Proceedings of Meetings on Acoustics

Volume 19, 2013 http://acousticalsociety.org/

!

!

ICA 2013 Montreal
Montreal, Canada
2 - 7 June 2013

Animal Bioacoustics
Session 3aAB: Perceiving Objects I

3aAB2. Recognizing objects from multiple orientations using dolphin echoes

Caroline M. DeLong*, Amanda L. Heberle, Kayla Mata, Heidi E. Harley and Whitlow W. Au

***Corresponding author's address: Dept. of Psychology, Rochester Institute of Technology, 18 Lomb Memorial Dr., Rochester, NY 14623, cmdgsh@rit.edu**

Object constancy, the ability to recognize objects despite changes in orientation, has not been well studied in the auditory modality. Dolphins use echolocation for object recognition, and objects ensonified by dolphins produce echoes that can vary significantly as a function of orientation. In four experiments, human listeners had to classify echoes from objects ensonified with dolphin signals. Participants were trained to discriminate among the objects using an 18-echo stimulus from a 10 degree range of aspect angles, then tested with novel aspect angles across a 60 degree range. In the first two experiments the three objects varied in material, size, and shape. Participants were typically successful recognizing the objects at all angles (M = 78%). In Experiment 3, the three objects had the same material but different shapes. Participants were often unsuccessful recognizing the objects at all angles (M = 46%). In Experiment 4, participants had to classify echoes from four fish species across a wider range of angles (330 degrees). Preliminary results show overall poor performance (M = 45%). These results suggest that object characteristics play a role in whether performance is more view-dependent or view-invariant. These studies can provide insight into the process dolphins use to identify objects.

Published by the Acoustical Society of America through the American Institute of Physics

Received 18 Jan 2013; published 2 Jun 2013

INTRODUCTION

The ability to recognize an object from any viewpoint is called object constancy. It has been well-studied in the visual modality (e.g., Jolicoeur & Humphrey, 1998), but rarely investigated in the auditory domain (DeLong, Bragg, & Simmons, 2008). Bottlenose dolphins can identify objects using echolocation, and encounter objects from varying viewpoints as they move in both the horizontal and vertical plane. For example, a dolphin needs to be able to recognize a particular fish species from the head, tail, or broadside. Only two prior studies have explored auditory object constancy in dolphins (Au & Turl, 1991; Nachtigall, Murchison, & Au, 1980). In these studies, a dolphin was trained to discriminate among objects that varied in material or shape at one aspect angle then tested with a few novel aspect angles. Nachtigall et al.'s (1980) study suggests that the dolphin only successfully discriminated among objects at novel aspect angles when the echoes of the two objects were highly dissimilar. However, neither study examined in depth the potential individual echoic acoustic features used by the dolphin for object recognition.

One way to determine salient acoustic features in echoes is to use a comparative approach. Humans and dolphins share many similarities in auditory perception. Both humans and dolphins can discriminate between sounds that differ in intensity by 1 dB (Green, 1993; Vel'min et al., 1975 c.f. Au, 1993) and the frequency discrimination abilities of humans and dolphins for tonal stimuli are comparable in the range of best hearing for each species (Herman & Arbeit, 1972; Thompson & Herman, 1975; Wier et al., 1977). Human listeners presented with slowed-down echoes are capable of discriminating among objects and reporting potential auditory cues (e.g., Au & Martin, 1989; DeLong, Au, Harley, Roitblat, Pytka, 2007; Helweg, Roitblat, Nachtigall, Au, & Irwin, 1995). Human listeners and dolphins can discriminate among the same set of objects and their performance can be compared. Whether salient echo features reported by humans overlap with those used by dolphins can be assessed, in part, by comparing errors made by humans and dolphins. Matching error patterns would suggest use of similar features whereas mismatching error patterns would imply the use of different features.

Several prior studies have shown the usefulness of this comparative approach. In one study, human listeners were presented with echoes of aspect-dependent objects that had been used in a dolphin discrimination task (DeLong et al., 2007). The objects varied in size, shape, material, and/or texture. For these object sets, echoes from different orientations of the same object varied significantly more than the echoes from different objects (DeLong, Au, Lemonds, Harley, & Roitblat, 2006). In two experiments, the human listeners performed as well or better than the dolphin at discriminating objects and reported the salient acoustic cues. The error patterns of the humans and the dolphin were compared to determine which acoustic features reported by the humans were likely to have been used by the dolphin. The results indicated that the dolphin did not appear to use overall echo amplitude but that it attended to the pattern of changes in the echoes across different object orientations. Having information in multiple echoes gleaned from multiple object orientations appeared to be particularly advantageous when discriminating among objects that varied in shape.

In a recent study, DeLong, Kannyo, Benoit-Bird, and Au (2013) presented human listeners with echoes recorded from four fish species, including sea bass (*Dicentrarchus labras)*, pollack (*Pollachius pollachius)*, grey mullet (*Chelon labrosus)*, and Atlantic cod (*Gadus morhua*). These fish echoes were recorded using both dolphin-like signals and porpoise-like signals, and all lateral aspect angles of the fish were collected. In Experiment 1, participants were asked to discriminate among the fish species using an echo sequence of 150 echoes that included all orientations of the fish (360 degrees) in half the trials, and in the other half of the trials participants had to discriminate among the fish using multiple echoes from only a single randomly-selected aspect angle of the fish. Participants performed well on the trials where they had access to all orientations of the fish, but they performed very poorly when asked to discriminate among the fish given echoes from only one aspect angle. In Experiment 2, participants were asked to discriminate among the fish species using echoes that spanned 90 degrees (head to broadside aspects) and 180 degrees (head to tail aspects). There was no significant difference in performance between the 90 degree condition and 180 degree condition, which suggests that there was sufficient information for discrimination in the 90 degree condition. In Experiment 3, participants were asked to discriminate among the fish species using unaltered echoes and amplitude-normalized echo trains. Participants were able to succeed with both types of echoes, which indicates that they were able to utilize cues other than loudness. This study suggests that hearing echoes from multiple object orientations is clearly important for listeners to identify these fish species, and that the multiple auditory features they used for discrimination could be modified based on available information.

The purpose of the current study was to investigate discrimination cues used by human listeners in an auditory object constancy task. The experiment was similar to other object constancy studies with dolphins (Au & Turl, 1991; Nachtigall et al., 1980), except that the objects varied along multiple traits instead of a single trait. The participants were presented with echo sequences from three objects that varied in size, shape, and material. Participants were trained to discriminate the objects using echoes from a single aspect angle and then tested on four novel aspect angles. Participants were interviewed after each training and testing session to determine potential discriminatory cues. We hypothesized that the participants would be able to discriminate among objects at all angles and be able to report many of the cues they relied on for discrimination.

METHOD

Participants

Twenty-six participants (14 females and 12 males) volunteered to be tested. Participants ranged in age from 19 to 49 years (M = 22.6 years). All participants were students or staff at Rochester Institute of Technology. All participants were tested for normal hearing with the Home Audiometer 2.0 hearing test (Esser Audio, 2009) and all had normal sensitivity in the frequency range of the stimuli. Participants received either extra credit for a psychology course or were paid $10.

Materials and Procedure

Echo Recordings

The stimuli used in this experiment were pre-recorded echoes from three objects; a ceramic cross, a copper figure 8, and a wooden rectangle (see Figure 1). These objects were used in a previous echolocation study with a bottlenose dolphin (DeLong et al., 2006). Echoes from the three objects were recorded using a typical bottlenose dolphin click (70 μs long with a peak frequency of about 120 kHz and a 60-kHz bandwidth; see Au, 1993). Echo recordings were made from multiple orientations of each object in which one echo was captured for each aspect (1.3 degrees apart) between -30 degrees and +30 degrees, where 0 degrees is the front face of the object. Nine echo train measurements were collected for each object. A detailed description of the echo recording apparati and procedure can be found in DeLong et al. (2006).

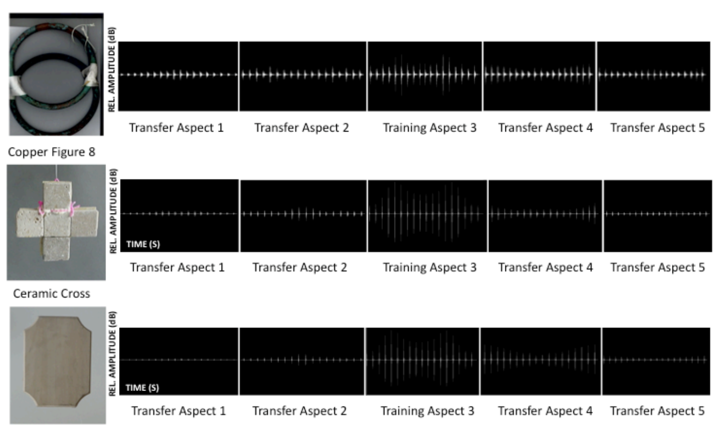

FIGURE 1. The three objects and the stimulus echo trains for each of the five aspects, including the training aspect (3) and the four transfer aspects (1, 2, 4, and 5). The 18-echo trains include two sweeps of nine echoes to mimic a dolphin sweeping its head back and forth.

Echo Formatting

The frequency span of the stimuli were slowed down to shift the spectra of the echoes into the human hearing range using Avisoft-SAS Lab Pro version 5.0.11 (Avisoft Acoustics, 2010). The original echoes were digitized at 1MHz and had center frequencies around 125 kHz. The echoes were all time stretched by a factor of 50 by converting the echoes from digital to analog at 20 kHz. The time-stretched echoes had center frequencies around 2.5 kHz. Other studies in which dolphin echoes were presented to human listeners use the same factor of 50 (Au & Martin, 1989; Helweg et al., 1995).

Echo Stimuli

Each echo recording was an echo train consisting of 45 echoes ranging from -30 to +30 degrees. The echo trains were divided into five aspects. Aspect 1 contained echoes from approximately -19 to -30 degrees, aspect 2 contained echoes from approximately -7 to -18 degrees, aspect 3 contained echoes from approximately -6 to +6 degrees (including 0 degrees, the front of the object), aspect 4 contained echoes from approximately +7 to +18 degrees, and aspect 5 contained echoes from approximately +19 to +30 degrees. Each aspect contained nine echoes, but dolphins rarely use so few echoes to discriminate among objects in a laboratory setting (Au, 1993). To increase the ecological validity of the task, the echo stimuli were formatted to mimic a dolphin sweeping its head back and forth so that there was a total of 18 echoes in each sequence of echoes. Echoes were presented consecutively in each stimulus echo train with no delays introduced between echoes (although individual echoes were audible in the stimuli). The total duration of each stimulus echo train was approximately 1 s. Since there were nine echo recordings for each object, nine exemplars of stimulus echo trains were produced for each of the five object aspects (three objects x nine exemplars x five aspects = 135 total echo train stimuli).

Training and Testing

The participants were tested with an iMac desktop computer (Apple, 2009) that channeled the echo stimuli via two Bose on-ear headphones (Bose, 2009) to the participant and experimenter. The participant sat with their back to the experimenter and was not able to see the computer screen or the experimenter. The echo stimuli were played using Quicktime version 7.6.6 (Apple, 2010). Participants were tested individually by a single experimenter in a quiet room. After participants passed a brief hearing test, they heard a set of instructions and were given a sound vocabulary sheet with terms to describe the echo stimuli. The participants were instructed to use these terms and their definitions and were allowed to reference the sound vocabulary sheet throughout testing. Five terms and their operational definitions were given: “loudness” (how loud or soft the sound is, related to the sound’s amplitude or intensity), “pitch” (how high or low the sound is, related to the sound’s frequency), “length” (how long each echo is in duration), “timbre/sound quality” (the property in musical tones that makes it possible to distinguish one instrument from another), and “pattern of change in echoes across the echo train” (the change in the echoes within the echo train may form a different pattern for different objects). This ensured that all participants had the same set of descriptive tools to describe differences between echo stimuli in the interviews that followed each training and test session. Participants were told they may also use other terms that were not listed on the vocabulary sheet.

Participants were then played three pure tone sounds with the same pitch and timbre demonstrating a loudness difference (the tones rose in volume), two pure tone sounds demonstrating a difference in pitch (880 Hz tone vs. 220 Hz tone played at the same volume by a piano), three tones demonstrating a differences in timbre (middle C played by a french horn, trumpet, and soprano sax at the same volume). Participants viewed photographs of the ceramic cross (6.3 cm x 8.1 cm), copper figure 8 (7 cm x 8.2 cm), and wood rectangle (6.7 cm x 8.2 cm) throughout the study. All participants listened to the echo stimuli at the same overall volume, with the exception of one participant who preferred to listen to the echoes at two-thirds volume (for all participants, the original amplitude relations between echoes in an echo train and between objects was retained).

First, participants received three training sessions. In each training trial, the experimenter played a stimulus echo train, the participant gave a verbal response (the name of the object), and the experimenter provided feedback (correct/incorrect). If the participant was correct, s/he would move onto the next trial. If the participant was incorrect, s/he was allowed to give a second answer and would again receive feedback (correct, on incorrect and the correct choice). In each trial, the participant was allowed to listen to the stimulus echo train as many times as they liked before they made their first answer (but they could not hear it again before their second answer). Participants usually requested to hear each stimulus echo train once or twice. The average number of times in a trial the participant requested to hear each stimulus across three training sessions is as follows: copper fig 8. = 1.03, wooden rectangle = 1.03, cross = 1.07.

Only aspect 3 for each object was used in the three training sessions. Each training session consisted of four blocks, with nine trials per block (three trials for each of the three objects) for a total of 36 trials per session. Stimulus echo train exemplars 1, 2, and 3 were used in the first two training sessions. The first 18 trials of training session three contained stimulus echo train exemplars 4, 5, 6 and the last 18 trials contained exemplars 7, 8, and 9. The trial order within each training block was randomized separately for each participant. There was no pause between the trials.

Participants completed the two test sessions after the training was complete. Each test session consisted of 36 trials (four blocks with nine trials per block). Each block contained the training aspect (aspect 3) and two transfer aspects; a near aspect (2 or 4), and a far aspect (1 or 5) for each of the three objects (only exemplar 1 was used in the test sessions). Test session 1 contained the following transfer aspects: cross aspect 1 and 2, copper aspect 4 and 5, and wood aspect 1 and 4. Test session 2 contained the following transfer aspects: cross aspect 4 and 5, copper aspect 1 and 2, and wood aspect 2 and 5. Half of the participants were randomly assigned to receive test session 1 first, and the other half received test session 2 first. The trial order within each test block was randomized separately for each participant. Participants were provided with feedback in each trial and allowed to select a second answer in the same manner as in the training sessions. Participants could listen to the echo train stimulus more than once before the first answer. The average number of times in a trial the participant requested to hear each stimulus across two test sessions is as follows: copper = 1.02, wood = 1.04, cross = 1.05.

Participants completed an interview immediately after each training and testing session. Participants were asked to describe how the auditory features of the echo stimuli (e.g., loudness, pitch, timbre) differed between the three objects, to report the feature that was most important to them when discriminating among the objects, and to describe whether their performance changed as the session progressed. A SONY IC-Recorder ICD-PX720 (2009) was used to record interviews. Each session and interview took approximately 10-15 minutes. All participants completed the entire experiment in approximately 60 to 80 minutes.

RESULTS

Training Sessions

Figure 2 shows the discrimination performance in the training sessions. Chance choice accuracy is 33%, because the participants could choose from among three alternatives. To determine whether the participants' performance was above chance in both the training and test sessions, a t-test was performed separately for each object in each block using one score for each participant (average performance on the block) and comparing the participants' scores to a value of 0.33. The participants' performance was significantly above chance for all objects in all blocks except for the ceramic cross in block 1. Performance accuracy was 80% or higher on the last three training blocks for all objects.

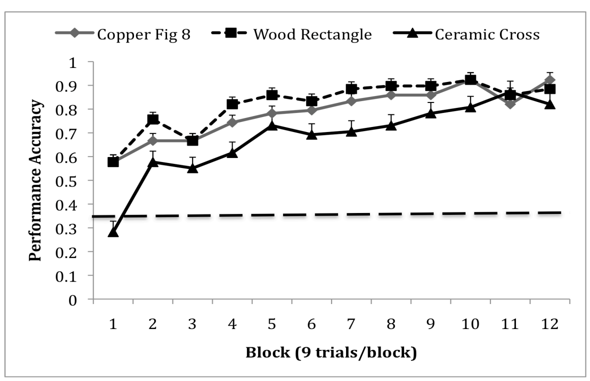

FIGURE 2. Discrimination performance in the training sessions. Training session 1 included blocks 1-8, training session 2 included blocks 9-10, and training session 3 included blocks 11-12 (all training trials included only aspect 3). Error bars show standard error. The dotted line shows chance performance (33%).

Test Sessions

The participants successfully transferred their discrimination to the four novel object aspects. Figure 3 shows the discrimination performance in the test sessions for the aspect participants were trained on (3) and the four transfer aspects (1, 2, 4, 5). The participants' performance was significantly above chance ($p <$ 0.001) for all objects in all five aspects.

A 2 (order: order 1 is test session 1 first, order 2 is test session 2 first) × 2 (test session) × 3 (object: ceramic cross, copper fig. 8, wooden rectangle) × 3 (aspect: training aspect, transfer near aspect, transfer far aspect) analysis of variance (ANOVA) with the last three factors as repeated measures, was conducted on the proportion of correct answers made by the participants. There was a significant effect of object, $F(2, 48) = 9.42$, $p < 0.001$, a significant effect of aspect, $F(2, 48) = 7.99$, $p < 0.01$, and an interaction between object and aspect, $F(4, 96) = 5.73$, $p < 0.001$. Post-hoc analyses revealed that performance for the copper fig. 8 on the transfer far aspect (70%) was significantly worse than on the transfer near aspect (91%) and the training aspect (91%). Choice accuracy did not differ significantly between the training aspect, transfer near aspect, and transfer far aspect for the ceramic cross (61%, 72%, 71%) and the wooden rectangle (83%, 82%, 72%; Newman-Keuls tests, $p < 0.05$).

There was also a significant effect of test session, $F(1, 24) = 6.34$, $p < 0.05$, and a three-way interaction between test session, object and aspect, $F(4, 96) = 2.72$, $p < 0.05$. Post-hoc analyses show that performance for the copper fig. 8 on transfer aspect 5 (62%) was significantly worse than on training aspect 3 (94%) and transfer aspect 4 (96%) in session 1. In session 2, there was no significant difference in performance for the copper fig. 8 on training aspect 3 (87%), transfer aspect 1 (77%) and transfer aspect 2 (86%). Choice accuracy did not differ significantly between the training aspect, transfer near aspect, and transfer far aspect for the ceramic cross and the wooden rectangle in test session 1 or test session 2 (Newman-Keuls tests, $p < 0.05$).

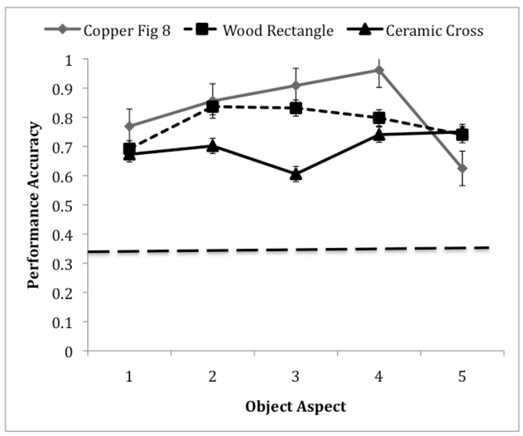

FIGURE 3. Discrimination performance in the test sessions. Aspect 3 was the training aspect, and aspects 1, 2, 4, and 5 were the transfer aspects. Error bars show standard error. The dotted line shows chance performance (33%). Performance was significantly above chance ($p < 0.001$) for all three objects in all five aspects.

Errors

Figure 4 shows the types of errors participants made in each session. Errors were divided into three types by pairing the objects; copper-wood confusions, wood-cross confusions, and copper-cross confusions. For example, when participants were presented with an echo stimulus from the copper object but reported the wood object or vice versa, those errors were classified as copper-wood confusions. The predominant error type in the test sessions was a wood-cross confusion. In the training sessions, wood-cross and copper-cross confusions were approximately equal in likelihood. There were very few copper-wood confusions in the experiment.

Reported Auditory Features

Table 1 shows the auditory features reported by human listeners after each session was completed. The participants reported using between one and five features to discriminate among the objects in the sessions. Participants reported using more features in the test sessions ($M = 3.83$) compared to the training sessions ($M = 3.21$). The features reported most often in the training sessions was pitch and timbre, whereas in the test session the features reported most often were pitch, timbre, and loudness. A majority of the participants (60%) reported pitch to be the most important feature in the training session (e.g., some participants reported ranking the objects in terms of pitch; copper had the highest pitch, ceramic was intermediate, and wood had the lowest pitch). There was no consensus on the most important feature in the test sessions;

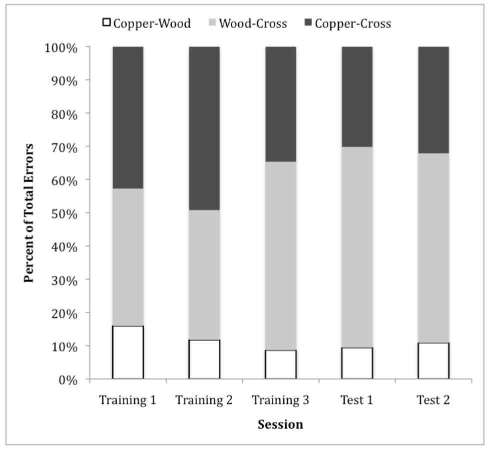

FIGURE 4. Errors (object confusions) shown for all sessions. Participants could make three types of errors: confusing the copper figure 8 and wooden rectangle, confusing the wooden rectangle and ceramic cross, or confusing the copper figure 8 and the ceramic cross. The predominant error type in the test sessions was a confusion between the wooden rectangle and ceramic cross.

pitch (36%), timbre (33%), loudness (19%), and pattern (8%) were all reported by some participants. The majority of participants (88%) reported the echo stimuli from the novel aspects sounded different than the training aspect and many noted using a different strategy in the testing sessions as compared to the training sessions. Many participants reported using more auditory features (e.g., the addition of pattern and loudness) than the ones reported for the training sessions (e.g., pitch and timbre).

TABLE 1. Echoic cues reported by participants during the interview phase after each session

Echoic Cue	Training Session 1	Training Session 2	Training Session 3	Test Session 1	Test Session 2
Loudness	12	13	14	23	25
Pitch	26	26	25	25	25
Length	8	8	10	13	10
Timbre	20	20	24	23	23
Pattern of change in echo train	14	12	14	15	16
Average number of cues reported	3.12	3.12	3.38	3.81	3.85

Note. Each cell contains the number of participants (N = 26) who reported using each cue for each session. Participants could report multiple cues for each session.

DISCUSSION

The human listeners were able to learn to discriminate among the ceramic cross, copper figure 8, and wooden rectangle using echoes from aspect three, and they successfully generalized that discrimination to four novel transfer aspects. Thus, the human listeners showed object constancy in this study. In a subsequent study (DeLong, Heberle, Harley, & Au, 2013), a different group of human listeners was presented with echoes from three objects that were all made of the same material but had different shapes. Participants had much more difficulty with the task compared to the current study and were often unsuccessful at recognizing the objects at novel test angles. This suggests that the characteristics of the objects play a role in determining whether participants achieve object constancy. The material of the object appears to generate some echoic features which are aspect-invariant. However, the shape of the object generates echoic features which are aspect-dependent.

Another factor that may influence success on an auditory object constancy task is the angular distance between the training aspect and the novel test aspects. In visual object constancy studies, performance on the transfer aspects near to the training aspect typically exceeds performance on transfer aspects far from the training aspect (e.g., Kirkpatrick, 2001). In the current study, performance did not follow the typical bell-shaped generalization gradient. This may be because the “far” aspects are only 13-24 degrees away from the training aspect. A steeper generalization gradient would be expected if the transfer far aspects were 90 or 180 degrees away from the training aspect. To resolve this issue, a future study can be conducted where participants have to classify echoes from objects across a wider range of angles.

The auditory features reported by participants suggest that they used some combination of amplitude and frequency over the course of echo trains to identify objects. However, whether these reported features are actually sufficient for identifying objects or how these features might be used to identify objects remains unclear. The auditory features that identify an object might vary across objects or across aspects, such that amplitude cues are useful for identifying a particular object at a particular aspect, but not other objects at that same aspect. Similarly, participants that reported using patterns of frequency and amplitude over time could have been using those patterns for some echoes, but not for others. They might also have been assessing variations in frequency and amplitude over the course of the entire echo train or for only portions of the train. To more closely evaluate the acoustic information relevant for object identification within echo trains, artificial neural networks can be trained to classify objects based on their echoes, and then examined how the networks accomplished this task (Wisniewski, Mercado, DeLong, & Heberle, 2013). The results of these simulations suggest that one auditory feature may not be sufficient to distinguish an object from all other objects, and that feedback (learning) is important for generalizing identification to novel aspects.

The most fruitful approach to elucidating the mechanisms underlying auditory object constancy is to combine human listening studies, artificial neural network simulations, and behavioral studies with animal subjects. Each type of experiment provides complementary insights into representational processes underlying object representation in organisms.

ACKNOWLEDGMENTS

This project was supported by a Research Seed Funding Grant to CMD from the Rochester Institute of Technology Office of the Vice President for Research, and a Faculty Development Grant to CMD from the Rochester Institute of Technology College of Liberal Arts.

REFERENCES

Au, W.W.L. **(1993)**. *The Sonar of Dolphins*. (Springer-Verlag, New York).

Au, W.W.L. and Martin, D.W. **(1989)**. “Insights into dolphin sonar discrimination capabilities from human listening experiments,” J. Acoust. Soc. Am. **86**, 1662-1670.

Au, W. W. L., and Turl, C. W. **(1991)**. “Material composition discrimination of cylinders at different aspect angles by an echolocating dolphin,” J. Acoust. Soc. Am. **89,** 2448-2451.

DeLong, C.M., Au, W.W.L., Harley, H.E., Roitblat, H.L., and Pytka, L. **(2007)**. “Human listeners provide insights into echo features used by dolphins to discriminate among objects,” J. Comp. Psych. **121**, 306-319.

DeLong, C.M., Au, W.W.L., Lemonds, D.W., Harley, H.E., and Roitblat, H.L. **(2006)**. “Acoustic features of objects matched by an echolocating bottlenose dolphin,” J. Acoust. Soc. Am. **119**, 1867-1879.

DeLong, C.M., Bragg, R., and Simmons, J.A. **(2008)**. “Evidence for spatial representation of object shape by echolocating bats (*Eptesicus fuscus),”* J. Acoust. Soc. Am. **123***,* 4582-4598.

DeLong, C.M., Heberle, A.L., Harley, H.E., and Au, W.W.L. **(2013)**. “View-invariant and view-dependent recognition of objects from multiple orientations by human listeners using dolphin echoes,” Manuscript in preparation.

DeLong, C.M., Kannyo, I., Benoit-Bird, K. J., and Au, W.W.L. **(2013)**. “Discrimination of fish prey by human listeners using dolphin and porpoise echolocation signals,” Manuscript in preparation.

Green, D.M. **(1993)**. “Auditory intensity discrimination,” in: *Human Psychophysics*, edited by W.A. Yost, A.N. Popper, and R.R. Fay (Springer, New York), pp. 13-55.

Helweg, D.A., Roitblat, H.L., Nachtigall, P.E., Au, W.W.L., and Irwin, R.J. **(1995)**. “Discrimination of echoes from aspect-dependent targets by a bottlenose dolphin and human listeners,” in: *Sensory Systems of Aquatic Mammals,* edited by R.A. Kastelein, J.A. Thomas, and P.E. Nachtigall (De Spil Publishers, Woerden, The Netherlands), pp. 129-136.

Herman, L.M., and Arbeit, W.R. **(1972)**. “Frequency discrimination limens in the bottlenose dolphin: 1-70Ks/c,” J. Aud. Res. **2**, 109-120.

Jolicoeur, P., and Humphrey, G.K. **(1998)**. “Perception of rotated two-dimensional and three-dimensional objects and visual shapes,” in: *Visual Constancies: Why Things Look As They Do,* edited by V. Walsh and J. Kulikowski (Cambridge University Press, London), pp. 69-123.

Kirkpatrick, K. **(2001)**. “Object recognition,” in: *Avian Visual Cognition* [Online] edited by R. G. Cook. Retrieved from www.pigeon.psy.tufts.edu/avc/kirkpatrick/

Nachtigall, P.E., Murchison, A.E., and Au, W.W.L. **(1980)**. “Cylinder and cube discrimination by an echolocating blindfolded bottlenose dolphin,” in: *Animal Sonar Systems,* edited by R.G. Busnel and J.F. Fish (Plenum Press, New York), pp. 945-947.

Thompson, R.K.R., and Herman, L.M. **(1975)**. “Underwater frequency discrimination in the bottlenose dolphin (1-140 kHz) and the human (1-8 kHz),” J. Acoust. Soc. Am. **57**, 943-948.

Wier, C., Jesteadt, W., and Green, D. **(1977)**. “Frequency discrimination as a function of frequency and sensation level,” J. Acoust. Soc. Am. **61**, 178-184.

Wisniewski, M.G., Mercado, E., III, DeLong, C.M., and Heberle, A.L., **(2013)**. “Exploring the capacity of neural networks to recognize objects from dolphin echoes across multiple orientations.” Paper presented at the 21st International Congress on Acoustics, Montreal, Quebec, Canada.