




Contents lists available at ScienceDirect

Behavioural Processes

journal homepage: www.elsevier.com/locate/behavproc



Border collie comprehends object names as verbal referents

John W. Pilley^{a,**}, Alliston K. Reid^{b,*}

^a 101 Seal St., Spartanburg, SC 29301, USA

^b Department of Psychology, Wofford College, 429N. Church St., Spartanburg, SC 29303, USA

ARTICLE INFO

Article history:

Received 6 August 2010
Received in revised form
16 September 2010
Accepted 30 November 2010

Keywords:

Referential understanding
Inferential reasoning by exclusion
Exclusion learning
Border collie
Dog
Receptive language

ABSTRACT

Four experiments investigated the ability of a border collie (Chaser) to acquire receptive language skills. Experiment 1 demonstrated that Chaser learned and retained, over a 3-year period of intensive training, the proper-noun names of 1022 objects. Experiment 2 presented random pair-wise combinations of three commands and three names, and demonstrated that she understood the separate meanings of proper-noun names and commands. Chaser understood that names refer to objects, independent of the behavior directed toward those objects. Experiment 3 demonstrated Chaser's ability to learn three common nouns – words that represent categories. Chaser demonstrated one-to-many (common noun) and many-to-one (multiple-name) name–object mappings. Experiment 4 demonstrated Chaser's ability to learn words by inferential reasoning by exclusion – inferring the name of an object based on its novelty among familiar objects that already had names. Together, these studies indicate that Chaser acquired referential understanding of nouns, an ability normally attributed to children, which included: (a) awareness that words may refer to objects, (b) awareness of verbal cues that map words upon the object referent, and (c) awareness that names may refer to unique objects or categories of objects, independent of the behaviors directed toward those objects.

© 2010 Elsevier B.V. All rights reserved.

1. Introduction

In an article in *Science* Kaminski et al. (2004) reported that a 9-year-old border collie (Rico) knew the names of more than 200 items. Their first experiment demonstrated that Rico's acquisition of the names of toys was indeed genuine – not a “Clever Hans” phenomenon, in which the successful retrieval of toys would be due to subtle cues other than the words. Their second experiment demonstrated Rico's “exclusion learning” – inferring the name of an object based on its novelty in the midst of familiar objects (Carey and Bartlett, 1978). As an example, Carey and Bartlett arranged a scenario during which 3 and 4-year-old children were shown two trays and asked to retrieve the “chromium, not the red tray.” Despite the children's lack of knowledge of the word chromium as a shade of green, the children correctly inferred that the teacher wanted the green tray. Carey and Bartlett dubbed this rapid linking of a proper-noun label upon an object as “fast mapping.”

Markman and Abelev (2004) found the report of Rico's apparent exclusion learning to be fascinating. Demonstrations of word

learning by exclusion in children have usually led researchers to conclude that the child has learned Baldwin's (1993) elements of referential understanding. Baldwin's work with children led her to conclude that if learning is limited to “associative factors” alone, learning would be slow. She identified two elementary factors of referential understanding that she believed to be necessary to expedite rapid word learning by children: (a) awareness that words may refer to objects, and (b) awareness of social cues that enable the mapping of the words upon the referent. Testing her hypotheses, Baldwin tried to eliminate associative factors by pitting temporal contiguity against her two elements of referential understanding. An experimenter presented two opaque containers to infants. Looking inside the first container, the experimenter exclaimed, “It's a *modi*.” Immediately thereafter, the experimenter withdrew the toy from the second container and gave it to the child for play. After a 10-s delay, the object in the first container was also given to the infant for play. Baldwin assumed that if associations based on temporal contiguity alone were critical for learning, the infants would identify the object from the second container as “*modi*.” However, despite the 10-s delay, the infants chose the object in the first container as “*modi*” – indicating that awareness of the reference cue influenced choice more than simple temporal contiguity. Baldwin concluded that fast word learning is mediated by referential understanding as opposed to associative mechanisms. Thus, the conclusion by Kaminski et al. (2004) that a border collie is able to learn words rapidly by exclusion invites intriguing

* Corresponding author. Tel.: +1 864 597 4642; fax: +1 864 597 4649.

** Corresponding author. Tel.: +1 864 576 7880.

E-mail addresses: pilleyjw@wofford.edu (J.W. Pilley), reidak@wofford.edu (A.K. Reid).

questions about the differences in word learning between dogs and children.

Markman and Abelev (2004) were unable to accept Rico's data as compelling evidence for exclusion learning because they identified two potential difficulties with the study: (a) lack of control for baseline novelty preference; and (b) reward after the exclusion choice response could have mediated the subsequent exclusion learning test trial. Thus, they questioned the validity of Rico's (Kaminski et al., 2004) demonstration of exclusion learning.

Bloom (2004) also considered the Rico data to be less than compelling. He acknowledged the possibility that Rico's learning of the names of objects may be qualitatively similar to that of a child, but may differ only in degree. However, he questioned the conclusion that Rico's words actually referred to objects. Did Rico treat the sound "sock" as a sock or did Rico treat the sound as a command to fetch a sock, and nothing more? If Rico treated the sound as a one-word proposition "fetch-the-sock," then his performance may have had little to do with language learning in the human sense. In addition, Bloom argued that words for children become symbols that refer to categories of things in the external world. "They appreciate that a word can refer to a category, and thereby can be used to request a sock, or point out a sock, or comment on the absence of one" (p. 1605).

We obtained our border collie, Chaser, soon after the publication of the Kaminski et al. (2004) study. The article led us to focus our research on the questions resulting from their intriguing research (Bloom, 2004; Fischer et al., 2004; Markman and Abelev, 2004). Experiment 1 investigated Chaser's ability to learn the names of over 1000 proper-noun objects over a 3-year period of intensive daily training. We wanted to know whether Rico's acquisition of over 200 words represented an upper limit for border collies, or whether an intensive training program with abundant rehearsal could teach a more extensive vocabulary. Experiment 2 tested Bloom's (2004) concern that words actually refer to objects, independent in meaning from the given command relative to the object. Experiment 3 explored the degree to which Chaser could learn several common nouns – names that represent categories of objects, in addition to the previously learned proper-noun names. Experiment 4 measured Chaser's ability to learn nouns by inferential reasoning by exclusion – inferring the name of an object based on its novelty in the midst of familiar objects, logically excluding the familiar alternatives. The four studies were designed to allow us to address concerns posed by Bloom (2004) and Markman and Abelev (2004) in their critiques of Kaminski et al. (2004) and to evaluate whether a dog can acquire referential understanding of nouns, as defined by Baldwin (1993), when all proper control conditions are included.

2. Experiment 1: investigating the ability of a border collie to learn proper nouns

Kaminski et al. (2004) reported that a 9-year-old border collie (Rico) knew the names of more than 200 objects. We wanted to know whether Rico's vocabulary represented an upper limit for border collies, or whether an intensive training program providing abundant rehearsal could teach a more extensive vocabulary. Experiment 1 provided 4–5 h of daily training over a 3-year period to teach Chaser the names of more than 1000 proper-noun objects.

2.1. Materials and method

2.1.1. Subject

The subject was a registered female border collie, Chaser, born in May 2004. We acquired Chaser when she was 8 weeks old, and she lived in our home primarily as a pet as well as a research subject. She exhibited the usual characteristics of her breed: intense visual

focus/concentration, instinct to find, chase, herd, attentive to auditory cues even during complex visual stimulation (such as herding sheep), responsive to soft levels of praise and verbal correction, and boundless energy.

2.1.2. Materials

Over a period of 3 years, we obtained 1038 objects for Chaser to identify and fetch. Objects were toys for children or dogs obtained mostly from second-hand discount stores, consisting of over 800 cloth animals, 116 balls, 26 "Frisbees," and over 100 plastic items. There were no duplicates. Objects differed in size, weight, texture, configuration, color, design, and material. Despite some similarities, each object contained unique features enabling discrimination. We gave 1022 of the objects a distinctive proper name consisting of 1–2 words (e.g., "elephant," "lion," "tennis," "Santa Claus"). [We duplicated the names for 16 objects, so these 16 objects were not used in this experiment, leaving 1022 objects with unique proper names]. We wrote this name on the object with a permanent marker to ensure that all trainers used the correct name consistently in all training sessions. Because Chaser handled each object with her mouth, we washed the objects when necessary to eliminate acquired odors and to maintain sanitary conditions. Fig. 1 shows a photograph of 42 of these objects with their associated names. Photographs of all 1038 objects with their associated names are available as online supplemental material.

2.1.3. Procedures

We initiated simple obedience and socialization training, 4–5 h daily, as soon as Chaser was brought into our home at 8 weeks of age. Behaviors and discriminations were taught by means of associative procedures, such as classical and operant conditioning including shaping procedures. Gradually, we began to provide training for herding, agility training, and tracking behaviors. These behaviors are not relevant to the focus of this paper, so we limit our discussions to the teaching and learning of nouns, with the exception of those commands that were essential to the teaching and testing of nouns. In Chaser's fifth month, we began to focus more of our time on word learning. Incentives and rewards used to encourage desired behaviors were petting, attention, and providing opportunities to engage in enjoyable activities (e.g., tugging, ball chasing, toy shaking, Frisbee play, agility play, walks, search by exploration, outdoor tracking, and stalking). For Chaser these types of incentives and rewards were more powerful than the traditional use of food. Furthermore, they were less distracting than food and more resistant to satiation. We used food to shape behavior only when food served as a lure, such as turning around in a circle. Throughout the paper, we use the word "novel" to refer to the objects for which Chaser had not learned a unique name. The word "familiar" denotes objects for which she had learned a unique name.

2.1.4. Specific training procedures for proper-noun objects

We taught Chaser one or two proper-noun names per day. Several trainers taught Chaser using the same procedures. All trainers were consistent in their use of the correct proper nouns because the name of each object was written on the object. Most of the training took place in our home and front and back yards. Each time we gave Chaser a new name to learn, we held and pointed to the object to be associated with the name and always said, "Chaser, this is Pop hide. Chaser find" No other objects were available on the floor for retrieval, so errors were unlikely. A 3–5 min play rehearsal period followed retrieval. During trials and play rehearsal periods, we repetitively verbalized the name of the object 20–40 times each session in order to facilitate the association of the name and object.

Following the initial training in the absence of other objects, the newly learned object was placed on the floor among other



Fig. 1. This photograph displays 42 of the 1022 objects used in Experiment 1, along with their corresponding names. Photographs and names of all objects are available as online supplementary material.

objects that had been recently learned – our working group. Over a period of 2–4 weeks, we gave Chaser daily rehearsal testing and play with these items, during which we repeatedly paired the name with the object, along with reinforcing play. As Chaser's vocabulary for proper-nouns grew, ultimately 50 or more newly learned objects were usually available in open Tupperware tubs for play and rehearsal testing, which could be initiated by Chaser or a trainer. As new names were learned, they were phased in for daily rehearsal and the older objects were phased out of the working group until monthly tests were given. When Chaser failed to retrieve an object upon command, we removed the other objects and gave Chaser additional training trials until she met our learning criterion (described below). Subsequently, the working group of objects was returned, and play with rehearsal testing continued. No object was removed from the working group unless Chaser fulfilled our learning criteria.

The procedure of teaching the names of humans, dogs, cats, locations, and stationary objects was similar to that used in teaching the names of objects, except that we told Chaser to “go to” the designated target both during learning and during testing. Rewards were praise and the opportunity to engage in enjoyable activities, such as “playing catch” with one of her toys.

We adjusted our daily training procedures, the duration of our sessions, and the amount of rehearsal to adapt to Chaser's ability to learn and retain new words. Each time Chaser made an error with any name–object pair, we provided additional training with that word until Chaser again completed the learning criterion for that name–object pair. Therefore, the majority of most training sessions involved rehearsal of previously learned words and many tests of retained knowledge. This procedure allowed us to measure Chaser's cumulative knowledge of proper nouns and to ensure that Chaser's vocabulary was increasing systematically, rather than new knowledge replacing previous knowledge. Because we adjusted our daily procedures to meet Chaser's retention and to keep her interested in the tasks, our training procedures did not allow us to assess the maximum rate of word learning or the maximum number of words that Chaser could learn. We were not concerned with measuring Chaser's innate abilities independent of training; rather, our primary goal was to discover what language accomplishments Chaser might achieve when given daily intensive training over years.

2.1.5. Design and statistical analysis

Language training, ongoing assessment, and formal blind tests of knowledge involve separate procedures that must maintain the subject's interest in participating. Herman et al. (1984) faced similar constraints as they trained and assessed dolphins' acquired language. They introduced terminology to help distinguish between the many informal tests of knowledge carried out in the highly social environment of training, and the rigorous formal tests of knowledge carried out under highly controlled conditions. They used the term “local” to describe the frequent tests of knowledge carried out during the training sessions, which normally lacked the rigorous experimental control necessary to rule out the influence of unintended cues. In contrast, the term “formal” described the formal assessments of acquired knowledge using the rigorous experimental control necessary to rule out the influence of unintended cues by the experimenter. We adopted these two terms to describe our various tests of word–object mapping.

2.1.6. Daily local tests during training

We first established a learning criterion that would provide clear evidence of name–object mapping of each name–object pair. This criterion was implemented as a testing procedure, used throughout training, that required Chaser to select the single correct object out of a varying collection of eight familiar objects (which we called a “local 1-of-8 test”, with probability of success = .125) without error eight times in a row (“local 8-of-8 binomial test”). The binomial probability of completing this difficult task due to random chance was $p = 5.96E - 8$ for each word, which we used as our learning criterion for each of the 1022 names. If Chaser made an error before selecting the object correctly eight times in a row, then the test ended, and we provided additional training until she successfully completed the local 8-of-8 test. During each local 1-of-8 test, Chaser was asked to select each of the remaining objects by name, without replacement, in order to rehearse prior learning. If Chaser made an error selecting any of the seven familiar distracter objects, then that object would be put aside for further training at a later time, and the local 8-of-8 test would continue. We often varied the seven distracter objects during these tests to provide increased rehearsal of prior learning.




Fig. 2. Filled circles depict the total number of proper-noun names learned over the months of training. Learning of each word was demonstrated by local 8-of-8 binomial tests (see text). Open circles display the minimum number of individual 1-of-8 word tests completed accurately over the months of training. Errors resulted in an increased number of tests.

Fig. 2 shows the minimum number of local 1-of-8 tests that Chaser had to complete as training progressed over the 3-year period, assuming she never made an error. Each noun-object pair was used in a minimum of eight 1-of-8 tests, yielding more than 8000 separate tests over the course of training. Similarly, Fig. 3 shows the minimum number of local 8-of-8 binomial tests that Chaser had to complete assuming no errors were made [note the logarithmic scale]. Of course, Chaser did make errors, so the actual number of local tests was much higher than these figures depict. Recall that our learning criterion (8 of 8 correct consecutive selections) was equivalent to a binomial test yielding $p = 5.96E - 8$. Thus, over the course of training, over 1000 different 8-of-8 tests were completed; or equivalently, Chaser rejected the null hypothesis over 1000 times in individual binomial statistical tests during informal training. Because the number of tests during training sessions was so high, we did not record the total number of additional 8-of-8 tests required when Chaser made errors during training, nor the individual errors made. Instead, we adjusted our daily training sessions to ensure the learning criterion was met for each noun-object pair to encourage over-learning. We separately

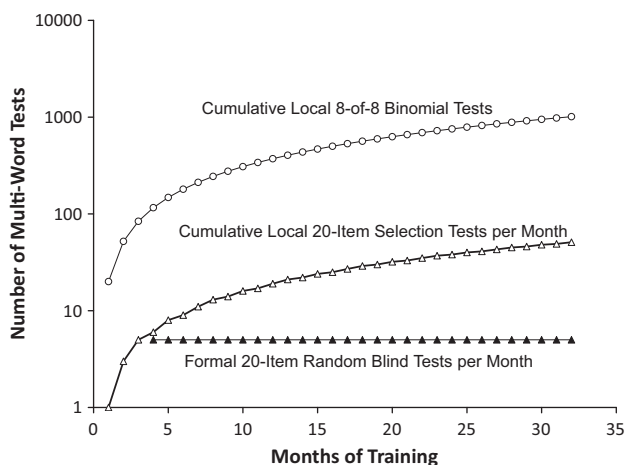


Fig. 3. This figure displays the minimum number of tests completed involving 8 or 20 names over the months of training. Note the logarithmic scale on the y-axis. The curves with open symbols depict local tests carried out in the presence of the trainer. The curve with filled triangles depicts formal tests carried out when the trainer was not present.

assessed noun-object mapping with formal tests outside of the training context, described below.

2.1.7. Monthly local tests during training

Once Chaser was 5 months old, we gave monthly local tests of Chaser's retention for all the proper nouns that she had previously learned. These tests also served to provide additional rehearsal of all noun-object mappings she had learned, distributed throughout her training. Each object was randomly placed upon the floor in sets of 20 items. We asked Chaser to retrieve each item in succession by its unique name. Objects were not replaced after retrieval, reducing test time and enabling immediate correction of error responses. As Chaser learned more and more names in each succeeding month, it was necessary later in training to space out the monthly rehearsal and testing over a 2-week period. Interspersed between the retrieval of toys, we asked Chaser to find or go to stationary objects, humans, cats, dogs, and locations. We inserted interim play between sets of 20 trials in order to maintain interest. Because these procedures produced selection without replacement, we replaced the binomial test with the hypergeometric test, designed for this purpose. As with all local tests, these 20-item local selection tests occurred in the presence of the trainer. Fig. 3 depicts the number of 20-item local selection tests carried out each month as training progressed. For example, once Chaser had learned 1000 nouns, 50 separate 20-item tests were required each month to assess and rehearse her learning.

2.1.8. Formal monthly 20-item random blind tests

All of the daily and monthly local tests described above were carried out during training in the presence of the trainer. Therefore, they may not have had the rigorous experimental control to ensure that Chaser selected objects exclusively on the basis of the verbalized name of the object. Even though multiple trainers provided training, it is possible that Chaser learned to attend to subtle visual or social cues that guided her behavior. Therefore, we carried out formal monthly tests of Chaser's long-term retention in a controlled situation in which the trainer and Chaser were out of visual contact in separate rooms. Chaser could not see the trainer, and the trainer could not see the location of the objects nor Chaser as she made her selection, until Chaser returned to the trainer with the selected object in her mouth and placed it into the empty Tupperware tub.

Every month, a random sample of 100 objects was selected from the total number of objects that Chaser had learned at that date (i.e. successfully met our learning criterion). These objects were randomly divided into five groups of 20 items, and five lists of the names of the 20 items were created, each in random order. Each 20-item group was dumped from a Tupperware container in random order on the floor of a distal bedroom. The 20 objects were then dispersed by moving the objects on the top to an outside perimeter until two concentric circles were formed, so that all objects were visible and did not overlap. The order of objects on the floor was always random and was not correlated with the order of the name on the list. Each trial began with Chaser standing beside the trainer in one room (the living room), and Chaser was asked to retrieve the objects on the list from the bedroom, one at a time without replacement. Chaser would enter the bedroom, select the object, and return to the living room to place the object in the Tupperware tub. Trials were considered correct if the name of the object (as written on the object) matched the name on the list. Trials were considered errors if the names did not match. Each of the 20 objects was retrieved in this fashion without replacement. Once this 20-item test was completed and the results recorded, the remaining 20-item tests were carried out using the same procedure. The time between these tests varied considerably each month, from minutes to hours, in order to maintain Chaser's interest, eat meals, or engage in other activities.




Fig. 4. This figure shows how the theoretical probability of making correct choices due to random chance, when selection without replacement is used, depends upon the total number of objects learned at that point.

Because these 20-item tests were carried out without replacement, hypergeometric tests (not binomial tests) were used to calculate the probability of selecting objects correctly from the available options. With hypergeometric distributions the probability of making, say, 18 of 20 correct choices due to random chance depends upon the total number of objects learned at that point (equivalently, the number of objects from which the 20 were selected). Fig. 4 shows how this probability changed as the number of learned objects changed. Over the 3-year duration of this experiment, these formal monthly 20-item random blind tests were completed 145 times, each yielding independent randomized measures of Chaser's retention of noun/object mapping.

2.1.9. Formal public double-blind demonstration

At the end of 3 years of training and several months after the formal tests above had terminated, we set up a public demonstration of Chaser's acquired vocabulary in a college auditorium. Although it was not as comprehensive as the formal tests described above, it was designed as a formal controlled, double-blind test in front of about 100 psychology students. Out of sight of the trainer, five students randomly selected 10 objects (50 total) from the mass of 1022 objects piled on the auditorium floor. The students wrote the names of each object on paper and provided each list to the trainer. The 10 objects in each group were placed upon the floor in random order by the students behind and out of sight of the trainer. Chaser was then requested to retrieve each of the 50 objects in the order provided by the students' lists, in 5 consecutive sets of 10 objects. The audience determined the accuracy of Chaser's selection, and the entire demonstration was videotaped.

2.2. Results and discussion

2.2.1. Number of words learned

Over the course of training, Chaser learned all 1022 name/object pairs, correctly completing our local 8-of-8 binomial tests ($p = 5.96E - 8$) for every word. Recall that when errors were made, Chaser received additional training until she again completed our learning criterion. This over-learning procedure required Chaser to repeat the local 8-of-8 binomial tests several times for some names over the course of training, and it helped ensure that names were not forgotten as new names were learned. Fig. 2 shows the number of words learned over the months of training. It is important to recognize that the linear learning rate depicted was produced by our procedure of training only one or two words per day, and it does not reflect innate cognitive abilities or limitations that Chaser

might have. Chaser may have been able to learn at a faster rate, but our procedure did not provide the opportunity. We observed no decrease in speed of learning over the course of training, even as she learned over 1000 words. Therefore, we are not able to extrapolate how long the linear learning rate would continue, nor the maximum number of words she would be able to learn. We stopped training after 3 years, not because Chaser had reached some cognitive limit, but because we could no longer invest the 4–5 h per day training her.

2.2.2. Local monthly 20-item test results

For 32 consecutive months, Chaser was tested for her retention of all the names of objects that she had previously learned. As her vocabulary grew, the number of 20-item tests given per month increased from 1 in the first month to 51 in the 32nd month, resulting in a total of 838 independent tests. Recall that each object was randomly placed upon the floor in sets of 20 items. We asked Chaser to retrieve each item in succession by its unique name, and objects were not replaced after retrieval. Because these 20-item tests were carried out without replacement, hypergeometric tests (not binomial tests) were used to calculate the probability of selecting objects correctly from the available options. With hypergeometric distributions the probability of making, say, 18 of 20 correct choices due to random chance depends upon the total number of objects learned at that point (equivalently, the number of objects from which the 20 were selected). Fig. 4 shows how this probability changed as the number of learned objects changed. During all 838 independent tests, Chaser successfully recalled the names of 18, 19, or 20 objects out of the many sets of 20 items – that is, in no test did Chaser fail to recall the name of at least 18 items correctly in each set of 20 items. The probability of selecting 18 of 20 correctly due to random chance asymptotes at a maximum value of $p = .00016$; 19 of 20 at $p = .000016$; and 20 of 20 at $p = .0000008$. Therefore, Chaser's exceptional accuracy demonstrates that her overall vocabulary size increased to 1022 nouns, rather than new words replacing previously acquired vocabulary.

2.2.3. Formal monthly 20-item random blind test results

Because the local 20-item tests described above were carried out in the presence of the trainer, the formal 20-item random blind tests were designed to ensure that the trainer could not have inadvertently provided visual or social cues that influenced Chaser's selections. Thus, these formal tests controlled for the presence of, and ability to see, the trainer. If accuracy was lower in these tests than when the trainer was present, then we would be suspicious of a Clever Hans effect.

Recall that these monthly tests randomly selected 100 objects from the entire set of learned objects, and divided them into five 20-item random blind tests. Thus, 145 independent tests were completed over the course of training. During all 145 independent tests, Chaser successfully recalled the names of 18, 19, or 20 objects out of the many sets of 20 items – that is, in no test did Chaser fail to recall the name of at least 18 items correctly in each set of 20 items. Therefore, these tests duplicated the results of the local 20-item tests. The presence of the trainer did not increase Chaser's exceptional accuracy, so there was no Clever Hans effect.

2.2.4. Formal public double-blind demonstration results

Recall that the double-blind public demonstration tested Chaser's ability to correctly retrieve 50 objects randomly selected from all 1022 objects she had learned. Thus, this test was not as comprehensive as the tests described above. Nevertheless, Chaser successfully retrieved from the five sets of 10 objects, the following numbers of objects: 10, 9, 9, 8, 10, resulting in a total of 46 out of 50 or 92% of the total items (hypergeometric test, $p < .0000001$). This impressive accuracy was demonstrated in an auditorium with

about 100 noisy, distracting students, in a double-blind procedure carried out months after training was terminated.

In summary, in no case was Chaser unable to distinguish between verbal labels given as names for her objects. Her learning and retention of more than 1000 proper nouns revealed clear evidence of several capacities necessary for learning receptive human language: the ability to discriminate many nouns phonetically, the ability to discriminate many objects visually, a sizable vocabulary, and a sufficient memory system.

3. Experiment 2: testing independence of meaning of nouns and commands

Recall that in the Rico study (Kaminski et al., 2004), Bloom (2004) was not convinced that words for Rico conveyed *reference* – that labels actually *referred* to objects, independent from the meaning of an associated command. For example, when Rico was told to “fetch sock,” did Rico comprehend that the label “sock” *referred* to a specific object *and* separately comprehend that the word “fetch” meant that he should produce a specific behavior involving that specific object? If Rico actually treated the label “sock” as a command to “fetch sock” only, then it would not be evidence that he understood *reference*. That is, Rico may not have understood that the label “sock” referred to a specific object, independent of a behavior directed toward the sock. Thus, Bloom argued that Rico “might not understand *reference* at all and might be limited to associating the word spoken by the owner with a specific behavior” (p. 1604). If so, then Rico’s word learning may have little to do with language learning as exhibited by humans.

In essence, Bloom’s concern addresses the question as to whether Rico understood that the phrase “fetch sock” represented two independent morphemes – that objects are independent in meaning from the activity requested that involved that object. Thus, the primary purpose of Experiment 2 was to test whether Chaser treats the name of an object independent in meaning from the given command. A secondary purpose of the study was to investigate Chaser’s ability to combine nouns and commands and demonstrate understanding of the combination of the two words, even though the command and noun had never been used in combination prior to the testing condition. In 14 independent trials, we randomly paired one of three commands with one of three different names of her objects to see whether Chaser would demonstrate independence of meaning between the nouns and commands.

3.1. Method

3.1.1. Subject and materials

The subject used in this study, Chaser, was the same as that used in Experiment 1. We arbitrarily selected three of Chaser’s objects that had been used in Experiment 1, so Chaser had already learned their proper-noun names. All three objects were similarly sized cloth toys. The object named “Lips” resembled human lips. The object named “ABC” was a cloth cube with the letters A, B, and C written on its sides. “Lamb” was a stuffed toy resembling a lamb. Each object had its assigned name written on its side in permanent ink.

3.1.2. Procedure

In order to demonstrate that command–noun combinations were indeed two independent units of meaning, Chaser was asked to produce appropriate behaviors when three different commands (take, paw, and nose) were randomly combined with three different objects (Lips, ABC, Lamb) across 14 independent trials in a double-blind procedure. None of the three objects had previously been paired with any of the three commands. To guarantee independent and random selection of each object and each command, the

Table 1

Sequence of trials in Experiment 2 testing independent meanings of three names and three commands presented in random pairs.

	Instruction given:		Trial	Cumulative probability of success	
	Command	Name			
Trial 1	Take	Lamb	Correct	(1/9) ¹	0.11111
Trial 2	Paw	Lamb	Correct	(1/9) ²	0.01235
Trial 3	Take	Lips	Correct	(1/9) ³	0.00137
Trial 4	Nose	Lamb	Correct	(1/9) ⁴	0.00015
Trial 5	Paw	Lips	Correct	(1/9) ⁵	0.000017
Trial 6	Paw	Lamb	Correct	(1/9) ⁶	0.000002
Trial 7	Nose	Lips	Correct	(1/9) ⁷	2.09E–07
Trial 8	Nose	ABC	Correct	(1/9) ⁸	2.32E–08
Trial 9	Take	ABC	Correct	(1/9) ⁹	2.58E–09
Trial 10	Nose	Lamb	Correct	(1/9) ¹⁰	2.87E–10
Trial 11	Paw	ABC	Correct	(1/9) ¹¹	3.19E–11
Trial 12	Take	Lamb	Correct	(1/9) ¹²	3.54E–12
Trial 13	Paw	Lamb	Correct	(1/9) ¹³	3.93E–13
Trial 14	Take	ABC	Correct	(1/9) ¹⁴	4.37E–14

three objects and three commands were each assigned different numbers, and a table of random numbers was used to select each object and each command for each of the 14 trials. Table 1 shows the details of each trial.

Testing was carried out in a large carpeted room at Wofford College. The three cloth objects were placed upon a soft pad (1.5-m long) positioned about 38 cm apart in a row, 1.5-m in front of and parallel to a vertical cloth barrier. The cloth barrier (1-m high and 1.4-m wide) separated the experimenter from the objects, and the experimenter knelt behind the barrier before giving Chaser instructions and until the trial was over. Neither Chaser nor the experimenter could see the other during execution of the choice task, and the experimenter could not see which object was selected nor the actual command carried out. A confederate, seated 4.3 m from the objects, was able to observe Chaser’s behavior and signal to the experimenter (by silently waving her hand) that the trial was over. Neither the experimenter nor the confederate measured the accuracy of Chaser’s behavior or the object selected. The confederate was used only to indicate when the trial was over, since the experimenter could not see Chaser during the trial. When each trial ended, the experimenter called out “Good dog!” (independent of the accuracy of Chaser’s behavior), thus terminating the trial for Chaser and beginning a brief play session with a rubber ball. Play between trials was essential in order to maintain interest in the assigned tasks. Following each trial, objects were returned to their original ordinal position. Fig. 5 shows the testing scenario, including the room, the cloth barrier, the experimenter, Chaser, placement of the three objects, and the location of the confederate.



Fig. 5. This photograph shows the testing scenario used in Experiment 2, including the room, the cloth barrier, the experimenter, Chaser, placement of the three objects, and the location of the confederate. A video of this procedure is available as online supplementary material.

The entire testing session was videotaped with sound recording (available as online [supplemental material](#)). Three independent raters were recruited to measure independently the accuracy of each command and the accuracy of each object selected. We read them a definition of the actions required for each command, but explained that their judgment was final. “Take” was defined as picking up the object with the mouth. “Paw” was defined as an actual touching of an object with the front paw or a close sweep of the paw over the object. “Nose” was defined as actually touching the object with the nose without using the mouth or a nose-oriented movement very close to the object. [Note: This lenience in defining “paw” and “nose” responses resulted from Chaser’s previous experience with abrasive objects, such as trees and pinecones, which encouraged Chaser to produce paw and nose responses slightly short of actual touch.] Each rater first watched the videotape with the sound turned off (so the rater could not know which instructions were given to Chaser), and recorded which command was actually executed towards which object. Once all 14 trials were rated, the rater then watched each trial again with the sound turned on in order to assess whether Chaser’s behavior accurately matched the command and object instructed by the experimenter. Inter-rater reliability was assessed.

3.2. Results and Discussions

Inter-rater reliability for the three raters across the 14 trials for both commands and objects was 100%. All raters judged Chaser to be 100% accurate across the 14 trials, performing the correct command to the correct object as instructed. Table 1 displays the orders of the pairings of the three commands and three nouns, Chaser’s accuracy in each trial, and the obtained cumulative probability of correct responses over the 14 trials in the testing session. With the combination of three commands and three objects, the probability of success on a given trial was $1/9 = .111$, and the probability of failure was $8/9 = .889$. As indicated in the table, judgments by the three raters clearly showed that Chaser successfully demonstrated appropriate responses to the correct objects on all 14 trials ($P = (1/9)^{14} = 4.37E - 14$).

3.2.1. Independent meanings of commands and nouns

Thus, Chaser produced a unique response oriented to each object depending upon the meaning of the associated command. She responded as though the commands and the proper-noun names were independent entities or morphemes. Thus, in effect, Chaser treated phrases like “fetch sock” as though the “sock” was a sock and not “fetch sock” – indicting that her nouns referred to objects. Thus, Bloom’s (2004) concern that Rico may not have understood the difference between “sock” and “fetch the sock” is ruled out in this study. These results clearly support the conclusion that Chaser understood *reference* – that the verbal noun of an object referred to a particular object with distinct physical features independent of actions directed toward that object.

3.2.2. Combinatorial understanding

Not only did Chaser demonstrate that she ascribed independent meanings to commands and names, she also demonstrated that she could combine the two meanings accurately without explicit training to do so. None of the three objects had been paired previously with any of the three commands prior to this experiment, and she never received differential reinforcement for correct pairings (even during testing). Chaser had received extensive experience with these three commands directed toward other familiar objects, but she was able to combine these novel command–noun phrases accurately on the first trial without additional training (perhaps due to response generalization). Thus, after learning the meanings of the commands and the nouns, Chaser was able, as

do children, to understand the novel combination of two-word phrases.

4. Experiment 3: learning common nouns

Bloom (2004) noted that Kaminski et al. (2004) did not provide evidence that words for Rico represented categories, as nouns often do for children: “For a child, words are symbols that refer to categories and individuals in the external world” (p. 1606). The important distinction here is the difference in reference between a proper noun and a common noun. A proper noun is a symbol that refers to a particular object, as demonstrated in Experiment 1. A common noun, however, is a symbol that refers to a category of objects. Categorization in animals has been widely studied in animal cognition in the last few decades, and a few studies have demonstrated that dogs can form categories. For example, Range et al. (2007) demonstrated that dogs are able to categorize visual stimuli, treating pictures of dogs differently from pictures of landscapes, and the dogs successfully extended these categories to novel stimuli. Similarly, Heffner (1975) demonstrated that dogs could form categories of auditory stimuli, differentiating between “dog” versus “non-dog” sounds, and extending these categories to novel stimuli. To our knowledge, no one has demonstrated that dogs can form categories identified by the names of objects.

This study assessed Chaser’s ability to learn three common nouns (“toy,” “ball,” “Frisbee”) that represented different categories of objects. Chaser had already learned the proper-noun name of all of the objects in each category (see Experiment 1). Because objects in the categories “ball” and “Frisbee” had similar shapes (generally round or disk-shaped, respectively), we predicted that these physical features would form the basis for categorization. Objects in the “toy” category, however, differed widely in their physical features, but they were similar in function: Chaser had been allowed to play with all of the toys, but she was not allowed to play with the many available non-toys in the house that had similar appearance.

4.1. Materials and method

The subject, Chaser, and materials used in this study were the same as those used in Experiment 1.

4.1.1. Categories

Chaser had already learned the proper-noun names of all objects used in this study. We selected three category labels that the experimenters commonly used (when talking to each other) to represent many of Chaser’s objects (“toy,” “ball,” “Frisbee”). Chaser had many toys, many balls, and many “Frisbees.”

4.1.1.1. “Toy”. “Toy” represented a category for all 1022 objects for which Chaser had learned a unique name and was allowed to play. “Non-toys” were physically similar objects that belonged to family members with which play was forbidden or they were out of reach, such as certain cloth animals, balls, dolls, shoes, socks, and a variety of small personal and household objects. These were located throughout the house on shelves, desks, tables, and closets. Thus, the physical characteristics of toys and non-toys were not discriminative features of the category, but the categories did differ reliably in function: Chaser learned that she could play with toys; she could not play with non-toys. We selected eight non-toys (see Fig. 6) that had been located within Chaser’s sight but out of reach. She had never been expressly forbidden to touch these particular non-toys, but she had never been given the opportunity to treat them as her toys.



Fig. 6. This photograph displays the 16 objects (eight toys and eight non-toys) used to test Chaser's comprehension of the common noun "toy" in Experiment 3.

4.1.1.2. "Ball". "Ball" represented a category of 116 round, bouncy objects, which we identified from all of the objects Chaser had learned to identify. Thus, "ball" was a subset of Chaser's "toys," and Chaser had previously learned proper names for all balls and toys in Experiment 1.

4.1.1.3. "Frisbee". "Frisbee" represented a category of 26 objects that had disk-like qualities, which we identified from all of the objects Chaser had learned to identify. Thus, "Frisbee" was a subset of Chaser's "toys," and Chaser had previously learned proper names for all Frisbees and toys in Experiment 1.

Thus, for each of the three categories, the name of each category was a second or third label for Chaser's objects: (1) they were identified by the unique proper-noun label; (2) they were all categorized by the common noun "toy"; and, (3) "Balls" and "Frisbees" were further categorized with these common-noun labels, while still being identified as "toy" and by the unique proper-noun name.

4.2. Procedures

4.2.1. Generalization and discrimination training

We taught the common-noun name for an object representing a particular category using the same methods as those followed in Experiment 1 in teaching proper nouns. For generalization training of a common noun, we placed upon the floor eight exemplars of the particular common noun. We asked Chaser to retrieve each of the common-noun items successively in response to the given name of the category ("fetch another toy"). We replicated this training procedure at least three times, using different exemplars of the common noun in each of the three replications. Since no other items were available on the floor, errors in selection were unlikely.

Following generalization training, we initiated discrimination training during which eight non-exemplars of the category were randomly placed on the floor among eight exemplars that had not been used during generalization training (such as eight balls and eight non-balls). Arrangement of the 16 objects formed an inner circle of 8 objects and an outer circle of 8 objects. In both circles, exemplars of the, say, ball and non-ball categories were randomly placed in alternating positions in order to inhibit choice of items based on the object's position. We asked Chaser to retrieve each of the exemplars representing the category (e.g., "fetch a ball" or "fetch another ball"). If she retrieved correctly, she was reinforced; if she chose a non-exemplar, she was gently told, "No, that is not ___ (name of common noun)." This procedure of correcting her wrong choices immediately with a soft "no" was effective in that Chaser's

selection of non-exemplars decreased. Discrimination training continued until Chaser successfully retrieved eight exemplars of the common nouns upon command without retrieving any of the eight non-exemplars. Our criteria for the learning for each of the three categories required eight consecutive demonstrations of both generalization and discrimination without error.

4.2.2. Testing comprehension of common nouns

During testing, all objects were placed in an adjoining room, outside of the vision of the tester, to ensure that the tester could not provide inadvertent cues influencing the selection of the objects. As in training, arrangement of the eight exemplar and eight non-exemplars of each category formed an inner circle of eight objects and outer circle of eight objects. In each circle, exemplars and non-exemplars of the category were randomly placed in alternating positions to prevent selection of an object based on its position.

Each trial began with Chaser standing beside the trainer in one room (the living room). Chaser was asked eight times to retrieve an exemplar from the bedroom, one at a time without replacement (e.g., "fetch a ball" or "fetch another ball"). Chaser would enter the bedroom, select the object, and return to the living room to place the object in the Tupperware tub. Trials were considered correct if the object was a member of the set of objects described by the common noun (an exemplar). Trials were considered errors, if a non-exemplar was selected. Each of the eight exemplars was retrieved in this fashion without replacement, leaving all eight non-exemplars in the bedroom. Videotaped demonstrations of Chaser completing this procedure for the three categories are available online as [supplemental material](#).

Because the number of exemplars available in the toy and ball categories was high, we were able to test Chaser's comprehension of those two common nouns without using the exemplars and non-exemplars that had been used during training. With only 26 Frisbees, however, we had to use some of the same Frisbees during both training and testing.

4.3. Results and discussion

4.3.1. Toy

The photograph in Fig. 6 displays the 16 objects (eight toys and eight non-toys) used to test Chaser's comprehension of the common noun "toy." When Chaser was asked to retrieve a toy eight times in succession out of the 16 objects available, she correctly selected a toy without error in every trial (i.e., without selecting a non-toy). Hypergeometric tests were used to calculate the probability of selecting objects correctly without replacement from the available options. The probability of making eight correct choices out of the 16 options due to random chance was $P(\text{toy}=8)=7.77\text{E}-5$. Thus, Chaser comprehended the common noun "toy" as a label for the toy category, even though Chaser also knew each object by its proper-noun name.

4.3.2. Ball

The photograph in Fig. 7 displays the 16 objects used to test Chaser's comprehension of the common noun "ball." When Chaser was asked to retrieve a ball eight times in succession out of the 16 objects available, she correctly retrieved a ball without error in every trial. The hypergeometric probability of making eight correct choices out of the 16 options due to random chance was $P(\text{ball}=8)=7.77\text{E}-5$. Thus, Chaser comprehended the common noun "ball" as a label for the ball category, even though Chaser also knew each object by its proper-noun name and by the common noun "toy."



Fig. 7. This photograph displays the 16 objects (eight balls and eight non-balls) used to test Chaser's comprehension of the common noun "ball" in Experiment 3.

4.3.3. Frisbee

The photograph in Fig. 8 displays the 16 objects used to test Chaser's comprehension of the common noun "Frisbee." When Chaser was asked to retrieve a Frisbee eight times in succession out of the 16 objects available, she correctly retrieved a Frisbee without error in every trial. The hypergeometric probability of making eight correct choices out of the 16 options due to random chance was $P(\text{Frisbee} = 8) = 7.77E - 5$. Thus, Chaser comprehended the common noun "Frisbee" as a label for the Frisbee category, even though Chaser also knew each object by its proper-noun name and by the common noun "toy."

4.3.4. Discrimination of categories

Thus, Chaser learned that names of objects may represent categories with many exemplars – for each of the three common nouns, she mapped one label onto many objects. Membership in two of the categories, "ball" and "Frisbee", could be discriminated based on common physical properties. For example, all balls had similar round shapes, and all Frisbees were shaped like a disk. The toy category required a more abstract discrimination. The "toy" and "non-toy" objects did not differ in terms of common physical characteristics. Thus, discrimination of whether an object is a toy or not was unlikely to be based upon identifying physical "toy-like" fea-



Fig. 8. This photograph displays the 16 objects (eight Frisbees and eight non-Frisbees) used to test Chaser's comprehension of the common noun "Frisbee" in Experiment 3.

tures. However, Chaser learned that she could play with toys, but she could not play with non-toys – they differed in functionality. While this study was not designed to assess how Chaser created categories, it seems likely that the common noun "toy" reflected a more abstract type of categorization than did "ball" or "Frisbee."

4.3.5. Noun–object mapping

By forming categories represented by common nouns, Chaser mapped one label onto many objects. Chaser also demonstrated that she could map up to three labels onto the same object without error. For example, Experiment 1 demonstrated that Chaser knew the proper-noun names of all objects used in this study. Chaser also mapped the common noun "toy" onto these same objects. Her additional success with the two common nouns "ball" and "Frisbee" demonstrates that she mapped a third label onto these objects.

Bloom (2004) had commented that Kaminski et al. (2004) did not provide evidence that words for Rico represented categories. Chaser, however, has demonstrated that words may represent categories. Her demonstrations of one-to-many and many-to-one noun/object mappings extend our understanding of the referential nature of words in border collies.

5. Experiment 4: learning words by exclusion

Markman and Abelev (2004) asserted that the most striking finding of Rico's word learning (Kaminski et al., 2004) was Rico's apparent demonstration of learning new words by exclusion – inferring the referent of a novel word by logically excluding other potential alternatives. Successfully learning new names by exclusion consists of two or more separate abilities. The procedure involves a choice trial in which the subject is provided a novel name and expected to choose a novel object located among a group of familiar objects that already have associated names. This choice condition requires inferential reasoning by exclusion, defined as "the selection of the correct alternative by logically excluding other potential alternatives" (Aust et al., 2008, p. 587). This choice response cannot be based on associative learning mechanisms because the name and object referent are not presented together in this single trial. Aust et al. (2008) demonstrated inferential reasoning by exclusion in about half of their dogs in a non-verbal visual discrimination task. Similarly, Erdőhegyi et al. (2007) demonstrated that dogs can find hidden objects using inferential reasoning by exclusion. Following Kaminski et al. (2004), the current experiment explores whether the names of novel objects can be learned by exclusion.

Retention of the novel name–object mapping may be a separate ability beyond that of the original inferential reasoning that led to the selection of the object in that first choice trial. Retention may depend upon additional factors, including associative learning such as the reinforcing consequences of selecting that novel object. Thus, it is possible that a dog could correctly select a novel object in response to a novel name, but be less able to retain this name–object pairing over minutes, days, or weeks. Nevertheless, Kaminski et al. (2004) claimed that Rico was able to select the novel object and retain this name–object knowledge over time, comparable to the performance of 3-year-old toddlers.

Markman and Abelev were unwilling to accept Rico's data as compelling evidence of word learning by exclusion because they perceived potential flaws in the design: (a) a lack of control for baseline novelty preference, and (b) the possibility that positive feedback after the exclusion trial could have confounded the subsequent retention test. Fischer et al. (2004) replied that Rico was asked to retrieve two familiar objects before being asked to retrieve the novel object, revealing that correct retrieval was not due to preference for novel objects alone. Rico showed clear impulse con-

trol by retrieving the requested novel object only after retrieving two requested familiar objects. We replicated this procedure with Chaser, but we also measured Chaser's baseline probabilities for the choice of objects based on novelty alone, prior to the first opportunity to learn words by exclusion. We also measured Chaser's retention of word–object mapping after several delay intervals.

5.1. Materials and method

The subject, Chaser, used in this study was the same as in Experiment 1. The materials consisted of both novel and familiar objects. Each of the familiar objects had been used in Experiment 1. The novel objects were 64 additional children's toys and stuffed animals that had not been used in the previous experiments.

5.2. Procedures

5.2.1. Measuring baseline novelty preference

In order to control for choice of an object based on novelty alone, we placed two novel and eight familiar objects in a separate room. Chaser was asked successively to retrieve each of the eight familiar objects by its unique name, one at a time, without replacement. If Chaser had a novelty preference over familiar objects, then she should retrieve novel objects. We repeated this procedure seven more times using different novel and familiar objects in each trial. Novel objects were not given names during this procedure.

5.2.2. Testing inferential reasoning by exclusion

Testing procedures for inferential reasoning by exclusion were similar to those used with Rico (Kaminski et al., 2004). We placed a single novel object among seven familiar objects in an adjacent room. On the first two trials, we asked Chaser to retrieve familiar toys by name. On the third trial, we asked her to retrieve the novel item by its novel name. Upon successful retrieval of the novel item, Chaser was told "Good dog."

5.2.3. Testing immediate and delayed name retention

We immediately followed this novel choice trial with a retention test. This novel item was placed in an adjoining room among three novel and four familiar objects. Objects were arranged in a circle, randomly alternating the placement of familiar and novel objects. On the first two trials, we asked Chaser to retrieve familiar toys by name. On the third trial, we asked her to retrieve the newly learned novel item by its new name. We repeated this retention test after an additional 10-min delay, and again 24-h later. We completed eight replications of this four-part procedure over a period of 8 days, using different novel and familiar items in each replication.

5.2.4. Retesting at 5 years

Two years later, when Chaser was 5-year old, we retested Chaser's ability for inferential reasoning by exclusion and name retention by repeating the procedures above with a new group of novel items.

5.3. Results and discussion

5.3.1. Baseline rates for choice based on novelty alone

When we measured the baseline rates for choice of objects based upon novelty alone, Chaser retrieved each of the eight familiar items in succession without error, ignoring the two novel items, in each of the eight separate tests. Thus, her baseline rate for choosing objects based on novelty alone was zero.

5.3.2. Inferential reasoning by exclusion

Early in Chaser's training, we corrected her immediately if she retrieved an incorrect object. As the result of this training, if she

did not remember the name of an object, she would simply stand among all the objects without selecting one. Thus, on the first exclusion learning trial when she did not return with an object, a quick look revealed that she was simply standing among the objects. Only with repeated encouragements to fetch the newly named novel object, did she finally retrieve the novel object. Special encouragement was also necessary on the second replication. On the remaining six replications, she continued to retrieve novel objects successfully without further encouragement or error. Thus, Chaser was successful in eight successive replications, conducted over 8 days, in retrieving the novel objects (binomial test, $P(X=8) = 5.96E - 8$).

Recall that Markman and Abelev (2004) were unwilling to accept Rico's data as compelling evidence of word learning by exclusion because they were concerned about inadequate control for baseline novelty preference. Chaser's baseline rate of selecting novel items was zero. In addition, as with Rico, Chaser first retrieved two familiar objects upon command before the request to retrieve the novel object, thereby demonstrating impulse control for retrieving novel objects. Therefore, her demonstration of inferential reasoning by exclusion was not due to preference for selecting novel items.

5.3.3. Immediate and delayed name retention

After each of the eight tests of inferential reasoning by exclusion, we tested Chaser's retention immediately, after a 10-min delay, and after 24-h delay. Immediately after correctly mapping the novel name to the novel item, the item was randomly placed among three other novel objects and four familiar objects. Chaser was asked to retrieve two familiar items before the request to retrieve the newly named object. In each of the eight replications, Chaser was successful in retrieving the newly named object when tested immediately after the exclusion choice response (binomial test with $\text{PROB}(\text{success}) = 1/6$, $P(X=8) = 5.95E - 7$). When tested after 10-min delay, she was successful in retrieving five of the eight novel items (binomial test, $P(X=5) = .0042$). With 24-h delays, Chaser displayed little retention of the names of the novel items, successfully retrieving only one of the eight items (binomial test, $P(X=1) = .37$).

5.3.4. Re-testing at 5-year old

Approximately 2 years later when Chaser's was 5-year old, we retested Chaser's ability for inferential reasoning by exclusion and name retention by repeating the procedures above with a new group of novel items. Again, in each of the eight replications, Chaser correctly chose the novel objects without error (binomial test, $P(X=8) = 5.95E - 7$). When tested for memory retention immediately after the exclusion choice trial, Chaser selected the newly named object in six of the eight tests (binomial test, $P(X=6) = .0004$). When tested at the 10-min delay interval, she selected the newly named object in four of the eight tests (binomial test, $P(X=4) = .026$).

A central issue is whether inferential reasoning by exclusion can lead to word learning in border collies. Chaser demonstrated accurate name–object mapping when tested immediately following the exclusion choice trial. However, although statistically significant, retention was reduced when tested at the 10-min delay and essentially non-existent when tested at the 24-h delay. Informal observations indicated that additional pairings with the name (playing with the object) were necessary for retention intervals of an hour or more. Markman and Abelev (2004) were concerned that positive feedback given to Rico after the exclusion trial could have affected the follow-up retention test 4 weeks later. Chaser also received minimal positive feedback following correct choices, so we cannot know what role associative learning mechanisms contributed to Chaser's successful, but transient, word learning. It is likely that rehearsal play with the objects would provide

additional opportunities for associative learning to produce longer retention, but this retention would no longer reflect properties of inferential reasoning by exclusion. Working with toddlers, Horst and Samuelson (2005, 2006) and Horst et al. (2006) reported that the exclusion choice trial was not sufficient to produce long-term retention without additional rehearsal. Kaminski et al. (2004) reported Rico's successful word learning by exclusion with retention intervals up to 4 weeks – intervals comparable to those sometimes observed in toddlers. However, Chaser did not show such retention unless we provided additional rehearsal with the object during play periods (Carey and Bartlett, 1978; Golinkoff et al., 1992; Horst et al., 2006; Horst and Samuelson, 2005, 2006; Markman and Wachtel, 1988).

6. General discussion

6.1. Experiment 1: investigating the ability of a border collie to learn proper nouns

Experiment 1 showed that over a 3-year period, Chaser demonstrated impressive abilities to learn the names of proper nouns, and her extensive vocabulary was tested under carefully controlled conditions. She met our rigorous learning criteria repeatedly for 1022 proper nouns. She accumulated and maintained this knowledge of nouns over a lengthy 32-month period. In each month, when we tested Chaser on her entire vocabulary, she was able to respond correctly to more than 95% of the 1022 proper-noun objects that she had learned. She revealed no difficulty in discriminating between the many different sounds of the nouns given to her as names for objects. Likewise, she revealed no problems in distinguishing objects visually, despite their similarity and varying positions (often partially occluded) on the floor. Her ability to learn and remember more than 1000 proper nouns, each mapped to a unique object, revealed clear evidence of several capacities necessary for learning receptive human language: the ability to discriminate between 1022 different sounds representing names of objects, the ability to discriminate many objects visually, a sizable vocabulary, and an extensive memory system that allowed the mapping of many auditory stimuli to many visual stimuli.

6.2. Experiment 2: testing independence of meaning of nouns and commands

In Experiment 2 Chaser demonstrated she understood the unique combinations of three commands (take, paw, nose) with three proper-noun names (Lips, Lamb, ABC), producing the correct behavior towards the correct object in each trial. Chaser responded to each object depending upon the meaning of the associated command – such that Chaser responded as though the commands and the proper-noun names represented independent morphemes. Chaser understood that names refer to particular objects, independent of the activity requested involving that object. Thus, Bloom's (2004) concern that Rico may not have understood the difference between “sock” and “fetch-the-sock” is ruled out in this study.

Chaser did not require special training before producing the appropriate response when the three commands (take, paw, nose) were paired with the three names (Lips, Lamb, ABC). Because Chaser had previously learned the meaning of the three commands and the names of the three objects, she was able to respond correctly, even on the first trial, when each of the three commands and the three nouns were randomly paired together in verb–noun phrases – demonstrating “combinatorial understanding” of two-word phrases. Chaser was able to understand the combination of the two words and behave appropriately in each trial, even though these commands and nouns had never been used in combination.

6.3. Experiment 3: learning common nouns

Experiment 3 showed that Chaser was successful in learning three common nouns—words that represent categories. Chaser learned the names of two concrete categories, “balls” and “Frisbees.” For each of these concrete categories, exemplars of each category displayed physical characteristics that may have helped define the category, such as all “balls” being relatively round. Chaser also successfully learned the more abstract concept “toy” that seemed to be based on function, rather than shared physical characteristics. “Toys” were objects for which she had learned unique proper-noun names. “Non-toys” were similar objects belonging to family members. The physical characteristics of toys and non-toys were not reliable discriminative features of the category – instead, toys and non-toys differed in functionality. Chaser could play with toys; she could not play with non-toys.

By forming categories represented by common nouns, Chaser mapped one label onto many objects. Conversely, Chaser also demonstrated that she could map up to three labels onto the same object without error. For example, Experiment 1 demonstrated that Chaser knew the proper-noun names of all objects used in this study. Chaser also mapped the common noun “toy” onto these same objects. Her additional success with the two common nouns “ball” and “Frisbee” demonstrates that she mapped a third label onto these objects. Her demonstrations of one-to-many and many-to-one noun/object mappings reveal flexibility in the referential nature of words in border collies.

6.4. Experiment 4: learning words by exclusion

Experiment 4 demonstrated that Chaser was able to use inferential reasoning by exclusion to map words upon objects, partially replicating Rico's performance (Kaminski et al., 2004). Chaser demonstrated strong retention of the name–object mapping when she was tested immediately after the exclusion choice trial. However, as indicated earlier, retention of the name of the object was reduced at the 10-min delay interval and essentially non-existent at the 24-h delay interval. These results support the conclusion of Horst and Samuelson (2005, 2006) that a single exclusion choice trial with toddlers may not be sufficient to create durable memory traces of word learning. Rehearsal or play that provides additional name–object pairings may be necessary for strong long-term retention, as we provided in Experiment 1.

6.5. Referential understanding of nouns

Recognizing that infants and young children often fall short of skilled referential understanding of words, Baldwin (1993) identified two basic factors that define a child's referential understanding of words that enable rapid word learning: (a) awareness that words may refer to objects, and (b) awareness of social cues that enable the mapping of the words upon the referent. Bloom (2004) expanded the view of referential understanding by pointing out, “They [children] appreciate that the word may refer to a category, and thereby can be used to request a sock, or point out a sock, or comment on the absence of one” (p. 1605).

Together, our four experiments provide compelling evidence that Chaser understood that objects have names. She demonstrated sensitivity to our social cues that that enabled her to map words upon the referents. To our knowledge, Chaser learned more proper-noun names (1022) than any other dog, fulfilling our strict learning criteria and satisfying 1000's of statistical tests. Chaser demonstrated the one-to-many name–object mappings required for common nouns to represent categories. Conversely, she also understood that an object may have more than one name, such as

a unique proper-noun name as well as one or more common-noun names (“ball” and “toy”). By successfully responding to random combinations of names and commands, she demonstrated her understanding that names refer to objects, independent of the behaviors directed toward those objects. Thus, she understood the difference in meaning between names and commands and treated their combination the same way humans do. Combined, these experiments provide clear evidence that Chaser acquired referential understanding of nouns, an ability normally attributed to children.

Acknowledgements

We give special thanks to Wayne West, owner of Fleet Hill Farms, Pauline, SC, for his breeding of Chaser, wise advice, and opening his gates for Chaser to fulfill her passion for herding sheep. Special thanks to Robin Pilley for videotaping and photographing Chaser on so many occasions. We thank P. Bloom for his helpful discussions of this work, and Kara Bopp for comments on an earlier draft. We thank Sally and Robin Pilley for their daily support, nurture, and help in training Chaser. We thank two Wofford College students (Caroline Reid and Katie Grainger) for their work and play with Chaser.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at [doi:10.1016/j.beproc.2010.11.007](https://doi.org/10.1016/j.beproc.2010.11.007).

References

- Aust, U., Range, F., Steurer, M., Huber, L., 2008. Inferential reasoning by exclusion in pigeons, dogs, and humans. *Anim. Cogn.* 11, 587–597.
- Baldwin, D.A., 1993. Early referential understanding: infants' ability to recognize referential acts for what they are. *Dev. Psychol.* 29, 832–843.
- Bloom, P., 2004. Can a dog learn a word? *Science* 304, 1605–1606.
- Carey, S., Bartlett, E., 1978. Acquiring a single new word. *Proc. Child Lang. Dev.* 15, 17–29.
- Erdőhegyi, A., Topál, J., Virányi, Z., Miklósi, A., 2007. Dog-logic: inferential reasoning in a two-way choice task and its restricted use. *Anim. Behav.* 74, 725–737.
- Fischer, J., Call, J., Kaminski, J., 2004. A pluralistic account of word learning. *Trends Cogn. Sci.* 8, 481.
- Golinkoff, R.M., Hirsh-Pasek, K., Bailey, L.M., Wenger, N.R., 1992. Young children and adults use lexical principles to learn new nouns. *Dev. Psychol.* 28, 99–108.
- Heffner, H., 1975. Perception of biologically meaningful sounds by dogs. *J. Acoust. Soc. Am.* 58, S124.
- Herman, L.M., Richards, D.G., Wolz, J.P., 1984. Comprehension of sentences by bottlenosed dolphins. *Cognition* 16, 129–219.
- Horst, J.S., McMurray, B., Samuelson, L.K., 2006. Understanding fast mapping and word learning in a dynamic connectionist architecture. *Proceedings from the 28th Meeting of the Cognitive Science Society*.
- Horst, J.S., Samuelson, L.K., 2005. Slow down: understanding the time course behind fast mapping. *Poster Session Presented at the 2005 Biennial Meeting of the Society for Research in Child Development*. Atlanta, GA.
- Horst, J.S., Samuelson, L.K., 2006. Turning novel names into known names: understanding the processes of fast mapping and word learning. *Poster Presented at the XVth Biennial International Conference on Infant Studies*. Kyoto, Japan.
- Kaminski, J., Call, J., Fischer, J., 2004. Word learning in the domestic dog: evidence for “fast mapping.”. *Science* 304, 1682–1683.
- Markman, E.M., Abelev, M., 2004. Word learning in dog? *Trends Cogn. Sci.* 8, 479–480.
- Markman, E.M., Wachtel, G.F., 1988. Children's use of mutual exclusivity to constrain the meaning of words. *Cogn. Psychol.* 20, 121–157.
- Range, F., Aust, U., Steurer, M., Huber, L., 2007. Visual categorization of natural stimuli by domestic dogs. *Anim. Cogn.*, [doi:10.1007/s10071-007-0123-2](https://doi.org/10.1007/s10071-007-0123-2).