that dimensions are processed and used differently depending on how they are perceived by participants (Garner, 1976, 1978; Kemler & Smith, 1979). Whereas separable dimensions are easier to attend to selectively and thus, may benefit rule-based category learning, integral dimensions are more difficult to attend to selectively and may be detrimental for rule-based category learning. For instance, learning RB categories is more difficult across the integral dimensions of saturation and brightness than across separable dimensions like circle size and the angle of a radian inside the circle (McKinley & Nosofsky, 1996). However, other researchers found that participants had higher categorization accuracy when learning RB categories than II categories based on the integral dimensions of saturation and brightness (Ell, Ashby, & Hutchinson, 2012). The investigation of II and RB category learning across

integral dimensions has been limited and more research is needed to understand how the separable or integral nature of dimensions impacts II or RB category learning.

One study on auditory category learning with the integral dimensions of locations of spectral peaks in frequency space demonstrated that participants better learned II categories that required a negative integration in the stimulus space than a positive integration (Scharinger, Henry, & Obleser, 2013). However, these researchers did not compare this II category learning with typical RB category learning. Thus, it is not yet fully understood how auditory dimensions that are difficult to attend to selectively impact category learning across different category structures within the *same* acoustic space. It is necessary to examine acoustic dimensions that are integral and are difficult to verbalize to evaluate the ability of the COVIS perspective to accommodate the complexities of acoustic dimensions. The goal of the current study was to examine how categories within the same two-dimensional acoustic space are learned when selective attention to the dimensions is difficult and the dimensions are not easy to verbalize. This would allow us to investigate the applicability of the dual system perspective with complex acoustic dimensions that are similar to many complex acoustic dimensions that define auditory categories in real-world contexts, such as speech.

For acoustic dimensions that interact or are difficult to attend to selectively, we may expect that information-integration categories will be learned better than rule-based categories because the dimensions are difficult to separate perceptually. Likewise, the ability to learn different categories within this space may depend on the precise nature of the relationship between the acoustic dimensions and how this relates to internal *perceptual* representation of the dimensions, which may differ from the *acoustic* dimensions.

To test these predictions, we trained participants on auditory categories defined across two acoustic dimensions that past research suggests are difficult to attend to selectively (Holt & Lotto, 2006; Roark & Holt, 2018). We examined participants' accuracy across training and their ability to generalize to novel sounds. Additionally, we applied decision-bound computational models to assess participants' strategy use in category learning and their propensity to integrate or selectively attend to the dimensions in category decisions at each stage of learning.

Methods

We investigated two types of information-integration (II) category-learning problems and two types of unidimensional rule-based (RB) category-learning problems. Each of these category-learning challenges was defined across the same two acoustic input dimensions. Sampling this input space in four different ways allowed us to avoid assumptions about

which dimension, or combination of dimensions, would most impact learning. To anticipate, we found that participants tended to integrate across the dimensions, especially in a way that reflected a positive correlation between the dimensions.

Frequency-modulated nonspeech tones served as category exemplars across each of the four category-learning challenges. We used nonspeech stimuli to control as much as possible for participants' prior experience with this acoustic space. By using nonspeech sounds, we were able to carefully construct artificial categories and match different category exemplar distributions as much as possible. The acoustic input dimensions across which category exemplar distributions were sampled were center, or carrier, frequency (CF) and modulation frequency (MF). In a previous study of auditory category learning using these same dimensions (Holt & Lotto, 2006), participants were able to adjust perceptual weighting across the dimensions on the basis of what was required for the task. However, participants tended to place some weight on each dimension, even when the category-learning task required selective attention to a single dimension. In other words, selective attention to these acoustic dimensions is difficult, even when it is required by the task. We chose this particular pair of dimensions because perceptual reliance on the dimensions is malleable and, at the same time, the perceptual representation of these dimensions may not be entirely separable.

Participants

A total of 81 adults (38 females, 43 males) aged 18–24 years and affiliated with Carnegie Mellon University participated for partial course credit. Participants were randomly assigned to one of four conditions defined by the sampling of category exemplars in the acoustic input space. Three participants were excluded due to equipment error, leaving 78 subjects in the final analysis. There were 20 participants in the rule-based-CF (RB_{CF}) condition, 19 in the rule-based-MF (RB_{MF}) condition, 19 in the information-integration positive slope ($II_{Positive}$), and 20 in the information-integration negative slope ($II_{Negative}$) condition. All participants reported normal hearing.

Stimuli

Sound exemplars The two-dimensional acoustic space from which stimuli were sampled was defined by CF and MF. As in Holt and Lotto (2006), each stimulus was created from a sine wave tone with a particular CF modulated with a depth of 100 Hz at the corresponding MF. For example, if the CF was 760 Hz and the MF was 203 Hz, the tone was modulated from 710 to 810 Hz at a rate of 203 Hz. Each stimulus was 300 ms long. Exemplars were synthesized in

MATLAB (Mathworks, Natick, MA, USA) and matched for RMS energy.

Category distributions Two individual category distributions were created to define the category-learning challenge for each of the four conditions (Fig. 1). The information-integration conditions sampled acoustic space such that optimal performance would require integration across the two dimensions. The $II_{Positive}$ and $II_{Negative}$ conditions are mirror images, differing only in the nature of the correlation between CF, shown on the x-axis in Fig. 1, and MF, shown on the y-axis in Fig. 1. In the case of $II_{Positive}$, higher CF values were associated with higher MF values and, for $II_{Negative}$, higher CF values were associated with lower MF values. The rule-based categories sampled acoustic space such that they could be optimally differentiated by selectively attending to one of the two stimulus dimensions that define the categories. The RB_{CF} condition requires selective attention to the CF dimension. The RB_{MF} condition requires selective attention to the MF dimension.

Each category was defined by 100 distinct stimuli sampled from a bivariate normal distribution across the two input dimensions (Table 1). Half of the stimuli from each category were used during training and the other half were reserved for the generalization test. The exemplars defined as training and test were randomly selected, with consistent sampling across participants.

Procedure

Separate groups of listeners participated in each condition. The task was identical across conditions; only the sampling of stimulus distributions varied (Fig. 1). Participants were not informed about the nature of the dimensions.

During the training phase, participants completed four blocks of training (96 trials/block; 384 total trials), with a brief break in quiet separating blocks. The trial structure was largely the same across training and generalization phases of the experiment. On each trial, participants heard a single sound exemplar (300 ms) randomly selected from one of the two categories, repeated five times (50-ms silent inter-stimulus interval). Two boxes on the screen indicated response options corresponding to the "u" and "i" keys on a standard keyboard. Participants indicated which of two equally likely categories the sound belonged to by pressing a response button. A red X indicating the correct category decision appeared in one of the boxes 500 ms after response. Participants were instructed to use this feedback to inform future categorization decisions. A 1-s inter-trial interval followed the feedback.

After completing the training phase, participants completed the generalization test (100 trials). Participants were instructed that they would now be tested on what they learned during training and that there would no longer be
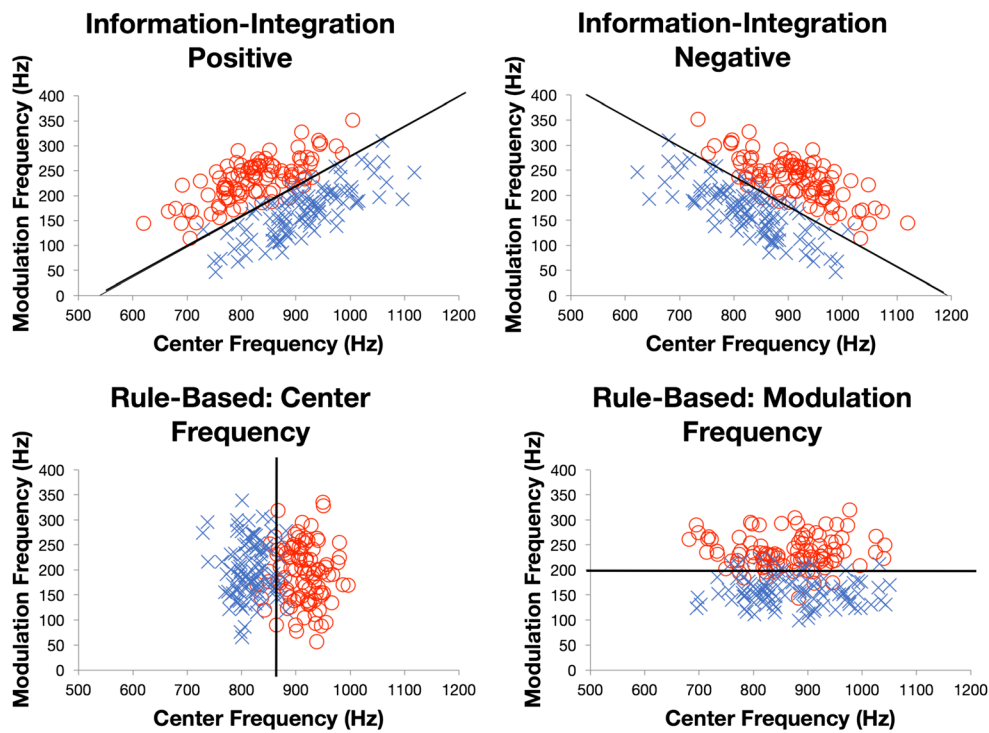
Fig. 1 Stimulus distributions for the four conditions in this study. The black line represents the optimal decision boundary that separates the two categories

feedback. Instead of a red X indicating the correct category decision, question marks appeared inside each of the two boxes on the screen. With the exception of the feedback, the trial structure was identical to the training phase. During the generalization test phase, participants encountered category exemplars (50/category, 100 total) that they had not encountered in training. Thus, the generalization

phase measured the ability to generalize category learning to novel exemplars – a hallmark of categorization.

The task was run in a sound-attenuated booth using E-Prime software (Psychology Software Tools, Inc., Sharpsburg, PA, USA), with stimuli presented diotically over Beyer DT-150 headphones at a comfortable listening level.

Results

The analyses focused on accuracy of categorization during training and as assessed in the generalization test. Additionally, we fit a series of decision-bound models to categorization responses across training in order to examine response strategies across conditions (for more detailed information about model applications see: Ashby & Maddox, 1993; Maddox & Ashby, 1993; Maddox & Chandrasekaran, 2014).

Behavioral results

Normalization Although we attempted to equate the stimulus distributions for RB and II for distance between the means and variance, there remained small differences in the overall overlap of the two categories across conditions. An ideal observer would be able to achieve 96% accuracy in the II conditions, 91% in the RB_{MF} condition, and 92% in the RB_{CF} condition. Thus, we computed a normalized accuracy score to account

Table 1 Category distribution information

	Mean (CF, MF)	Variance (CF, MF)	Covariance
Information-integration positive			
Category A	(818.3, 225.8)	(6328, 2278.2)	2260.5
Category B	(913.8, 168.3)	(6328, 2278.2)	2969.9
Information-integration negative			
Category A	(913.8, 225.8)	(6328, 2278.2)	-2260.5
Category B	(818.3, 168.3)	(6328, 2278.2)	-2969.9
Rule-based modulation frequency			
Category A	(866, 156.5)	(6770, 796.7)	53.4
Category B	(866, 237.5)	(6770, 796.7)	-32.3
Rule-based center frequency			
Category A	(914, 197)	(1020.2, 3281.8)	-214.5
Category B	(816, 197)	(1020.2, 3281.8)	-129.1

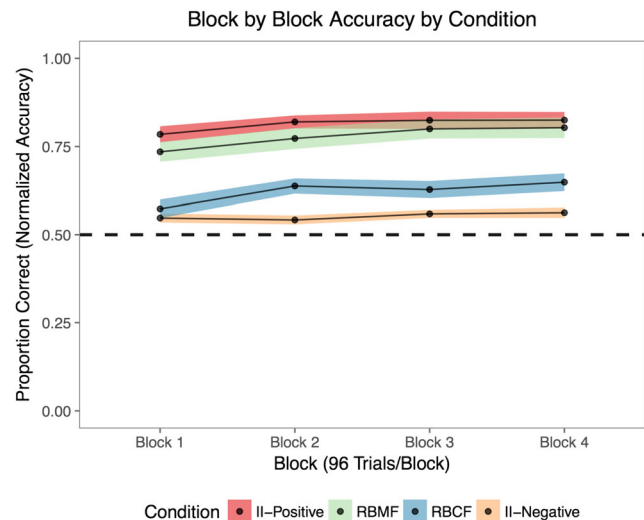*CF* center or carrier frequency, *MF* modulation frequency

for potential cross-condition differences (Normalized Accuracy = Raw Accuracy / Optimal Accuracy).[1] Below, we report only the normalized accuracies to give a more conservative measure of the differences among conditions.

Training accuracy We measured categorization response accuracy across the four blocks of training to determine the effect of stimulus distribution condition on training performance (Fig. 2). A mixed-model ANOVA with block as a within-subjects factor and condition as a between-subjects factor revealed a significant main effect of block ($F(3,222) = 8.53$, $p < .001$, $\eta_p^2 = .10$), a significant main effect of condition ($F(3,74) = 44.9$, $p < .001$, $\eta_p^2 = .65$), and no interaction ($F(9,222) = 0.96$, $p = .47$, $\eta_p^2 = .038$). We outline the results of the analyses for block and condition separately below.

> *Learning across blocks.* Bonferroni-corrected *post hoc* comparisons indicated that the majority of learning for all conditions occurred across the first two blocks. All blocks had significantly higher accuracy than Block 1 (Block 1 vs. Block 2 $p = .049$, $d = 0.30$, Block 1 vs. Block 3 $p = .004$, $d = 0.40$, Block 1 vs. Block 4 $p = .001$, $d = 0.45$), but Blocks 2, 3, and 4 did not differ from one another (all $ps > .05$, $ds < .20$). The majority of learning gains occurred within the first two blocks for all conditions. Learning occurred early and differences among conditions persisted throughout the experiment.
>
> *Differences among conditions.* Bonferroni-corrected *post hoc* comparisons revealed that participants learning the $II_{Positive}$ categories performed better than participants learning $II_{Negative}$ categories across training ($p < .001$, $d = 4.19$). Categorization accuracy was also higher for participants in the RB_{MF} condition relative to the RB_{CF} condition ($p < .001$, $d = 1.58$). Further, the $II_{Positive}$ and RB_{MF} categories did not differ in accuracy over the course of training ($p > .99$, $d = 0.37$) and neither did the RB_{CF} and $II_{Negative}$ ($p = .06$, $d = 1.06$). Participants in the $II_{Positive}$ condition had significantly better performance than RB_{CF} in all four training blocks ($p < .001$, $d = 2.25$). However, $II_{Negative}$ had significantly lower accuracy than RB_{MF} ($p < .001$, $d = 2.83$). In our data, training accuracy cannot be easily explained by classifying the category-learning challenge as a rule-based or information-integration category distribution. Instead, we found substantial differences between the two II conditions and differences between the two RB conditions, which we examine in greater detail below.

Fig. 2 Block-by-block average normalized accuracy, normalized according to ideal observer accuracy, for all conditions. Ribbon error bars reflect standard error of the mean. Dashed line represents chance accuracy (50%)

We found very striking differences in learning across the two II conditions, which had identical category distributions, but required integration across the dimensions in opposite directions. In the first block, average accuracy of both $II_{Positive}$ and $II_{Negative}$ conditions was above chance ($II_{Positive}$: 78.5%, $t(18) = 12.4$, $p < .001$, $d = 2.85$; $II_{Negative}$: 54.7%, $t(19) = 3.75$, $p = .001$, $d = 0.84$). However, recall that participants in the $II_{Positive}$ condition had significantly higher accuracy than participants in the $II_{Negative}$ condition throughout training ($p < .001$, $d = 4.19$). By the end of training in Block 4, average accuracy for each condition was still above chance ($II_{Positive}$: $t(18) = 14.0$, $p < .001$, $d = 3.22$; $II_{Negative}$: $t(19) = 4.23$, $p < .001$, $d = 0.95$), but participants in the $II_{Positive}$ condition reached 82.5% accuracy whereas participants in the $II_{Negative}$ condition achieved only 56.3% correct. Across training, we found much better performance for participants learning the II distribution that required an integration along the positive axis compared to participants learning the II distribution that required an integration along the negative axis.

We also found significant differences in performance across the two RB conditions. The primary difference between these two distributions is the dimension that distinguishes the categories. While both RB_{CF} and RB_{MF} performed above chance even in the first block (RB_{CF}: 57.3%, $t(19) = 2.79$, $p = .012$, $d = 0.62$; RB_{MF}: 73.5%, $t(18) = 8.36$, $p < .001$, $d = 1.92$), the RB_{MF} condition outperformed the RB_{CF} condition ($p < .001$, $d = 1.58$). This pattern remained throughout training, and by Block 4, participants in the RB_{MF} condition reached 80.4% accuracy and participants in the RB_{CF} condition reached only 64.9% accuracy. Block 4 accuracy remained significantly greater than chance for both the RB_{CF} ($t(19) = .590$, $p < .001$, $d = 1.32$) and RB_{MF} conditions ($t(18) = 10.3$, $p < .001$, $d = 2.36$). Whereas participants in both RB conditions

---

[1] Compared to the non-normalized data, the only differences were that with the normalized accuracy RB_{CF} was above chance in the first block ($t(19) = 2.788$, $p = .012$) and there were no statistically significant differences between $II_{Positive}$ and RB_{MF} in the training blocks (whereas with the non-normalized scores, $II_{Positive}$ and RB_{MF} were significantly different in the first and second training blocks).

were performing above chance throughout training, those learning categories that required a distinction based on MF outperformed those learning categories that required a distinction based on CF throughout training.
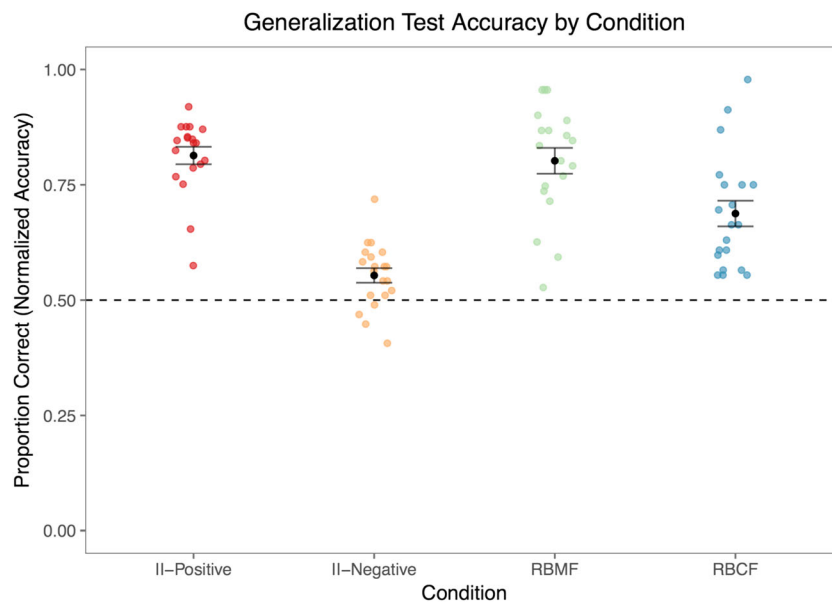
Generalization test accuracy After training, participants engaged in a generalization test that involved categorizing novel sound exemplars drawn from the distributions experienced in training without feedback (Fig. 3). Participants in all four conditions exhibited generalization performance greater than chance (50%) accuracy, indicating category learning ($II_{Positive}$ $t(18) = 16.5$, $p < .001$, $d = 3.80$; $II_{Negative}$ $t(19) = 3.38$, $p = .003$, $d = 0.76$; RB_{CF} $t(19) = 6.72$, $p < .001$, $d = 1.50$; RB_{MF} $t(18) = 10.8$, $p < .001$, $d = 2.47$). Generalization test accuracy varied across conditions ($F(3,74) = 27.20$, $p < .001$, $\eta_p^2 = .52$). Participants in the $II_{Positive}$ condition accurately categorized novel sound exemplars 81.4% on average; participants in the $II_{Negative}$ condition reached 55.4%, participants in the RB_{MF} condition reached 80.2%, and participants the RB_{CF} condition reached 68.8% correct. According to Bonferroni-corrected *post hoc* comparisons, the overall pattern of generalization mirrors the patterns of learning during training: $II_{Positive}$ generalization accuracy was greater than $II_{Negative}$ ($p < .001$, $d = 3.38$) and RB_{MF} was greater than RB_{CF} ($p = .006$, $d = 0.93$). $II_{Positive}$ and RB_{MF} had statistically equivalent performance ($p = .74$, $d = 0.11$), but $II_{Negative}$ was significantly worse than RB_{MF} ($p < .001$, $d = 2.48$). RB_{CF} generalization performance was significantly worse than $II_{Positive}$ ($p = .002$, $d = 1.19$). In general, generalization performance patterned with relative performance across conditions in training. The only difference in the overall pattern of results

compared to training is that in the generalization test, participants in the RB_{CF} condition performed significantly better than participants in the $II_{Negative}$ condition ($p = .001$, $d = 1.32$).

Computational modeling

Rationale Categorization accuracy across training and in the generalization test provides a relatively coarse measure of performance that does not reveal why differences between the II conditions and between the RB conditions persist. To obtain a better understanding of what participants learned over the course of this experiment, we applied and fit decision-bound models to each block of each participant's data (Ashby, 1992a; Ashby & Maddox, 1992, 1993; Maddox & Ashby, 1993). Decision-bound models are derived from General Recognition Theory (GRT, Ashby & Townsend, 1986), a multivariate application of signal detection theory (e.g., Green & Swets, 1966). These models have been applied extensively in the dual-systems literature with both auditory and visual categories (e.g., Ashby & Maddox, 2005, 2011; Chandrasekaran, Yi, et al., 2014; Maddox, Chandrasekaran, Smayda, & Yi, 2013; Scharinger, Henry, & Obleser, 2013). We provide a brief description of the models applied to the data; more specific details of these models, including the proposed neural instantiation of the models, can be found elsewhere (Ashby & Maddox, 1993; Ashby, Paul, & Maddox, 2011; Maddox & Ashby, 1993; Maddox & Chandrasekaran, 2014).

Model details Each model assumes participants create decision boundaries to separate the stimuli into two categories. Our model-based approach involves applying four classes

Fig. 3 Average generalization test accuracy, normalized based on optimal accuracy for each condition. The dashed line represents chance accuracy (50%). Error bars reflect standard error of the mean around the black dot, which represents the mean. Each individual point is an individual participant's average accuracy

of models, with multiple instantiations possible within a class. We fit a unidimensional model based on decision bounds across the CF dimension (UD_{CF}), a unidimensional model based on decision bounds across the MF dimension (UD_{MF}), an integration model (GLC) with decision bounds based on both CF and MF dimensions, and a random responder model (RR).

Unidimensional rule-based models. Two unidimensional models instantiate a unidimensional decisional bound that is optimal for either the RB_{CF} or RB_{MF} conditions. The unidimensional model has two free parameters – the decision boundary (vertical (90°) for UD_{CF} and horizontal (0°) for UD_{MF}) and the variance of noise (both perceptual and criterial). An example of a unidimensional rule based on CF might be: "If the tone's CF is greater than 866 Hz, it belongs to category A; if it less than 866 Hz, it belongs to category B." Optimal performance in the RB_{CF} requires a UD_{CF} decision bound whereas RB_{MF} requires a UD_{MF} decision bound.

Information-integration model. The general linear classifier (GLC) also assumes a linear decision boundary but, in contrast to the unidimensional rule-based models, it requires linear integration of the CF and MF dimensions and is therefore optimal for the $II_{Positive}$ and $II_{Negative}$ conditions. For the $II_{Positive}$ condition, the optimal decision boundary has a positive slope (45°), whereas for the $II_{Negative}$ condition, the optimal decision boundary has a negative slope (−45°). The specific weight a listener places on one dimension can vary, even when fit by the same GLC model. Thus, we also examine the angle of the decision boundaries in the CFxMF input space as an estimate of the perceptual weight of CF versus MF in categorization decisions. The model has three free parameters: the slope and intercept of the decision boundary and the variance of noise (perceptual and criterial).

Random responder model. The random responder model assumes that the participant guesses on each trial.

Model fitting For each of the four experimental conditions, we fit the models separately to each participant's data from each of the four training blocks and the generalization test. The model parameters were estimated using a maximum likelihood procedure (Ashby, 1992b; Wickens, 1982) and the goodness-of-fit statistic was Akaike's information criterion (AIC) = $2r - 2lnL$ where r is the number of free parameters and L is the likelihood of the model given the data (Akaike, 1974). The AIC allows comparison of model fits because it penalizes a model for extra free parameters such that the smaller the AIC, the closer the model is to the "true" model, regardless of the number of free parameters. To find the best-fit model, we computed AIC values for each model and chose

the model associated with the smallest AIC value. We separately replicated the model fit analyses using the Bayesian information criterion (BIC) as the model selection criterion, which gives steeper penalties for extra free parameters. The qualitative pattern of results was not different with the AIC and BIC model fits and so we focus report results based on AIC selection criterion.

Modeling results To better understand the pattern of learning across the different conditions, we examined the proportion of participants' best fit by each computational model (Fig. 4) and a more detailed measure of the boundaries participants drew between the categories in the generalization test (Figs. 5 and 6).

Proportion of participants using each strategy. Figure 4 shows the proportion of participants whose categorization decisions were best fit by the information-integration (GLC), UD_{CF}, UD_{MF}, and RR model, separately for each condition. Note that none of the participants in our study were best fit by the RR model in any block. We found that the strategy participants used in the first block was not independent from condition ($X^2(6, N = 78) = 17.4, p = .008$). We also examined proportions of participants using different strategies in the generalization test block and found the relation between strategy and condition was significant ($X^2(6, N = 78) = 17.4, p = .008$).

A majority of participants in the $II_{Positive}$ (78.9%) and RB_{MF} (78.9%) conditions were best fit by the integration strategy in the first block. The tendency to integrate across the dimensions for the $II_{Positive}$ and RB_{MF} conditions emerged early on and persisted throughout training. In Block 4, 100% of participants in the $II_{Positive}$ condition and 68.4% of participants in the RB_{MF} condition were best fit by the integration strategy. Integration was a successful strategy for these participants, as learning was most robust in these conditions. It is of note that this was a successful strategy for RB_{MF} participants, even though integration was a suboptimal strategy for this RB category type. Even though many participants in the RB_{MF} condition used an integration strategy rather than an optimal unidimensional strategy, their accuracy was still quite high.

In contrast, more participants in the RB_{CF} condition were best fit by one of the unidimensional (70%) strategies, compared to the integration (30%) strategies, in the first block (Bonferroni-corrected comparison, $p < .05$). However, many participants were fit by the suboptimal UD_{MF} strategy, indicating reliance on the MF dimension, which is poorly diagnostic of category membership in this condition. This helps to account for the poor categorization accuracy observed in training and generalization in the RB_{CF} condition.

Proportion of Participants fit by Each Strategy Across Blocks
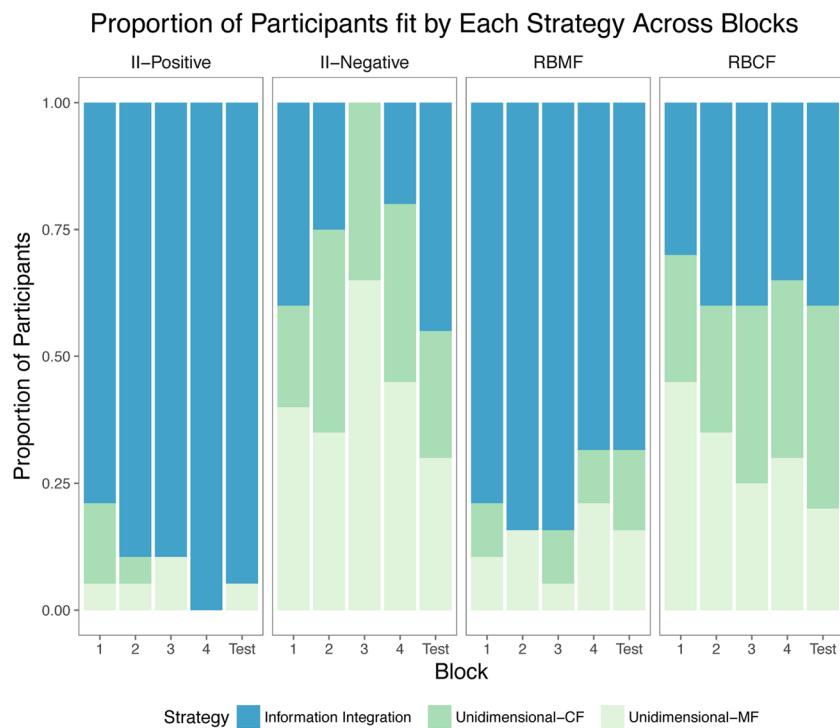
Fig. 4 Proportion of participants fit by each modeling strategy across all four training blocks and the generalization test. None of the participants were best fit by the Random Responder model, so it is not included in the graph

We also found that, in the first block across the $II_{Negative}$ condition, more participants' categorization responses were fit by a unidimensional strategy (60%) compared to an integration strategy (40%) (Bonferroni-corrected comparison, $p < .05$). Over the course of learning, participants in this $II_{Negative}$ condition were most often fit by a unidimensional
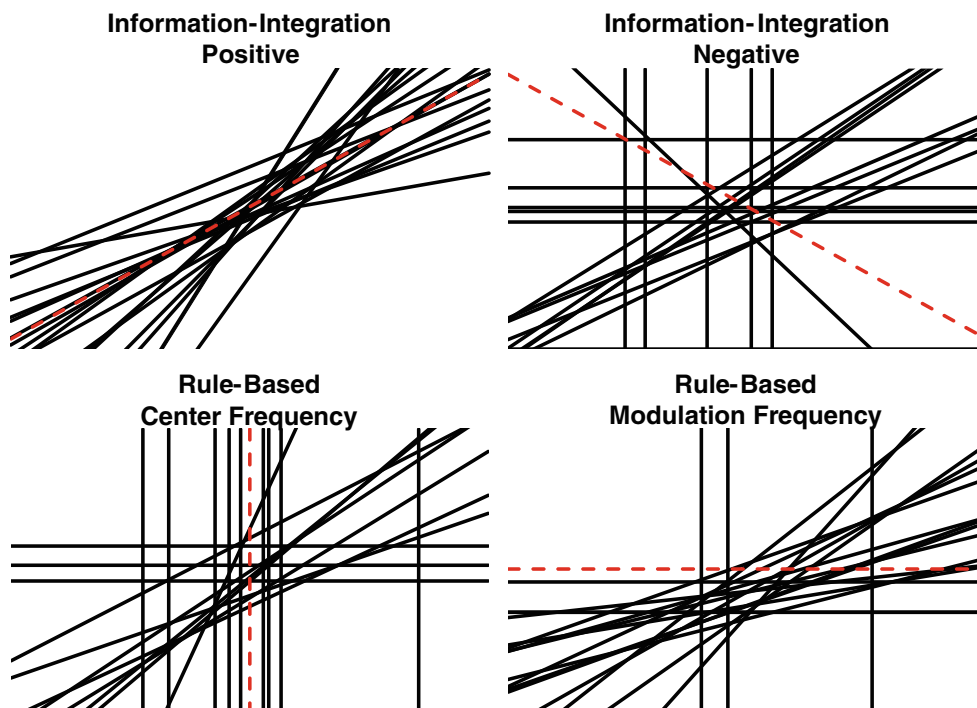
Fig. 5 Individual decision boundaries for each participant in the generalization test (after all training blocks). The optimal decision boundary for each category is shown as the red dotted line on each plot.

The x-axis represents the Center Frequency dimension, and the y-axis represents the Modulation Frequency dimension

Difference in Decision Bound Angles Relative to Optimal
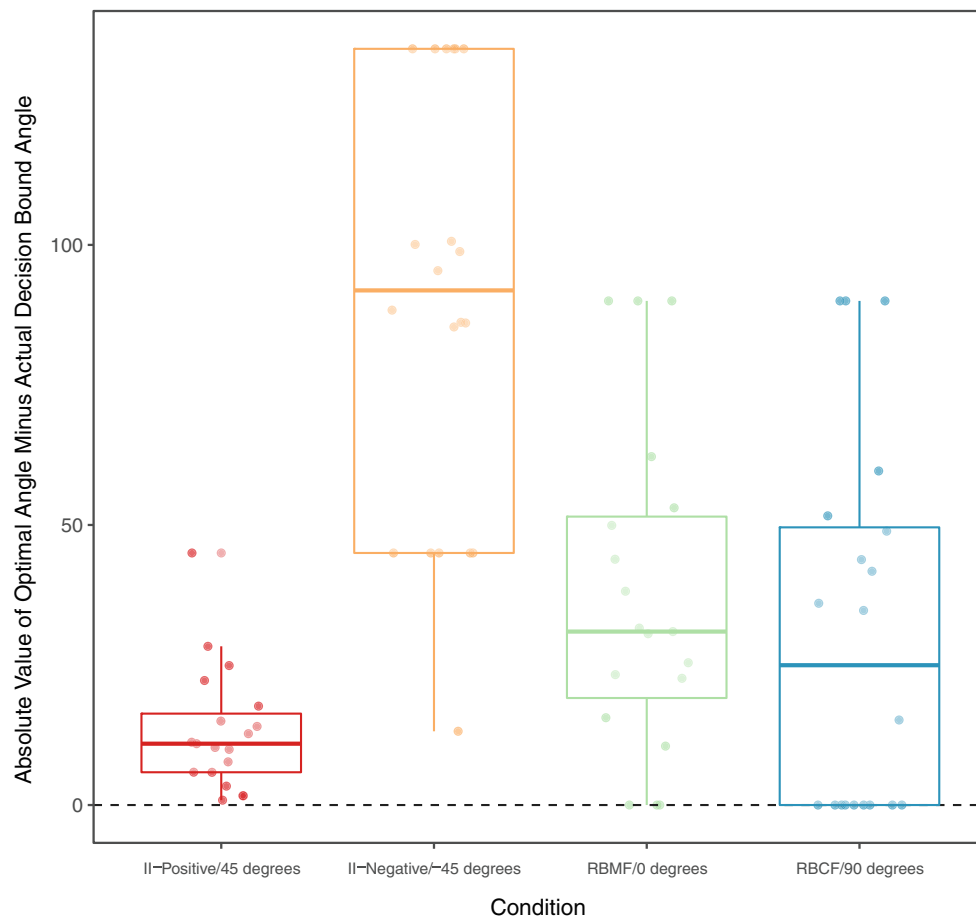
Fig. 6 Box plots of absolute value difference in participants' best fit decision-bound angles relative to the optimal decision boundary. The optimal decision-boundary angle is listed for each condition next to its name and is represented by the dashed line at 0. Each dot represents an individual participant value

strategy, rather than the optimal integration strategy. This is in sharp contrast to much greater adherence to the optimal integration strategy among participants in the $II_{Positive}$ condition, for which the only difference from the $II_{Negative}$ condition was the angle of the decision bound through MFxCF input space.

Decision boundaries in the generalization test. Examining the proportion of participants in each condition best fit by each of the models gives us a general sense of participants' category decision strategies in learning and generalization. However, only examining the proportion of participants using a given strategy does not present a full picture because many decision boundaries are possible within each class. Depending on how well a participant's decision boundary matches the optimal decision boundary, even within a class, there can be different effects on categorization accuracy. Figure 5 shows the individual best-fit decision boundaries for each participant

according to condition in the generalization test. The dashed line in each panel represents the optimal decision boundary.

By the generalization test phase, many participants in each condition were best fit by the integration model. This constituted the majority of participants for the RB_{MF} and $II_{Positive}$ conditions (68% and 95%, respectively), and also around 40% of the participants in the RB_{CF} and 45% of participants in the $II_{Negative}$ conditions. Visual inspection of the individual decision boundaries in Fig. 5 demonstrates that when participants were best fit by an integration strategy, it was along the positive axis, with the exception of a single participant in the $II_{Negative}$ condition. In the generalization test, there is a bias to integrate along these two dimensions in the positive direction, rather than selectively attending to either dimension (even when that is optimal for category learning) or integrating in the negative direction (even when that is optimal for category learning).

There are especially stark differences between the two II conditions. These two category types differ only in the direction of integration along the dimensions that is required by the category boundary. Nearly every participant in the $II_{Positive}$ condition used a nearly optimal decision boundary. In contrast, participants in the $II_{Negative}$ condition used a mixture of strategies in the generalization test. Among those using the integration strategy in the $II_{Negative}$ condition, all but one participant used a decision boundary with a positive slope between CF and MF, rather than the optimal negative slope. Thus, even participants best fit by the so-called optimal strategy (as in Fig. 4) were not optimally integrating across the two dimensions. This was especially true in the $II_{Negative}$ condition.

To better quantify the relative weight that participants placed on each dimension in the different conditions during the generalization test, we also computed the angle of the decision boundary for each participant. We compared the individual angle values to the optimal angle of the decision boundary for each of the conditions ($II_{Positive} = 45°$, $II_{Negative} = -45°$, $RB_{MF} = 0°$, $RB_{CF} = 90°$). Figure 6 shows the absolute values of these differences for each participant in each condition. The closer the participant's decision-boundary angle is to the optimal decision-boundary angle, the more they were optimally attending to the dimensions appropriate for the categories they were learning. This visualization helps to better understand how participants using the integration strategy differently weighted the two input dimensions in categorization decisions and provides more fine-grained information to quantify how close to optimal participants' strategies were in the generalization test.

We found that the vast majority of participants in the $II_{Positive}$ condition had a decision boundary with an angle very close to the optimal decision-boundary angle (median difference from optimal is 11.0°). In contrast, participants in the $II_{Negative}$ condition were very far from optimal (median difference from optimal is 91.9°). Thus, it may not be surprising that participants in the $II_{Negative}$ condition performed much worse than participants in the $II_{Positive}$ condition. Participants in the $II_{Negative}$ condition were less able to find the optimal integration strategy and even when they were best fit by the integration model, they were applying a decision boundary along the positive slope, opposite to what is optimal for the category distributions. Thus, integration alone is not enough for successful categorization and generalization – participants must integrate optimally.

In examining the differences between the two RB conditions, approximately equal numbers of participants in the RB_{MF} and RB_{CF} conditions had decision-boundary angles close to the optimal decision-boundary angle (Fig. 6). The median absolute difference from the optimal angle for RB_{MF} was 31.0° and for RB_{CF} was 25.0°. Just looking at the difference between these decision-boundary angles and what is optimal, it is not clear why RB_{MF} participants would outperform

RB_{CF} by such a large margin. However, it is clear from looking at the actual decision boundaries in Fig. 5 that the decision boundaries for participants using the optimal strategy (unidimensional CF for RB_{CF} and unidimensional MF for RB_{MF}) are not fully overlapping with the optimal decision boundary. Thus, it is not the angle that is suboptimal for the RB_{CF} participants, but the placement of the decision boundary on the x-axis. Participants in the RB_{CF} condition were unable to place the decision boundary at the optimal position along the CF dimension, even if they were reliant on CF for their decision. So even though many RB_{MF} participants integrated across dimensions and had a decision-boundary angle that was far from optimal, it resulted in better category learning than suboptimally placed unidimensional boundaries.

Across all participants, there was a significant correlation between the absolute difference in the decision-boundary angle relative to optimal and the generalization test accuracy ($r = -0.62$, $t(76) = -6.97$, $p < .0001$). However, in examining this correlation within each condition, the correlation was only significant for the $II_{Positive}$ ($r = -0.64$, $t(17) = -3.42$, $p = 0.003$) and RB_{MF} conditions ($r = -0.88$, $t(17) = -7.47$, $p < .0001$). There was no significant correlation for the $II_{Negative}$ ($r = -0.05$, $t(18) = -0.21$, $p = 0.83$) or RB_{CF} conditions ($r = -0.16$, $t(18) = -0.70$, $p = 0.49$). Across the entire group of participants, decision boundaries closer to optimal were associated with higher categorization accuracy.

Discussion

We examined learning outcomes and participant strategies for auditory categories defined by dimensions that are difficult to attend to selectively (Holt & Lotto, 2006). The results emphasize the importance of considering not just the physical acoustic dimensions that define a categorization challenge, but also the way that acoustic dimensions are represented *perceptually.* The nature of dimensions' perceptual representation greatly affects how categories are learned. Dual systems accounts, developed for visual categorization, have emphasized the importance of the sampling of category exemplars from a stimulus space on category-learning outcomes. Our results demonstrate that the *dimensions* defining a stimulus space also play a fundamental role. We cannot assume that a particular sampling from a physical stimulus space maps linearly to a *psychological* or *perceptual* space. Few models of category learning address how prior experience shapes representations and how existing representations constrain category learning (but see models of infant and second-language speech category learning: Best, McRoberts, & Sithole, 1988; Kuhl, 1991). Instead, it is common to assume that participants will be able to conquer each learning challenge placed in front of them, shifting attentional weights or decision boundaries flexibly based on the requirements of the task. In the current study,

even when sampling was equated across the II$_{Negative}$ and II$_{Positive}$ conditions, participants demonstrated strikingly different learning outcomes and across the two RB conditions; participants were not easily able to disengage from the irrelevant dimension.

Integration strategies persisted even when they were suboptimal

We found that many participants integrated across the CF and MF dimensions even early in training and even when integration was suboptimal. Prior studies using separable dimensions have reported that participants demonstrate a tendency to selectively attend to the dimensions (e.g., Ashby et al., 1999; Huang-Pollock, Maddox, & Karalunas, 2011; Smith, Beran, Crossley, Boomer, & Ashby, 2010). In the current study, integration across the input dimensions could be described as the "default" strategy for participants. For the acoustic dimensions of the present study, integration emerged early on and produced the best outcomes in terms of categorization accuracy even when integration was not the "optimal" strategy predicted by the categorization challenge defined by the sampling of exemplars across the input dimensions.

Strikingly, the bias toward integration was present even for the two RB conditions, for which unidimensional strategies were optimal based on the sampling of stimuli from the input space. The pattern of strategies revealed by computational modeling indicates that participants in the RB conditions did not easily disengage from the irrelevant dimension, even when training feedback did not align with this strategy. This pattern is not typically observed in the existing research on the dual-systems theory, for which dimensions defining RB categories have tended to be separable (Ashby & Maddox, 1990, 2005, 2011; Chandrasekaran, Koslov, et al., 2014; Goudbeek, Cutler, & Smits, 2008; Goudbeek, Swingley, & Smits, 2009; Maddox & Ashby, 2004).

Because many auditory dimensions may be integral and difficult to selectively attend to, it is important to examine category learning where it is difficult or impossible to engage with only the relevant dimension during category learning. Many models of category learning, including the COVIS model, posit that participants shift their attention weights or selectively attend to individual input dimensions during learning (Ashby et al., 1998; Nosofsky, 1986). In the current study, many participants in the two RB conditions were able to categorize the exemplars reasonably well even without selectively attending to the dimensions. The ability to selectively attend to the dimensions was not required for above-chance performance in this task.

There are few studies comparing two RB category-learning challenges that differ only on the dimension to which selective attention is required (for two exceptions with visual categories see Ell et al., 2012 and Maddox & Dodd, 2003). Instead,

experimenters typically choose one of the dimensions and use it as the single representative RB category-learning challenge. We investigated two possible RB categories to avoid the assumption that participants may be equally likely to learn RB categories based on either dimension.

This proved to be informative. Learners in the RB$_{MF}$ condition outperformed learners in the RB$_{CF}$ condition throughout the entire experiment. We did not expect that the RB$_{CF}$ and RB$_{MF}$ conditions would demonstrate such strikingly different performance. In Holt and Lotto (2006), participants placed more perceptual weight on the CF dimension when either CF or MF alone could distinguish the categories. However, even when participants placed more weight on CF, they continued to rely upon both dimensions, suggesting it may be difficult to selectively attend to these dimensions.

Several other studies have found differences in the reliance upon various input dimensions during category learning. Although these studies differ in their details, they collectively demonstrate that some input dimensions are more likely to be relied upon in category learning than others (Ell, Ashby, & Hutchinson, 2012; Goudbeek et al., 2008, 2009; Holt & Lotto, 2006; Maddox & Dodd, 2003; Scharinger et al., 2013). In the current experiment, participants did not demonstrate a bias for one dimension over the other. Instead, participants demonstrated a bias to *integrate* across the two dimensions.

In the current experiment, participants were better able to learn the RB$_{MF}$ categories than the RB$_{CF}$ categories. Participants in the RB$_{MF}$ condition used integration strategies often and performed well, even though these strategies are suboptimal. Participants in the RB$_{CF}$ conditions used the optimal selective attention more frequently but applied these strategies suboptimally. Thus, even though RB$_{MF}$ participants were technically using suboptimal decision strategies, the actual decision boundaries they were placing were closer to optimal than the optimal strategy decision boundaries that RB$_{CF}$ participants were placing. This finding is counterintuitive, but consistent in that it further demonstrates how difficult it is to optimally selectively attend to these acoustic dimensions.

A unique feature of the current study is that we examined learning of distinct category structures within the same acoustic space defined by dimensions that are difficult to attend to selectively. In a study investigating visual category learning with the integral dimensions of saturation and brightness, researchers have demonstrated integration across the integral dimensions, but generally RB categorization performance was better than II performance (Ell et al., 2012). In a study of auditory category learning with perceptually integral dimensions that are closely related to pitch and timbre, researchers found that participants strongly relied upon unidimensional strategies compared to integration strategies (Scharinger et al., 2013). In contrast, in the current experiment, we found that participants were more often integrating

across the dimensions and there was no clear benefit for II categories over RB categories.

The present results highlight that it is important to consider the role of perceptual dimensions in category learning. Whereas separable, relatively easy to verbalize acoustic dimensions, such as pitch and duration, may behave in a manner aligned with the COVIS model (Chandrasekaran, Koslov, et al., 2014; Goudbeek et al., 2008, 2009), this may not be the case for integral, interacting, or difficult to verbalize acoustic dimensions. Nonetheless, it is important to acknowledge that the current experiment investigated two specific acoustic dimensions; the observed pattern of results may not be true for all acoustic dimensions. Especially in light of the conflict between the current results and those of Scharinger et al. (2013), other interacting and integral acoustic dimensions should be examined. It is important to investigate learning categories defined by complex and interacting acoustic dimensions because many acoustic dimensions, including those defining speech categories, interact (Francis et al., 2000; Grau & Kemler Nelson, 1988; Hillenbrand et al., 1995). The present results highlight the need for caution in assuming the psychological relationship among perceptual dimensions involved in category learning.

There was a bias to integrate across the dimensions in a way that reflected a *positive* correlation between the dimensions

Not only did the participants demonstrate a propensity to integrate across the dimensions in the current study, they did so in a manner that reflected a positive correlation between CF and MF. This propensity to integrate along the positive correlation axis had a particularly potent impact on the two II conditions. Whereas category learning was robust in the $\text{II}_{\text{Positive}}$ condition, it hovered near chance in the $\text{II}_{\text{Negative}}$ condition. This is notable inasmuch as the statistical sampling of the acoustic input space was identical, except that the category boundary was rotated 90°.

The nature of this interaction may stem from the physical dimensions and participants' prior experience. Although we used nonspeech stimuli to attempt to better control the acoustic environment and minimize participants' prior experience with category exemplars, it still could be the case that existing representations for these acoustic dimensions, and their relationship, influenced category learning. Listeners are sensitive to statistical correlations among acoustic dimensions (Holt & Lotto, 2006; Liu & Holt, 2015; McMurray, Aslin, & Toscano, 2009; Stilp, Rogers, & Kluender, 2010; Wade & Holt, 2005). It is possible that the positive integration bias we found may be due to a natural correlation between these two physical dimensions. For instance, it might be the case that due to the physics of sound, when a sound has a higher CF, more modulations can be added. This interrelation between the

dimensions may lead to a natural positive correlation between CF and MF and may have contributed to the bias to integrate along the positive axis that we observed here. The current study was not designed to address this particular relationship between the two dimensions, but our results do point to the need to clarify potentially pre-existing relationships between perceptual dimensions and how they might influence learning.

Additionally, while our nomenclature of positive and negative integration refers to the direction of the decision boundary, the *perceptual* distinction that is made between the categories is orthogonal to the decision boundary. For the $\text{II}_{\text{Positive}}$ category, the boundary has a positive slope, but the two categories require distinctions between categories across this boundary, which would be along the negative axis. This is an important distinction to make to understand how the perceptual information is used and processed by learners. This also further highlights the importance of considering the underlying perceptual representation of the categories being learned in addition to the decision boundary in the perceptual space.

The propensity to integrate along the positive axis observed in the current study specifically benefitted learning of $\text{II}_{\text{Positive}}$ categories and impaired learning of $\text{II}_{\text{Negative}}$ categories. Very few studies have specifically compared learning of two different information-integration conditions (but see Ell et al., 2012; Scharinger et al., 2013). Scharinger et al. (2013) found that participants were better able to learn II categories with a negative correlation than a positive correlation. They argued that performance in that condition was better because the negative correlation better matched the natural correlation found in speech. The dimensions in the current study are not directly comparable to speech dimensions. However, our results support Scharinger et al.'s (2013) conclusion that prior experience with similar or identical dimensions influences category-learning behavior. It could be the case that either prior experience or general familiarity with these or similar acoustic dimensions could be driving perceptual processing and learning.

In contrast, Ell et al. (2012) found that for the integral visual dimensions of saturation and brightness, the positive and negative axis II categories were learned equally well. They did not observe a difference in their two II conditions because participants demonstrated a bias toward strategies that required selective attention to brightness, while ignoring saturation. Because there was no strong tendency to integrate across dimensions, Ell et al. (2012) did not find differences in their two II conditions.

Many acoustic dimensions demonstrate similar interacting relationships, such as pitch and loudness and pitch and timbre (Melara & Marks, 1990; Neuhoff, 2004). Thus, investigating how categories defined by complex and interacting dimensions are learned is an important and understudied area of category learning. Instead, typical experiments – for the sake

of simplicity – define their categories based on simple, separable dimensions. Use of these simple, separable dimensions may lead to better experimental control, but it comes at a cost to generalizability to real-world dimensions, which are often more complicated.

Additionally, it is possible that individual differences in musical expertise, language experience, or hearing loss may affect how participants interact with acoustic information during learning. Because none of these factors were directly tied to our main question of interest, we did not collect these measures from participants and are thus unable to evaluate whether and which individual differences had any effect on performance. Future studies may be directed at understanding the factors of individual participants that may lead to differences in perceptual processing or category learning in general.

Implications for models of category learning

Our results demonstrate that the nature of perceptual dimensions, in terms of their perceptual interaction or non-separability, impacts category learning. The influence of dimensions was apparent in the course of learning, the strategies participants applied in learning, and in generalization of learning.

The current study demonstrates that for the COVIS dual-systems approach to be sufficiently expanded into the auditory domain, the role and influence of integral or interacting dimensions must be taken into account. It will be important to consider not only sampling across input dimensions, but also the mapping of these dimensions to *perceptual* dimensions. Acoustic dimensions, including those important for speech, are often highly integral and often are not easily verbalizable (Garner, 1974; Grau & Kemler Nelson, 1988; Melara & Marks, 1990). Speech categories are highly multidimensional, with dimensions that are difficult to describe verbally and are often perceptually integral (Francis et al., 2000; Grau & Kemler Nelson, 1988; Hillenbrand et al., 1995). Researchers have demonstrated that many acoustic dimensions that contribute to speech perception are perceptually integral (Kingston, Diehl, Kirk, & Castleman, 2008; Macmillan, Kingston, Thorburn, Walsh Dickey, & Bartels, 1999). This perceptual integrality of speech dimensions may be driven by physical constraints on articulatory mechanisms that renders acoustic input dimensions to be interdependent (Carré, 2009) in the manner of the "natural co-variation" that we posit to explain the relationship between CF and MF in the current study. Alternatively, this integrality may also be a consequence of some psychoacoustic similarity between the dimensions (Diehl, 2008; Kingston et al., 2008; Macmillan et al., 1999). The issue of how and why acoustic dimensions co-vary has been engaged with in speech perception and has produced strong debates about whether this co-variation arises from environmental co-variation or from similar underlying auditory properties (Diehl & Kluender, 1989; Fowler, 1989). While the dimensions in the current study are not speech dimensions, the theoretical issue is quite similar and general theories of auditory categorization might benefit from engaging with the evidence from speech.

Although we set out to test predictions of the dual-systems perspective with auditory dimensions that are difficult to selectively attend to, the most meaningful patterns in the data were not between II and RB categories. Instead, stark differences *within* each category type ($II_{Positive}$ vs. $II_{Negative}$, RB_{MF} vs. RB_{CF}) emerged that are not easily explained by existing dual-systems frameworks. Our results indicate that understanding category learning requires understanding how the dimensions that define the space in which categories are situated are represented perceptually. The present results caution that the categories laid out on the page as orthogonal dimensions (acoustic categories, in this case), may not align with traditional conceptualizations of RB and II categories. Instead, the defining factor in determining whether categories are better described as II or RB relies on the *perceptual* not the physical space. A potential consequence of this is that category exemplar distributions that may appear to be RB (or II) learning challenges in the physical space may not truly be RB (or II) problems after taking into account representations in the perceptual space.

It is important to note that category-learning theories apart from COVIS have focused attention on dimensions or how participants weight different cues during learning (Francis & Nusbaum, 2002; Goldstone, 1993, 1994; Nosofsky, 1986). However, these models still assume that learners are generally flexible and adaptable and can learn to shift attention within the input space to respond to the demands of the current category-learning challenge. Additionally, many of these models do not propose a neurobiologically plausible account of how attention to dimensions interacts with category learning. Future research will need to address how to integrate theories to advance understanding of the interplay of perceptual encoding, attention, and learning involved in acquiring new categories.

Conclusion

The dimensions that define categories affect the ability to learn those categories. Here, participants demonstrated a bias to integrate across acoustic dimensions in a way that reflected a positive relationship between dimensions. This led to high accuracy for II categories requiring positive integration but was detrimental for learning statistically equivalent sampling of II category exemplars that required negative integration across dimensions. Participants often integrated along the dimensions even when this strategy was suboptimal for learning, in the case of RB categories. These suboptimal integration strategies were not detrimental for learning in the RB_{MF}

condition. However, learning in the RB$_{CF}$ condition was worse than in the RB$_{MF}$ condition. Thus, the dimensions used to define categories and the relationship between those dimensions greatly affected participants' category-learning performance, the strategies they used during learning, and their ability to generalize category learning to novel exemplars. The interaction of dimensions in experience and perception impacts category learning in a way that is currently unexplained by the existing COVIS dual-systems framework and other models of category learning. Thus, we caution for the need to consider that the input space is not necessarily homologous with the perceptual space as this has important effects on category learning.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

References

Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, *19*, 716–723.

Ashby, F. G. (1992a). Multidimensional models of categorization. In *Multidimensional models of perception and cognition* (pp. 449–483). Retrieved from http://psycnet.apa.org/psycinfo/1992-98026-016

Ashby, F. G. (1992b). Multivariate probability distributions. In *Multidimensional models of perception and cognition* (pp. 1–34).

Ashby, F. G., Alfonso-Reese, L. A., Turken, A. U., & Waldron, E. M. (1998). A neuropsychological theory of multiple systems in category learning. *Psychological Review*, *105*(3), 442–481. https://doi.org/10.1037/0033-295X.105.3.442

Ashby, F. G., & Maddox, W. T. (1990). Integrating information from separable psychological dimensions. *Journal of Experimental Psychology: Human Perception and Performance*, *16*(3), 598–612. https://doi.org/10.1037/0096-1523.16.3.598

Ashby, F. G., & Maddox, W. T. (1992). Complex decision rules in categorization: Contrasting novice and experienced performance. *Journal of Experimental Psychology: Human Perception and Performance*, *18*(1), 50–71. https://doi.org/10.1037/0096-1523.18.1.50

Ashby, F. G., & Maddox, W. T. (1993). Relations between prototype, exemplar, and decision bound models of categorization. *Journal of Mathematical Psychology*, *37*, 372–400. https://doi.org/10.1006/jmps.1993.1023

Ashby, F. G., & Maddox, W. T. (2005). Human category learning. *Annual Review of Psychology*, *56*, 149–178. https://doi.org/10.1146/annurev.psych.56.091103.070217

Ashby, F. G., & Maddox, W. T. (2011). Human category learning 2.0. *Annals of the New York Academy of Sciences*, *1224*, 147–61. https://doi.org/10.1111/j.1749-6632.2010.05874.x

Ashby, F. G., Paul, E. J., & Maddox, W. T. (2011). COVIS. In E. M. Pothos & A. J. Wills (Eds.), *Formal approaches in categorization.*

Ashby, F. G., Queller, S., & Berretty, P. M. (1999). On the dominance of unidimensional rules in unsupervised categorization. *Perception & Psychophysics*, *61*(6), 1178–1199. https://doi.org/10.3758/BF03207622

Ashby, F. G., & Townsend, J. T. (1986). Varieties of perceptual independence. *Psychological Review*, *93*(2), 154–179.

Best, C. T., McRoberts, G. W., & Sithole, N. M. (1988). Examination of perceptual reorganization for nonnative speech contrasts: Zulu click discrimination by english-speaking adults and infants. *Journal of Experimental Psychology: Human Perception and Performance*, *14*(3), 345–360. https://doi.org/10.1037/0096-1523.14.3.345

Carré, R. (2009). Dynamic properties of an acoustic tube: Prediction of vowel systems. *Speech Communication*, *51*(1), 26–41. https://doi.org/10.1016/j.specom.2008.05.015

Chandrasekaran, B., Koslov, S. R., & Maddox, W. T. (2014). Toward a dual-learning systems model of speech category learning. *Frontiers in Psychology*, *5*(July), 1–17. https://doi.org/10.3389/fpsyg.2014.00825

Chandrasekaran, B., Yi, H.-G., & Maddox, W. T. (2014). Dual-learning systems during speech category learning. *Psychonomic Bulletin & Review*, *21*, 488–95. https://doi.org/10.3758/s13423-013-0501-5

Diehl, R. L. (2008). Acoustic and auditory phonetics: The adaptive design of speech sound systems. *Philosophical Transactions of the Royal Society, B: Biological Sciences*, *363*(1493), 965–978. https://doi.org/10.1098/rstb.2007.2153

Diehl, R. L., & Kluender, K. R. (1989). On the objects of speech perception. *Ecological Psychology*, *1*(2), 121–144. https://doi.org/10.1207/s15326969eco0102_2

Ell, S. W., Ashby, F. G., & Hutchinson, S. (2012). Unsupervised category learning with integral-dimension stimuli. *The Quarterly Journal of Experimental Psychology*, *65*(8), 1537–1562.

Fowler, C. A. (1989). Real objects of speech perception: A commentary on Diehl and Kluender. *Ecological Psychology*, *1*(2), 145–160. https://doi.org/10.1207/s15326969eco0102

Francis, A. L., Baldwin, K., & Nusbaum, H. C. (2000). Effects of training on attention to acoustic cues. *Perception & Psychophysics*, *62*(8), 1668–1680. https://doi.org/10.3758/BF03212164

Francis, A. L., & Nusbaum, H. C. (2002). Selective attention and the acquisition of new phonetic categories. *Journal of Experimental Psychology: Human Perception and Performance*, *28*(2), 349–366. https://doi.org/10.1037/0096-1523.28.2.349

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

Garner, W. R. (1976). Interaction of stimulus dimensions in concept and choice processes. *Cognitive Psychology*, *8*(1), 98–123. https://doi.org/10.1016/0010-0285(76)90006-2

Garner, W. R. (1978). Selective attention to attributes and to stimuli. *Journal of Experimental Psychology. General*, *107*(3), 287–308. https://doi.org/10.1037/0096-3445.107.3.287

Goldstone, R. L. (1993). Feature distribution and biased estimation of visual displays. *Journal of Experimental Psychology: Human Perception and Performance*, *19*(3), 564–579. https://doi.org/10.1037/0096-1523.19.3.564

Goldstone, R. L. (1994). Influences of categorization on perceptual discrimination. *Journal of Experimental Psychology: General*, *123*(2), 178–200.

Goudbeek, M., Cutler, A., & Smits, R. (2008). Supervised and unsupervised learning of multidimensionally varying non-native speech categories. *Speech Communication*, *50*(2), 109–125. https://doi.org/10.1016/j.specom.2007.07.003

Goudbeek, M., Swingley, D., & Smits, R. (2009). Supervised and unsupervised learning of multidimensional acoustic categories. *Journal of Experimental Psychology: Human Perception and Performance*, *35*(6), 1913–1933. https://doi.org/10.1037/a0015781

Grau, J. W., & Kemler Nelson, D. G. (1988). The distinction between integral and separable dimensions: Evidence for the integrality of pitch and loudness. *Journal of Experimental Psychology: General*, *117*(4), 347–370.

Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics.* New York: Wiley.

Hillenbrand, J., Getty, L. A., Clark, M. J., & Wheeler, K. (1995). Acoustic characteristics of American English vowels. *The Journal of the Acoustical Society of America*, *97*(5 Pt 1), 3099–3111. https://doi.org/10.1121/1.411872

Holt, L. L., & Lotto, A. J. (2006). Cue weighting in auditory categorization: Implications for first and second language acquisition. *The Journal of the Acoustical Society of America*, *119*(5), 3059. https://doi.org/10.1121/1.2188377

Huang-Pollock, C. L., Maddox, W. T., & Karalunas, S. L. (2011). Development of implicit and explicit category learning. *Journal of Experimental Child Psychology*, *109*(3), 321–35. https://doi.org/10.1016/j.jecp.2011.02.002

Kemler, D. G., & Smith, L. B. (1979). Accessing similarity and dimensional relations: Effects of integrality and separability on the discovery of complex concepts. *Journal of Experimental Psychology. General*, *108*(2), 133–150. https://doi.org/10.1037/0096-3445.108.2.133

Kingston, J., Diehl, R. L., Kirk, C. J., & Castleman, W. A. (2008). On the internal perceptual structure of distinctive features: The [voice] contrast. *Journal of Phonetics*, *36*(1), 28–54. https://doi.org/10.1016/j.wocn.2007.02.001

Kuhl, P. K. (1991). Human adults and human infants show a "perceptual magnet effect" for the prototypes of speech categories, monkeys do not. *Perception & Psychophysics*, *50*(2), 93–107.

Liu, R., & Holt, L. L. (2015). Dimension-based statistical learning of vowels. *Journal of Experimental Psychology: Human Perception and Performance*, *41*(6), 1783–1798. https://doi.org/10.1037/xhp0000092

Macmillan, N. A., Kingston, J., Thorburn, R., Walsh Dickey, L., & Bartels, C. (1999). Integrality of nasalization and F1. II. Basic sensitivity and phonetic labeling measure distinct sensory and decision-rule interactions. *The Journal of the Acoustical Society of America*, *106*(5), 2913–32. https://doi.org/10.1121/1.428113

Maddox, W. T., & Ashby, F. G. (1993). Comparing decision bound and exemplar models of categorization. *Perception & Psychophysics*, *53*(1), 49–70. https://doi.org/10.3758/BF03211715

Maddox, W. T., & Ashby, F. G. (2004). Dissociating explicit and procedural-learning based systems of perceptual category learning. *Behavioural Processes*, *66*(3), 309–32. https://doi.org/10.1016/j.beproc.2004.03.011

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*(4), 650–662. https://doi.org/10.1037/0278-7393.29.4.650

Maddox, W. T., & Chandrasekaran, B. (2014). Tests of a dual-systems model of speech category learning. *Bilingualism: Language and Cognition*, *17*(4), 709–728. https://doi.org/10.1016/j.biotechadv.2011.08.021.Secreted

Maddox, W. T., Chandrasekaran, B., Smayda, K., & Yi, H.-G. (2013). Dual systems of speech category learning across the lifespan. *Psychology and Aging*, *28*(4), 1042–56. https://doi.org/10.1037/a0034969

Maddox, W. T., & Dodd, J. L. (2003). Separating perceptual and decisional attention processes in the identification and categorization of integral-dimension stimuli. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *29*(3), 467–480. https://doi.org/10.1037/0278-7393.29.3.467

Maddox, W. T., Filoteo, J. V., Lauritzen, J. S., Connally, E., & Hejl, K. D. (2005). Discontinuous categories affect information-integration but not rule-based category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *31*(4), 654–69. :https://doi.org/10.1037/0278-7393.31.4.654

Maddox, W. T., Molis, M. R., & Diehl, R. L. (2002). Generalizing a neuropsychological model of visual categorization to auditory categorization of vowels. *Perception & Psychophysics*, *64*(4), 584–597. https://doi.org/10.3758/BF03194728

McKinley, S. C., & Nosofsky, R. M. (1996). Selective attention and the formation of linear decision boundaries. *Journal of Experimental Psychology. Human Perception and Performance*, *22*(2), 294–317. https://doi.org/10.1037/0096-1523.24.1.339

McMurray, B., Aslin, R. N., & Toscano, J. C. (2009). Statistical learning of phonetic categories: Insights from a computational approach. *Developmental Science*, *12*(3), 369–378. https://doi.org/10.1111/j.1467-7687.2009.00822.x

Melara, R. D., & Marks, L. E. (1990). Interaction among auditory dimensions: Timbre, pitch, and loudness. *Perception & Psychophysics*, *48*(2), 169–178. https://doi.org/10.3758/BF03207084

Morrison, R. G., Reber, P. J., Bharani, K. L., & Paller, K. A. (2015). Dissociation of category-learning systems via brain potentials. *Frontiers in Human Neuroscience*, *9*(July), 1–11. https://doi.org/10.3389/fnhum.2015.00389

Neuhoff, J. G. (2004). *Ecological Psychoacoustics.* (J. G. Neuhoff, Ed.). New York: Academic Press.

Newell, B. R., Dunn, J. C., & Kalish, M. (2011). Systems of category learning. Fact or Fantasy? In *Psychology of learning and motivation - Advances in research and theory* (1st ed., Vol. 54, pp. 167–215). Elsevier Inc. 10.1016/B978-0-12-385527-5.00006-1

Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, *115*(1), 39–57.

Roark, C. L., & Holt, L. L. (2018). Task and distribution sampling affect auditory category learning. *Attention, Perception & Psychophysics*, *80*(7), 1804–1822. https://doi.org/10.3758/s13414-018-1552-5

Scharinger, M., Henry, M. J., & Obleser, J. (2013). Prior experience with negative spectral correlations promotes information integration during auditory category learning. *Memory & Cognition*, *41*, 752–68. https://doi.org/10.3758/s13421-013-0294-9

Smith, E. E., & Grossman, M. (2008). Multiple systems of category learning. *Neuroscience and Biobehavioral Reviews*, *32*(2), 249–64. https://doi.org/10.1016/j.neubiorev.2007.07.009

Smith, J. D., Beran, M. J., Crossley, M. J., Boomer, J., & Ashby, F. G. (2010). Implicit and explicit category learning by macaques (Macaca mulatta) and humans (Homo sapiens). *Journal of Experimental Psychology: Animal Behavior Processes*, *36*(1), 54–65. https://doi.org/10.1037/a0015892

Stilp, C. E., Rogers, T. T., & Kluender, K. R. (2010). Rapid efficient coding of correlated complex acoustic properties. *Proceedings of the National Academy of Sciences of the United States of America*, *107*(50), 21914–9. https://doi.org/10.1073/pnas.1009020107

Wade, T., & Holt, L. L. (2005). Incidental categorization of spectrally complex non-invariant auditory stimuli in a computer game task. *The Journal of the Acoustical Society of America*, *118*, 2618–2633. https://doi.org/10.1121/1.2011156

Wickens, T. D. (1982). *Models for behavior: Stochastic processes in psychology.* San Francisco, CA: W. H. Freeman.