

Feature weighting in “chick-a-dee” call notes of *Poecile atricapillus*

Carly M. Nickerson, Laurie L. Bloomfield, and Michael R. W. Dawson
Department of Psychology, University of Alberta, Edmonton, Alberta T6G 2E9, Canada

Isabelle Charrier
Bioacoustics Team, NAMC-CNRS UMR 8620, Université Paris Sud, Bat.446, F-91405, Orsay, France

Christopher B. Sturdy^{a)}
Department of Psychology, Centre for Neuroscience, University of Alberta, Edmonton, Alberta T6G 2E9, Canada

(Received 12 September 2006; revised 16 July 2007; accepted 17 July 2007)

Artificial neural networks were trained to discriminate between different note types from the black-capped chickadee (*Poecile atricapillus*) “chick-a-dee” call. Each individual note was represented as a vector of summary features taken from note spectrograms and networks were trained to respond to exemplar notes of one type and to fail to respond to exemplar notes of another type. Following initial network training, the network was presented novel notes in which individual acoustic features had been modified. The strength of the response of the network to each novel and shifted note was recorded. When network responses were plotted as a function of the degree of acoustic feature modification and training context, it became clear that modifications of some acoustic features had significant effects on network responses, while others did not. Moreover, the training context of the network also played a role in the responses of networks to manipulated test notes. The implications of using artificial neural networks to generate testable hypotheses for animal research and the role of context are discussed. © 2007 Acoustical Society of America.
[DOI: 10.1121/1.2770540]

PACS number(s): 43.80.Ka, 43.80.Lb, 43.80.Jz [JAS]

Pages: 2451–2458

I. INTRODUCTION

A perceptron is a simple artificial neural network (ANN) that consists of input processing units that encode stimuli, as well as output processing units that represent responses to these input patterns. These two sets of processing units are linked by a set of modifiable, weighted, connections. The weights are initially random, but with the use of a learning rule they achieve values that mediate a desired stimulus-response mapping (Rosenblatt, 1962; for a more general introduction, see Dawson, 2004).

Perceptrons are very simple ANNs because they do not have intermediate processing units (called hidden units) that intercept and modify input unit signals before they reach the output units. As a result, a perceptron’s response is based solely upon the sum of the weighted signals from each input unit. In the absence of hidden units, there are many stimulus-response mappings that cannot be computed by a perceptron (Minsky and Papert, 1988). Because of this, many researchers prefer to use more powerful networks that include hidden units, such as the multilayered perceptron (e.g., Rumelhart *et al.*, 1986).

However, even though perceptrons cannot learn to compute every stimulus-response mapping, the subset of mappings that they can represent is informative and interesting. In particular, perceptrons have been shown to be important

models in the domain of discrimination learning (Dawson, 2005; Dawson and Spetch, 2005; Nickerson *et al.*, 2006; Yaremchuk *et al.*, 2005). The purpose of the current paper is to explore the utility of perceptrons in the context of a particular discrimination learning paradigm that explores the component notes of the “chick-a-dee” call of the black-capped chickadee (*Poecile atricapillus*).

A. Note discrimination by black-capped chickadees

Black-capped chickadees produce a chick-a-dee call that consists of four note types categorized as A, B, C, and D by bioacousticians (Ficken *et al.*, 1978; see Fig. 1 for exemplar note types). Each note type can be repeated or omitted, but the note types always appear in a fixed order within the call. For instance, A notes occur before B notes, B notes before C notes, and C notes before D notes. The fixed sequence in which these note types occur in the black-capped chickadee call is analogous to the syntax of human language (Hailman, 1985; Hailman and Ficken, 1986). That is, the call’s generative syntax allows for countless note combinations, which parallels the human language system in the ability to generate an infinite variety of sentences from a finite vocabulary.

Pursuing this analogy, different combinations of notes in the chick-a-dee call might convey different messages, with each combination representing a different “sentence.” In fact, Ficken *et al.* (1994) noted that the chickadee call syntax aids in species identity, mate choice, and territory defense, and the call composition differs according to context in Mexican

^{a)}Electronic mail: csturdy@ualberta.ca




FIG. 1. Sound spectrogram of a “chick-a-dee” call (note types AABCD-DDD) with exemplars of A, B, C, and D note types.

chickadees (*Poecile sclateri*). As well, Charrier and Sturdy (2005) found that the black-capped chickadee fails to respond to atypical note ordering, and concluded that a fixed order is required to elicit behavior. Finally, Templeton *et al.* (2005) showed that chick-a-dee calls and their notes vary with the perceived threat level of a predator. Clearly, accurate perception of chick-a-dee calls and note type composition are important to the survival of chickadees in the field.

Researchers have also studied the functions of the constituent notes in the chick-a-dee call, independently of note ordering within calls. Playback studies have indicated that the black-capped chickadee D note mediates information on species and flock identification (Mammen and Nowicki, 1981; Nowicki, 1983), while the C note reveals information about food location and availability (Freeberg and Lucas, 2002). In addition, individual notes within the call may offer information about the caller such as identity and gender (Charrier *et al.*, 2004). Statistical analyses have also shown that any constituent note (A, B, C, or D) carries enough information to determine whether the note was uttered by a black-capped chickadee or by another species, the mountain chickadee (*Poecile gambeli*; Dawson *et al.*, 2006).

Considering that both the syntax and the component notes of the chick-a-dee call appear to have functional roles, chickadees must be able to categorize these notes into note types (e.g., to process syntax). Indeed, black-capped chickadees can differentiate between the bioacoustically defined A-, B-, C-, and D-note types (Sturdy *et al.*, 2000). For instance, Sturdy *et al.* (2000) found that between-category learning was faster than within-category learning of these note types. This result has led to subsequent research that studies the specific features and mechanisms used by black-capped chickadees to perceive and subsequently categorize individual notes.

For instance, Charrier *et al.* (2005) more closely examined how black-capped chickadees categorized the non-D note types in an operant discrimination experiment using a GO/NO GO paradigm. Of particular interest was a second phase that was conducted after the birds had successfully completed the acquisition phase of the discrimination. Charrier *et al.* (2005) investigated note classification mechanisms by presenting trained birds with notes that had been acoustically manipulated. In particular, they took exemplar notes (A, B, and C) and shifted their spectral features up or down

by shifting the entire spectrogram of a note up or down a specific interval from its original starting frequency. Charrier *et al.* (2005) found that black-capped chickadees could accurately categorize note types based on frequency characteristics alone (e.g., start frequency and ascending frequency modulation), and that adjacent note types were misclassified when the notes were modified by these spectral features (e.g., A notes shifted down were confused with B notes, while B notes shifted up were confused with A notes). That is, shifting the entire spectrogram of a note produced a systematic pattern of misclassifications by the birds.

B. Note discrimination by perceptrons

While it was clear that overall note frequency controlled note-type classification, what remained to be determined was why chickadees miscategorize shifted notes in the particular fashion observed by Charrier *et al.* (2005). Nickerson *et al.* (2006) began to answer this question by creating a model of note discrimination. They trained perceptrons in a simulation of the Charrier *et al.* (2005) experimental paradigm. These perceptrons used nine input units to represent each note as a set of summary features taken from sound spectrograms. A single output unit was trained to turn on when the network was presented a note of one type (e.g., an A note) and to turn off when the network was presented a note of a different type (e.g., a B note). All of the discrimination learning conditions in the Charrier *et al.* (2005) study were successfully simulated in this fashion.

After training, featural representations of shifted notes were presented to the trained perceptrons, and the output unit responses were analyzed. Results from the perceptron response activities were compared to the classifications made by the black-capped chickadees, and strong parallels between the ANNs and the birds were found. For example, as A notes were shifted downwards in frequency both the chickadees and the perceptron responded as if they were B notes. Likewise, B notes shifted upwards in frequency were responded to as if they were A notes. Given the high degree of correspondence between perceptrons and chickadees when presented shifted notes, it appeared that the networks provided a solid first step in developing a model to determine how birds perform these classifications.

However, in both the Nickerson *et al.* (2006) simulation study, and in the Charrier *et al.* (2005) chickadee study, when novel notes were created by shifting frequencies, all of the frequency features were affected by the shift, and none of the duration features were modified at all. As a result, in these studies it was not possible to examine how individual summary features independently contributed to note-type recognition. This is unfortunate, because previous studies do indicate that different features are important for classifying different notes types. For example, Dawson *et al.* (2006) presented the non-D note types to a multilayer perceptron as a categorization task. They found that only four of the nine summary features were reliable identifiers of note type when considered across a number of predictor methods (ANN, multiple regression, and linear discriminant analysis).

There is another general and prevalent issue in animal learning and cognition that was not addressed by either the Nickerson *et al.* (2006) or the Charrier *et al.* (2005) study, namely the effect of learning context on performance. It has been well established that contextual cues have strong and pervasive effects on learning performance. For instance, as so elegantly predicted by Miller's (2006) comparative hypothesis, context, when salient, can control conditioned responding to such an extent that putative CSs are ignored by animals and the context of the conditioning apparatus itself controls responding. Further, Bouton's (2004) research on extinction has shown in a wide variety of experiments the effects that context can have on the recovery of a previously extinguished behavior.

Context has also been shown to be crucial to chickadee's processing of auditory stimuli. Lee *et al.* (2006) revealed that chickadees are sensitive to context when testing generalization in an absolute pitch discrimination task. Lee *et al.* (2006) found that birds tested with novel tones in the absence of training tones failed to maintain their previously learned discrimination and did not show evidence of transfer to the novel tones. In contrast, birds that were tested with both training and novel tones presented in the same session, showed solid evidence of generalization from training to novel tones. Clearly, as in other studies of learning, context plays a vital role in songbird auditory perception. The second aim, then, of the current research was to determine what, if any, were the effects of discrimination training context on perceptron performance both in the acquisition of the task, and in the generalization to novel, manipulated notes.

In summary, the purpose of the present simulation study is twofold. First, we wanted to more directly examine the contribution of each individual summary feature to note discrimination, including duration features. To do so we manipulated each acoustic feature independently in each note type in an effort to specify the features potentially used by birds to classify notes into types, and more specifically, to establish which features determine how the shifted notes were categorized by a perceptron. Second, we set out to determine whether learning context affected either acquisition or transfer to manipulated notes by the perceptrons. To accomplish our second aim, we included training contexts not used in either Charrier *et al.* (2005) or in Nickerson *et al.* (2006) and compared performance during initial training and during generalization tests.

II. METHODS

A. Training stimuli

Sixty notes (20 A, 20 B, and 20 C) previously utilized by Charrier *et al.* (2005) were used to create the training sets for the perceptrons. D notes were not included in this set because they are acoustically dissimilar from non-D notes, and cannot be represented using the same summary features used to represent non-D notes. In addition, D notes are rarely confused with other note types (Charrier *et al.*, 2004; Sturdy *et al.*, 2000).

The 60 training notes were randomly selected from 60 different calls recorded from birds from Colorado, and On-

tario, Canada. An equal representation of both locations was included, and all notes were of high quality. Call-note amplitude was equalized using SIGNAL v 4.0 (Engineering Design, Berkeley, CA). As well, each call was bandpassed filtered (1–10 kHz) to remove background noise.

Each note was represented as a vector composed of seven numerical values. These values were of seven summary features that have been used in several previous studies to represent chick-a-dee call notes (Dawson *et al.*, 2006; Nickerson *et al.*, 2006; Nowicki and Nelson, 1990). These features were start frequency (SF), peak frequency (PF), end frequency (EF), maximum (loudest) frequency (Fmax), ascending duration (AD), descending duration (DD), and total duration (TD). These particular features were selected because of their mutual independence. That is, any one of these features could be manipulated without affecting any of the other six features.

The spectral measurements (SF, PF, EF, Fmax) were quantified in digital spectrograms (window size 1024 points, frequency precision=43 Hz) using a cutoff amplitude of –35 dB relative to peak amplitude. We used the loudest harmonic of A and B notes (i.e., the carrier frequency), and the fundamental frequency of C notes in our measurements. The temporal measurements (AD, DD, TD) were likewise quantified in a digital spectrogram (window size 256 points, duration precision=5.8 Hz).

After obtaining the feature measurements, we converted each value into a z-score. This reduced the range of values presented to the input units of the perceptrons while maintaining the essential characteristics of the raw data (Dawson *et al.*, 2006).

B. Test stimuli

The 60 training stimuli were also used to create a large set of shifted stimuli (4200 notes) to be used to test a perceptron's responses after it had learned to discriminate training stimuli. Each training stimulus was used to create 70 different shifted notes. Each shifted note was identical to the training stimulus with the exception that the value of one of its features was shifted either upwards or downwards in value. Each feature was shifted by one of ten different amounts (± 0.5 , ± 1.0 , ± 1.5 , ± 2.0 , and ± 2.5 s.d.). The 70 shifted variants of each note resulted by applying each of these ten shifts to each of the training note's seven features.

When shifted notes were created, they were created from the raw feature summary vectors. That is, they were not created from the z-score representation of note features. The approach to creating the shifted notes was analogous to that described by Charrier *et al.* (2005). The s.d. values for the shifts of each feature were dependent on note type, and were based on values previously reported by Charrier *et al.* (2004; see Table I). Once the shifted notes were created from the raw feature values, they were converted to z-scores using the means and s.d. of the original training stimuli.

C. Network architecture

All of the perceptrons trained in the simulations reported in the following consisted of seven input units (one for each

TABLE I. Provided here are the s.d. values that were used for shifting individual features according to note and feature type. The value was multiplied by each of the s.d. increments, and then added to or subtracted from the original value measured from the spectrogram of that note depending on direction of shift. For example, the SF s.d. values for an A note were (824×0.5) , (824×1) , (824×1.5) , (824×2) , and (824×2.5) , and each of these products, added to and subtracted from the original measured A note SF value, gave the shifts across s.d.

Feature	A	B	C
SF (Hz)	824	518	144
PF (Hz)	381	660	938
EF (Hz)	631	451	126
Fmax (Hz)	397.6	647.3	1009
AD (ms)	13.1	5.7	5.1
DD (ms)	6.3	4.8	2.2
TD (ms)	17.4	8.5	4.4

note feature) and a single output unit. The output unit used the logistic equation to convert the sum of the weighted signals from the input units (the net input) into an activation value that ranged between 0 and 1. The logistic equation is a sigmoid-shaped function that is frequently used in ANNs to apply a nonlinear transformation to net inputs (e.g., Rumelhart *et al.*, 1986).

D. Network training on unshifted notes

The training set of 20 A, 20 B, and 20 C notes was used to replicate the training conditions used by Charrier *et al.* (2005). For example, the discrimination group A+B- was comprised of 20 A notes and 20 B notes. Perceptrons that served as subjects in this condition were trained to turn their output unit “on” when presented A notes and to turn their output unit “off” when presented B notes. This approach was used to create four different discrimination learning conditions that paralleled the four studied by Charrier *et al.* (2005): A+B-, A-B+, B+C-, and B-C+. Ten different networks served as “subjects” in each of these four conditions. Each subject was different from another because prior to learning, all of the connection weights in a perceptron were set to random values selected from the range between -0.1 and 0.1.

A gradient descent rule was used to train each perceptron (Dawson, 2004, 2005). This rule modifies connection weights in such a way that output unit error decreases as patterns are repeatedly presented. Connection weights were modified by this learning rule after each note was presented during training. Training proceeded by presenting patterns in “epochs.” In a single epoch, the network was trained on each of the 40 stimuli being used for a particular condition. The order of pattern presentation was randomized every epoch. The learning rate for the simulations was set at 0.5. The training was conducted with the Rosenblatt program (Dawson, 2005).

Training proceeded until a perceptron generated a “hit” for each of the 40 stimuli that made up its training set. A hit was defined as the output unit producing activity of 0.9 or higher when the desired response to a pattern was on, or producing activity of 0.1 or lower when the desired response

to a pattern was off. Every perceptron that was trained in these simulations successfully learned the correct pattern of discrimination responses to all 40 patterns that it was presented.

E. Network testing on shifted notes

After a perceptron had successfully learned to discriminate one type of note from another (e.g., A+B-), it was then presented each of the 4200 shifted notes. Instead of training the perceptrons on these notes, we merely recorded the output unit’s activity to each of these shifted stimuli. The responses of the perceptrons to these novel, shifted notes could then be used to explore the contribution of each individual feature to the overall responses that had previously been examined in birds (Charrier *et al.*, 2005) and in perceptrons (Nickerson *et al.*, 2006).

III. RESULTS

A. Training results

The networks trained in both AB discrimination learning conditions converged (i.e., successfully learned responses to all 40 input patterns) after an average of 818.5 epochs. The networks trained in both BC conditions converged after an average of 357.4 epochs. This difference is statistically significant ($t=365.31$, $df=18$, $p<0.001$), and is consistent with the fact that the visual appearance of the sound spectrograms (and thus the summary features) for A and B notes are more similar to one another than to the spectrograms of C notes. Because of this, it is not surprising to find that one discrimination task is harder to learn than another.

B. Testing results

1. Analysis of network responses to shifted notes

After discrimination training, each trained network was presented each of the 4200 shifted notes, and the activity produced in a network’s output unit was recorded. These activities were then averaged over the ten perceptrons that served as subjects in each condition. These averaged responses were plotted as a function of shifted feature and note type, producing 21 different graphs. Each graph depicts how a network that was trained in a particular discrimination condition responded to one of the two note types presented in that set. These responses were plotted as a function of the shifts of one of the note’s seven descriptive features. For instance, one graph would illustrate the responses of networks trained in the A+B- discrimination condition and other networks trained in the A-B+ condition when presented an A note that had its starting frequency (SF) shifted upwards and downwards.

One can easily convert the information derived from a graph to a brief qualitative description of the effects shifting the values of a note’s feature. Table II provides a summary of these qualitative descriptions for all of the graphs that were produced in the analyses of network responses to notes with shifted features. The notation used in Table II indicates whether a feature shift resulted in a note becoming more like

TABLE II. A qualitative account of specific training conditions and the corresponding network generalizations to each note type with individually shifted features.

		SF		PF		EF		Fmax		AD		DD		TD	
		-	+	-	+	-	+	-	+	-	+	-	+	-	+
A	A+B-	↓ A	...	↓ A	↓ A	...	↓ A	↓ A	↓ A	...
	A-B+	↑ B	...	↑ B	↑ B	...	↑ B	↑ B	↑ B	...
B	A+B-	...	↑ A	↓ A	↑ A	↑ A	↑ A	↑ A
	A-B+	...	↓ B	↑ B	↓ B	↓ B	↓ B	↓ B
	B+C-	↓ B	...	↓ B	...	↓ B	↓ B	...	↓ B	...	↓ B
	B-C+	↑ C	...	↑ C	...	↑ C	↑ C	...	↑ C	...	↑ C
C	C-B+	↑ B	↑ B	↑ B	↑ B	...
	C+B-	↓ C	↓ C	↓ C	↓ C	...

one of the note types (e.g., ↑ A indicates that a shift causes a note to be more A-like), less like one of the note types (e.g., ↓ A), or had no effect at all (---).

An examination of Table II reveals several different regularities. First, when a feature shift affected network responses, this effect depended upon the specific reinforcement condition that perceptrons were trained in. For example, when SF was shifted downward in an A note, this led to the note being less A-like for perceptrons trained in the A+B- condition, and to the note being more B-like for perceptrons trained in the A-B+ condition. A second property is that shifts of some features were more likely to affect network responses than were shifts of other features. For instance, shifts of PF produced changes in network performance in ten different cells of Table II, while modifications of DD only produced changes in two of the table's cells. A third property revealed by Table II is that the effect of a feature shift was context dependent, in the sense that the effect depended upon which two types of notes were involved in discrimination learning. For instance, manipulations of SF had effects on networks that were trained to discriminate A notes from B notes, but did not affect networks that were trained to discriminate C notes from B notes.

2. Analysis of network connection weights

Why do the perceptrons respond in the fashion that is summarized in Table II? Ultimately, the responses of a perceptron are governed by the weights of the connections between the input units and the output unit. An input feature that has a strong excitatory influence on the output unit will have a large positive weight. An input feature that has a

strong inhibitory influence on the output unit will have a large negative weight. An input feature that has little influence on the output unit will have a near-zero weight. An examination of the connection weights of the trained perceptrons should provide insight into the regularities that are apparent in Table II.

Table III presents the connection weights for each discrimination learning condition averaged over the ten perceptrons that served as "subjects" in the condition. The pattern of weights presented in Table III is consistent with the pattern of network responses that was summarized in Table II, as is detailed in the following.

First, Table III shows that SF and Fmax were important when differentiating between A and B note types because they had the largest absolute values of connection weights in the AB discrimination learning conditions. Similarly, AD and Fmax were the most important features used to differentiate between B and C note types. In general, though, Fmax and PF were the most significant indicators of note type because they had large connection weights in all of the discrimination learning conditions.

Second, several summary features were associated with near zero connection weights, and therefore had little importance for differentiating between note types. This was particularly true for DD in the AB note discrimination conditions and for SF in the BC note discrimination conditions. Overall, DD was the least relevant feature in differentiating between note types.

Third, there was a distinct reflection of connection weights depending upon which note was reinforced in a particular discrimination learning condition. For example, con-

TABLE III. Connection weight average values and their standard deviations in each training set for the association between an input unit and the output unit.

	SF	PF	EF	AD	DD	TD	Fmax
A+B-	7.86 (0.032)	3.47 (0.026)	-1.73 (0.023)	1.86 (0.029)	0.04 (0.017)	4.10 (0.008)	-6.79 (0.028)
A-B+	-7.83 (0.018)	-3.48 (0.011)	1.74 (0.018)	-1.85 (0.014)	-0.04 (0.014)	-4.10 (0.004)	6.82 (0.016)
B+C-	-0.23 (0.031)	3.71 (0.020)	2.70 (0.015)	-4.46 (0.041)	-1.28 (0.007)	-2.80 (0.020)	4.22 (0.031)
B-C+	0.21 (0.043)	-3.70 (0.018)	-2.69 (0.024)	4.45 (0.035)	1.28 (0.010)	2.80 (0.015)	-4.23 (0.037)

sider the PF feature in the A+B- condition, which produced an average connection weight of 3.47. In the complementary A-B+ condition, this feature had a weight of nearly the same size, but of opposite sign (i.e., -3.48). A similar pattern is evident for all of the features, in either the AB or the BC discriminations.

IV. DISCUSSION

Previously, Charrier *et al.* (2005) studied how black-capped chickadees respond to artificial notes that had their entire spectrogram shifted up or down in frequency. They reported a systematic pattern of note misclassifications by these animals. Nickerson *et al.* (2006) replicated this study in artificial neural networks, and found a very strong agreement between the behavior of their networks and the behavior of the birds. There were two limitations of both of these studies. First, by shifting the frequency of the entire note, all of the frequency features were manipulated at the same time. Second, by only shifting note frequency, none of the duration features of the spectrogram were manipulated at all. The purpose of the above-reported simulations was to examine the response of artificial neural networks that had every note feature (frequency and duration) manipulated independently. In the following discussion, we first consider the effects of manipulating the frequency features. We then consider the effects of altering the duration features. Finally, we reflect upon the implications of these simulated results.

A. Frequency manipulations

An examination of Table II reveals several interesting patterns of results concerning the independent manipulation of the frequency features of the notes. The most general of these is that changing any of the frequency features (SF, PF, EF, or Fmax) produced changes in the responses in at least some of the trained perceptrons. That is, none of the columns of Table II for these features are empty.

Of more interest is the fact that not all of these features were of equal importance. In particular, Table II demonstrates that ten different effects were observed when PF was manipulated, and eight different effects were detected when Fmax was changed. In contrast, altering SF and EF produced only four observable effects apiece.

Of most interest is the fact that the effect of manipulating a frequency feature was highly context dependent. That is, whether a manipulation had an effect, and the nature of the effect that was produced, depended critically upon the type of discrimination learning that a perceptron had been given prior to being presented the manipulated notes. For example, consider SF. When networks had undergone AB discrimination learning, decreases in SF made A notes more B-like, and increases in SF made B notes more A-like. However, when networks had undergone BC discrimination learning, manipulations of SF had no effect on network responses at all.

The context dependence of PF manipulations is much more striking. First, consider the context of AB discrimination learning. Decreases in PF made both A notes and B notes more B-like in this context, while increases in PF made

B notes more A-like. In the context of BC discrimination learning, decreases in PF made B notes more C-like, while increases in PF made C notes more B-like. Notice the striking context dependence on responses to B notes: In one context, decreases in PF produce network responses that indicate that B notes are more B-like, but in the other context these same decreases generate network responses that indicate that B notes are less B-like.

The context dependence of EF manipulations is also evident. In the AB context, increases in EF make A notes more B-like, but no other effects are evident. In particular, manipulating EF does not affect B notes in this context. However, in the BC context, decreases in EF make B notes more C-like.

Finally, consider the pattern of results produced by Fmax manipulations. In the AB context, increases in Fmax made A notes more B-like, and decreases in Fmax made B notes more A-like. In the BC context, increases in Fmax made C notes more B-like while decreases in Fmax made B notes more C-like. These results are the least context-dependent of the frequency manipulations, in the sense that there was a qualitatively consistent pattern of effects on B notes across the two learning contexts.

How do these simulated frequency results compare to results from black-capped chickadees? In their extensive bioacoustic analysis of the chick-a-dee call, Charrier *et al.* (2004) examined the extent to which each summary feature was important for determining a particular, and unmodified, note type (A, B, or C). They did this by computing the potential for note-type coding (PNTC) for each feature. The PNTC is an adaptation of the method for the potential for individual identity coding in animal vocalizations (Charrier *et al.*, 2004; Robisson *et al.*, 1993), and is computed for a feature by taking the ratio of its between-note-type variation to its within-note-type variation. A high PNTC value for an acoustic feature suggests that the feature is crucial for classifying the note as belonging to one of the three note types (A, B, and C). Charrier *et al.* (2004) found that SF had a particularly high PNTC value (2.7), and that EF had a moderately high PNTC value (1.6). They also found that PF and Fmax had PNTC values near unity (1.06 and 1.03, respectively), indicating that these features were not crucial for distinguishing types of notes. All of these results are at odds with the above-reported simulation results. However, one factor that is absent from PNTC calculations is a learning context. The simulation results strongly suggest that which features are important for note discrimination depends upon which two notes are being discriminated.

B. Duration manipulations

The general pattern of results revealed for frequency features can also be seen in Table II for the three different duration features of ascending duration (AD), descending duration (DD), and total duration (TD). First, manipulations of all three features produced changes in behavior in some of the perceptrons. Second, some of the duration features appeared to be more important than others: AD and TD both

reveal eight effects in Table II, while DD only reveals two. Third, the effects of manipulating a note's features once again appear to be context dependent.

The most striking example of context dependence is evident in manipulations of DD. In the AB context, altering this feature does not generate any changes in perceptron behavior. In the BC context, increases in DD cause B notes to become more C-like, but no other effects are evident.

Manipulations of the other two duration variables appear to be less context dependent, in the sense that these manipulations appear to have similar effects on B notes in both discrimination learning contexts. In the AB context, increases in AD cause B notes to become more A-like, while decreases in AD cause A notes to become more B-like. In the BC context, increases in AD cause B notes to become more C-like, while decreases in AD cause C notes to become more B-like. The identical pattern of results is revealed in the manipulations of TD in the two contexts.

In terms of relating these results to previous studies, Charrier *et al.* (2004) found that TD had a PNTC value of 1.6, AD had a PNTC value of 1.5, and DD had a PNTC value of 1.3. The fact that DD had a lower PNTC value than did the other two duration features is consistent with our finding that DD did not play as important a role for the perceptrons as did the other two features. However, as with the frequency effects, discussed earlier, these PNTC values did not take context dependence into account.

C. Implications

What are the implications of the simulation results that we have reported? One implication comes from a standard empiricist approach to evaluating models: A model is only as good as the number of new experiments that it generates. Our simulations have revealed some general regularities that need to be explored using birds as subjects. First, the perceptrons reveal that their behavior can be changed by altering individual spectrogram features. Second, the perceptrons suggest that some of these manipulations are more powerful—or at least more likely to produce effects—than are others. Third, the perceptrons indicate that the effect of manipulating a particular feature is strongly context dependent.

Here we do not claim that black-capped chickadees can independently modify the features in their call notes in exactly the same manner that we have in this report. In fact, there is a good likelihood that they may not be able to manipulate individual features in their notes without also affecting other features in the process due to anatomical and physiological constraints on the vocal apparatus (reviewed in Podos and Nowicki, 2004; Suthers, 2004). That being said, there is natural variation in all of the acoustic features in normal chickadee call notes that we manipulated here to produce modified stimuli for our networks (Charrier *et al.*, 2004). Therefore, in spite of the fact that birds may not necessarily be able to produce notes of the type that we are using to test network performance they do produce notes in a variable manner. Further, what we are attempting to do here is to determine which feature (or features), observed to vary in

nature during normal vocal production, is (or are) the most important for controlling note type discrimination and categorization. That being said, further studies on anatomical and physiological performance limits on vocal production in chickadees would be an interesting avenue for future research.

The second implication reveals another important contribution that simulation studies can make to empirical research. Animal research is both costly and time consuming. It would be highly advantageous to have *a priori* knowledge about what effects are likely to be revealed in a study that uses animals as subjects. The above-reported simulations provided hypotheses about the potential effects of a large number of feature manipulations—recall that the test set consisted of 4200 different notes. To the extent that the perceptrons provide an applicable model of note discrimination in birds (a fact that was previously established by Nickerson *et al.*, 2006), they have permitted a search of a staggeringly large set of manipulations in an attempt to identify experimental variables that are likely to produce results. Our current work involves testing the hypotheses generated by these simulations using real birds as subjects.

The third implication, that context plays a role in learning and generalization by the perceptrons, places this research into a broader framework of contextual effects on learning and cognition. As noted in Sec. I, context can overshadow CSs, and cause previously extinguished behavior to reemerge. Moreover, and important for the current research, songbirds performing auditory discriminations appear to be very sensitive to the context during generalization tests with *pure tonal* stimuli (Lee *et al.*, 2006), to the extent that without a familiar training context, birds fail to show generalization to novel stimuli. Here, modeling *call note* discriminations using perceptrons, we show a similar and striking effect of training context on acquisition, with significant differences among the different discrimination tasks, and of generalization, with differential responding to transfer notes contingent upon the discrimination context of the perceptrons. Whether this contextual effect is observed in birds performing these same discriminations presents an exciting future avenue of research.

ACKNOWLEDGMENTS

This research was supported by an NSERC discovery grant, start-up funding from the University of Alberta, and an Alberta Ingenuity New Faculty Grant to C.B.S. and NSERC and SSHRC research grants to M.R.W.D. C.M.N. was supported in part by STEP funding from the University of Alberta and the Province of Alberta and L.L.B. was supported by an Alberta Ingenuity Studentship. This research was approved by the University of Alberta Biological Sciences Animal Care Committee (Protocol No. 351201, originally approved on 14 November 2001) and the University of Calgary Life and Environmental Sciences Animal Care Committee (Protocol No. BI2001-028 originally approved on 29 November 2001). Birds were captured under an Environment Canada, Canadian Wildlife Service Scientific permit (Permit No. WSA-1-02, originally approved on 25 January 2002) and

Alberta Fish and Wildlife Research permits (Nos. 4619 GP, 4621 GP, and 8734 GP) and collection licenses (Nos. 088 CN, 089 CN, and 147 CN) that were originally approved on 21 December 2001. Correspondence can be addressed to C.B.S. The authors thank the editor and two anonymous reviewers for their comments and suggestions.

- Bouton, M. E. (2004). "Context and behavioral processes in extinction," *Learn. Memory* **11**, 485–495.
- Charrier, I., Bloomfield, L. L., and Sturdy, C. B. (2004). "Note types and coding in parid vocalizations I. The chick-a-dee call of the black-capped chickadee (*Poecile atricapilla*)," *Can. J. Zool.* **82**, 769–779.
- Charrier, I., Lee, T. T. Y., Bloomfield, L. L., and Sturdy, C. B. (2005). "Acoustic mechanisms of note-type perception in black-capped chickadee (*Poecile atricapillus*) calls," *J. Comp. Psychol.* **119**, 371–380.
- Charrier, I., and Sturdy, C. B. (2005). "Call-based species recognition in black-capped chickadees," *Behav. Processes* **70**, 271–281.
- Dawson, M. R. W. (2004). *Minds And Machines: Connectionism And Psychological Modeling* (Blackwell, Malden, MA).
- Dawson, M. R. W. (2005). *Connectionism: A Hands-On Approach*, 1st ed. (Blackwell, Oxford, UK).
- Dawson, M. R. W., Bloomfield, L. L., Charrier, I., and Sturdy, C. B. (2006). "Statistical classification of black-capped (*Poecile atricapillus*) and mountain chickadee (*Poecile gambeli*) call notes," *J. Comp. Psychol.* **120**, 147–153.
- Dawson, M. R. W., Charrier, I., Sturdy, C. B. (2006). "Using an artificial neural network to classify black-capped chickadee (*Poecile atricapillus*) call note types," *J. Acoust. Soc. Am.* **119**, 3161.
- Dawson, M. R. W., and Spetch, M. L. (2005). "Traditional perceptrons do not produce the overexpectation effect," *Neur. Info. Proc. — Letters and Reviews*, found at <http://nlp-lr.info/V07N01/V07N01.htm>, **7**, 11–17.
- Ficken, M. S., Ficken, R. W., and Witkin, S. R. (1978). "Vocal repertoire of black-capped chickadee," *Auk* **95**, 34–48.
- Ficken, M. S., Hailman, E. D., and Hailman, J. P. (1994). "The chick-a-dee call system of the Mexican chickadee," *Condor* **96**, 70–82.
- Freeberg, T. M., and Lucas, J. R. (2002). "Receivers respond differently to chick-a-dee calls varying in note composition in Carolina chickadees," *Anim. Behav.* **63**, 837–845.
- Hailman, J. P. (1985). "Ethology, zoosemiotic and sociobiology," *Am. Zool.* **25**, 695–705.
- Hailman, J. P., and Ficken, M. S. (1986). "Combinatorial animal communication with computable syntax—Chick-a-dee calling qualifies as language by structural linguistics," *Anim. Behav.* **34**, 1899–1901.
- Lee, T. T. Y., Charrier, I., Bloomfield, L. L., Weisman, R. G., and Sturdy, C. B. (2006). "Frequency-range discriminations and absolute pitch in black-capped chickadees (*Poecile atricapillus*), mountain chickadees (*Poecile gambeli*), and zebra finches (*Taeniopygia guttata*)," *J. Comp. Psychol.* **120**, 217–228.
- Mammen, D. L., and Nowicki, S. (1981). "Individual-differences and within-flock convergence in chickadee calls" *Behav. Ecol. Sociobiol.* **9**, 179–186.
- Miller, R. R. (2006). "Challenges facing contemporary associative approaches to acquired behavior," *Comparative Cognition & Behavior Reviews*, **1**, 77–93. Retrieved from <http://psyc.queensu.ca/ccbr/index.html>
- Minsky, M., and Papert, S. (1988). *Perceptrons*, 3rd ed. (MIT, Cambridge, MA).
- Nickerson, C. M., Bloomfield, L. L., Dawson, M. R. W., and Sturdy, C. B. (2006). "Artificial neural networks that discriminate notes from the chick-a-dee call of *Poecile atricapillus*: The effect of pitch transformations," *J. Acoust. Soc. Am.* **120**, 1111–1117.
- Nowicki, S. (1983). "Flock-specific recognition of chickadee calls," *Behav. Ecol. Sociobiol.* **12**, 317–320.
- Nowicki, S., and Nelson, D. A. (1990). "Defining natural categories in acoustic-signals—Comparison of 3 methods applied to chick-a-dee call notes," *Ethology* **86**, 89–101.
- Podos, J., and Nowicki, S. (2004). "Performance limits on birdsong," in *Nature's Music: The Science of Birdsng*, edited by P. Marler and H. Slabbekoorn (Elsevier, New York), pp. 318–342.
- Robisson, P., Aubin, T., and Bremond, J. C. (1993). "Individuality in the voice of the emperor penguin *Aptenodytes-Forsteri*—Adaptation to a noisy environment," *Ethology* **94**, 279–290.
- Rosenblatt, F. (1962). *Principles of Neurodynamics* (Spartan, Washington, DC).
- Rumelhart, D. E., Hinton, G. E., and Williams, R. J. (1986). "Learning representations by back-propagating errors," *Nature (London)* **323**, 533–536.
- Sturdy, C. B., Phillmore, L. S., and Weisman, R. G. (2000). "Call-note discriminations in black-capped chickadees (*Poecile atricapillus*)," *J. Comp. Psychol.* **114**, 357–364.
- Suthers, R. A. (2004). "How birds sing and why it matters," in *Nature's Music: The Science of Birdsng*, edited by P. Marler and H. Slabbekoorn (Elsevier, New York), pp. 272–295.
- Templeton, C. N., Greene, E., and Davis, K. (2005). "Allometry of alarm calls: Black-capped chickadees encode information about predator size," *Science* **308**, 1934–1937.
- Yaremchuk, V., Willson, L. R., Spetch, M. L., and Dawson, M. R. W. (2005). "The implications of null patterns and output unit activation functions on simulation studies of learning: A case study of patterning," *Learn Motiv* **36**, 88–103.