THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Competition between phonology and semantics in noun class learning

Citation for published version:

Culbertson, J, Gagliardi, A & Smith, K 2017, 'Competition between phonology and semantics in noun class learning', *Journal of Memory and Language*, vol. 92, pp. 343-358. <https://doi.org/10.1016/j.jml.2016.08.001>

Digital Object Identifier (DOI):

[10.1016/j.jml.2016.08.001](https://doi.org/10.1016/j.jml.2016.08.001)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Journal of Memory and Language

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Competition between phonological and semantic cues in noun class learning

Jennifer Culbertson*, Annie Gagliardi, and Kenny Smith
University of Edinburgh, Edinburgh, United Kingdom

*Corresponding author: tel.: +44 0131 6515510, email: jennifer.culbertson@ed.ac.uk, mail: Dugald Stewart Building, 3 Charles Street, Edinburgh EH89AD, United Kingdom.

Abstract

Learning noun classification systems, like gender, involves inferring a language-particular set of (often probabilistic) cues to class membership. Previous work has shown that learners rely disproportionately on phonological cues (e.g., Karmiloff-Smith, 1981; Gagliardi & Lidz, 2014). Surprisingly, this occurs even when competing semantic cues are more reliable predictors of class. We investigate two possible explanations for this: first, that phonological cues are more salient to learners than semantic cues, and second that phonological cues are generally available earlier than semantic cues. We show that adult learners' treatment of conflicting cues to noun class in a miniature artificial language depends on *both* cue saliency and early availability. Importantly, learners prioritize earlier-available cues even when they are less salient than competitor cues. Our findings suggest a possible mechanism for children's over-reliance on phonology: children start building their classifications systems very early, when phonological information is available, but word meanings are not.

Keywords: noun classification; gender; artificial language learning; cue competition; category learning; language acquisition

Introduction

Noun classification systems are grammatical devices, found in many languages, which serve to categorize nouns into classes according to some set of features. They include gender systems made up of a small set of classes (e.g., 'masculine'/'feminine' as in the Romance languages), larger noun class systems (with up to 10 classes, as in many Bantu languages), and numeral classifier systems (sometimes with hundreds of distinct categories, as in many East Asian languages). These systems differ in their formal realization, the locus of the cues to class (e.g., based on distributional or morphological information, the semantics of noun referents, or phonological properties of nouns themselves) and in the particular cue features which are relevant.

Interestingly, all noun classification systems exploit semantic cues to some extent; there are no known noun classification systems which are based on phonological cues alone (Aikhenvald, 2000). In fact, the set of semantic features used is often similar across languages: natural gender, animacy, and shape are very common (Denny, 1976; Dixon, 1986; Lakoff, 1987; Comrie, 1989, Aikhenvald, 2000; Senft, 2000). That said, many languages exhibit mixed systems, with semantic cues like natural gender along with a set of noun-internal phonological cues which are predictive of class. The particular set of phonological features used varies widely across languages. For example, a prototypical mixed gender

system can be seen in French. Natural gender is a highly reliable cue to class: nouns with female gender referents are typically found in one class, while nouns with male gender referents are in another class. Other semantic cues are less reliable, but nevertheless probabilistically cue class; nouns referring to paths or roads are typically feminine, and nouns whose referents are superordinate categories are typically masculine (Nelson, 2005). Alongside semantics, many phonological features of nouns probabilistically cue class; according to Surridge (1986), there are 34 distinct suffixes which cue one gender or the other in French (see also Lyster, 2006).

This combination of semantic and phonological cues, some more reliable than others, leads to a complex problem for both first and second language learners (MacWhinney, 1978; Braine, 1987; Levy, 1988; Frigo & MacDonald, 1998; Carroll, 1999; Kempe & Brooks, 2001; Taraban, 2004; Arnon & Ramscar, 2012; a.o.). Here, we focus on a recurring finding in the literature on first language acquisition of noun classes: the weak role of semantic cues relative to phonological cues. In her classic work on the acquisition of French gender, Karmiloff-Smith (1981) found, across a series of experiments, that children even up to age 10 determine the gender of novel nouns according to their phonological properties (e.g., word endings), rather than exploiting highly reliable semantic cues like natural gender. For example, in one experiment, children (age 3-12) were presented with two pictures of unfamiliar alien characters, one clearly female, the other male. These were given noun labels whose phonology either conflicted with or matched the natural gender of the referents. For example, they might see a male alien and hear the label *podelle*, where *-elle* is a cue to feminine grammatical gender, or a female alien labelled *bicron*, where *-on* is a cue to masculine grammatical gender. The children were subsequently asked to produce these nouns in contexts which required a determiner whose form is determined by grammatical gender (e.g., *le/la*, ‘the_{masc}/the_{fem}’). For example, the experimenter might put an object on one of the pictures, and ask *Qu’est-ce que j’ai fait?* ‘What did I do?’, and the child might answer *Vous avez mis un crayon sur la podelle* ‘You put a pencil on the_{fem} podelle’. Children up to the age of 10 consistently used the determiner which matched the phonological features of the noun, apparently ignoring the semantic cue (e.g., producing *la podelle* for a male alien or *le bicron* for a female alien).

These findings have been reproduced in several other languages including German (Mills 1985), Spanish (Perez-Pereira 1991; Mariscal 2009), Sesotho (Demuth, 2000; Demuth & Ellis, 2008), Russian (Rodina & Westergaard, 2012) and Tsez (Gagliardi, 2012; Gagliardi & Lidz, 2014). Both natural gender and the phonological endings used in Karmiloff-Smith (1981) are highly reliable cues to class in French—even though both are only relevant for a subset of nouns in the language (e.g. natural gender is irrelevant for inanimates, *-elle* and *-on* endings only appear on some nouns), when they are present an ideal observer could use either cue to guess the grammatical gender with high accuracy. Importantly, in the case of Tsez, Gagliardi and Lidz (2014) found that even when phonological cues were statistically *less* reliable than competing semantic cues, young children still preferentially use them to determine class membership. They report that this effect disappears in older children and adults. In contrast to this body of evidence suggesting a preference for phonological over semantic cues, we are aware of only one study which finds that children use semantic cues preferentially (in Icelandic; Mulford, 1985). Interestingly, Mulford suggests this result may

be due to the extremely low predictive power of phonological cues in Icelandic (e.g., relative to many such cues in French).

Given what we know about noun classification systems in general, this apparent over-reliance on phonological cues during learning is surprising: semantic cues to class are ubiquitous in language, are in some cases highly reliable, and often pick out very similar properties across languages. By contrast, only some languages make use of phonological cues, and these cues are often less predictive, and highly variable. Why then, would children rely on phonology over semantics? Several distinct mechanisms have recently been proposed to explain this puzzling finding. For example, learners may be actively biased against using external cues like semantics in forming grammatical categories, particularly when internal properties of nouns are available to cue class (Gagliardi, 2012; Culbertson & Wilson, 2013; Gagliardi & Lidz, 2014; Gagliardi et al., in press). Perhaps relatedly, semantic cues may simply be less salient to learners than phonological cues, at least in the context of acquiring noun classification systems (Gagliardi, 2012; Gagliardi et al., in press). Alternatively, learners may have access to phonological properties of nouns before semantic features of their referents (Polinsky & Jackson, 1999; Carroll, 1999; Demuth, 2000; Culbertson & Wilson, 2013; Gagliardi et al., in press), simply because children encounter linguistic forms prior to mapping those forms to their referents. Under this latter explanation, there is no inherent bias against the use of semantic cues (they may even be *more* salient). However, if learners build representations of categories first based on phonology, before they have acquired the relevant word-to-meaning mappings, these initial representations may persist—either because representations initially built on purely distributional properties take some time to incorporate external cues, or because of an asymmetry in the relative amount of data learners have for each type of cue, with more data being available for early-available cues.

Which of these hypotheses is correct has clear implications for theories and models of the acquisition of noun classification systems. Here we conduct the first experimental tests of the proposed mechanisms outlined above, focusing on salience and early availability of cues. We use an artificial language learning paradigm in order to maintain complete control over the reliability and frequency of different cues. Following previous studies of artificial noun class learning (Braine et al., 1990; Brooks et al., 1993; Frigo & MacDonald, 1998; Williams, 2004; Culbertson & Wilson, 2013), we use adult learners: the hypotheses we are testing do not specifically distinguish child from adult learners, and if certain cues are more salient than others, we expect they will be for learners of any age. Our laboratory learning task allows us to make some cues available to learners earlier than others, re-creating a learning environment for adults that would be similar to first language acquisition in this key aspect.

Previous work has shown that properties like relative salience, availability (or frequency), and reliability of cues are at work in L1 and L2 acquisition of morphological patterns (e.g., Bates & MacWhinney, 1989; Goldschneider & DeKeyser, 2001). A number of studies have used artificial language learning to investigate these properties in the context of noun class learning. While adult learners do not acquire completely arbitrary systems of classification—that is, systems in which there are no semantic or phonological cues to class (Braine et al., 1990)—they can exploit cues which are present on only a subset of nouns, or which only probabilistically (non-deterministically) cue class membership (Brooks et al., 1993; Frigo & MacDonald, 1998). Further, several studies have manipulated the particular

features used to cue class in order to establish whether particular phonological and semantic cues differ in their salience. Frigo and MacDonald (1998) taught adult learners an artificial language in which different words for greetings (e.g., ‘good morning’ and ‘good evening’) served as class markers, used with different noun referents. Phonological properties of the nouns provided the only cues to class, and these consisted of either a prefix, a suffix, or both. Moreover, these markers were more or less substantial in form—some were made up of a single segment, others of a CVC syllable. For example, in the CVC suffix condition, participants were shown pictures of people, and heard a greeting (either *jai* or *fow*) and label for that picture, e.g., *jai chagor* (where *-gor* was a recurring cue to noun class). Frigo and MacDonald found that adults were most successful at learning systems with more salient cues: having both a prefix and a suffix was better than having only one or the other. Similarly, classification was better with a marker consisting of a CVC syllable rather than a single segment. Interestingly, a prefix alone was better than a suffix alone, contrary to previous research suggesting learners prefer suffixal grammatical markers (Hawkins & Cutler, 1988; Hupp et al., 2009; Bruening et al., 2012). Because nouns followed greetings in their task, this preference for prefixes may be due to increased salience of adjacent relationships compared to non-adjacent ones in the input (Newport & Aslin, 2004; Gomez & Maye, 2005).

Culbertson and Wilson (2013) taught adults an artificial numeral classifier system, in which nouns appearing with a numeral were accompanied by a marker determined completely by semantic properties of the noun referents. The stimuli were familiar objects, paired with their English labels so that no new phonology or semantic knowledge was required: learners simply had to map objects into classes based on the available semantic cues. For example, they might see a picture of a pencil, described as *one-ka pencil*, or a picture of two blankets, described as *two-po blanket* (where *po* and *ka* are class markers). Shape-based classification systems, in which narrow objects (e.g., a pencil) were distinguished from flat objects (e.g., a blanket) were acquired relatively easily; in contrast, systems based on the flexibility of a referent (e.g., rigid objects like a pencil distinguished from flexible objects like a chain) were less accurately learned (Culbertson & Wilson, unpublished results). This is consistent with the typology of classifier systems, where shape is the most common basis of classification for numeral classifier systems (Lee, 1987), and is frequently featured in noun class systems as well (Dixon, 1986; Senft, 2000). Furthermore, shape often serves as the primary basis of classification while flexibility is almost exclusively a secondary feature, for example further distinguishing objects already categorized by shape (Adams & Conklin, 1973).¹

In Experiment 1, we exploit the findings of Frigo and MacDonald (1998) and Culbertson and Wilson (2013) in order to compare learning of systems in which phonological *and* semantic cues of varying salience directly conflict with each other. This allows us to test whether learners’ reliance on phonological or semantic cues is dependent on the relative salience of the cues in question. In Experiment 2 we then investigate the role of cue

¹ The frequent use of shape to classify objects cross-linguistically is consistent with the shape bias documented in object categorization and labeling tasks more generally (Smith et al., 2002; Landau et al., 1988).

availability, in order to test whether the effects of salience can be altered by early access to particular cues.

Experiment 1: Cue Salience

We constructed artificial noun class systems in which novel nouns (referring to concrete objects) appeared with one of two class markers (introduced as words for ‘the’). These markers served as a distributional cue to categorize the nouns into two non-overlapping classes. However, the cue determining class membership was, by design, ambiguous during training; each class was made up of nouns consistently marked by a distinct phonological cue (e.g., a particular suffix), and which had referents in a single semantic category (e.g., animate objects). This setup created multiple cues to class—one phonological, and one semantic—which were equally frequent and reliable, with no evidence of differing ‘conflict validity’ (Bates & MacWhinney, 1989). After exposure to this ambiguous data, at test participants were required to classify stimuli in which these cues conflicted, in order to assess which cue more strongly determined their classification decisions. We manipulate cue salience by using two phonological cue types and three semantic cue types which previous experiments (Frigo & MacDonald, 1998; Culbertson & Wilson, 2013, unpublished data) have shown to produce different learning outcomes. The phonological cues we use are a prefix and suffix combination (high salience), and a suffix alone (low salience). The semantic cues we use are animacy (high salience), shape (medium salience), and flexibility (low salience). While animacy has not been shown to be higher salience than shape, there are a number of good reasons to believe it may be. First, animacy is the most commonly used feature across all types of noun classification systems (Dixon 1986). Second, sensitivity to animacy is apparent very early in child grammars (e.g., Becker, 2009; Bunker & Lidz, 2006). This setup mirrors conceptually the natural language acquisition experiments conducted by Karmiloff-Smith (1981) and others, while manipulating the particular cues in question.

Participants

Participants were 192 English-speaking adults recruited via Amazon Mechanical Turk (32 per condition). Participants were paid \$1, the experimental session lasted approximately 10-15 minutes.

Stimuli/materials

Visual stimuli consisted of pictures of real-world objects drawn from one of six semantic categories: animate or inanimate (animacy cue), narrow or flat (shape cue), and flexible or rigid (flexibility cue). Example visual stimuli can be found in Appendix A. Where possible the same objects (and corresponding pictures) were used to instantiate multiple categories (e.g., ‘ribbon’ was in the narrow category for shape cue, and also the flexible category for the flexibility cue). Each category included 48 pictures. The lexicon of the language was made up of 96 CVC(V) nonce roots (see Appendix B for full list, roots conformed to English phonotactics, as assessed using an online phonotactic probability calculator: Vitevich & Luce, 2004). Phonological cues were added to these roots: either a prefix and suffix together (*di*+root+*te* or *ba*+root+*po*), or a suffix alone (root+*te* or root+*po*). Words were displayed

both orthographically and auditorily to participants. Auditory stimuli were created using the MacTalk speech synthesizer (OS 10.9, voice “Alex”). The non-word class markers in the language were two chosen from the set (*kuh* [kʌ], *shae* [ʃæ], *muh* [mʌ], *gae* [gæ]), with the combination used counterbalanced across participants. These class markers were distinct from the nouns and were post-nominal (and therefore adjacent to the suffix). For example, participants might see a picture of a ribbon accompanied by the auditory stimulus *divokte kuh*.

Design

Each condition featured a semantic cue (animacy, shape, or flexibility) *and* a phonological cue (a combination of prefix and suffix, or a suffix alone, these are referred to throughout as *affixes* and *suffix* respectively). This gives a total of six conditions, as illustrated in Table 1. The experiment used a design in which training data are ambiguous between hypotheses of interest, with held-out disambiguating data presented at test (sometimes called ‘Poverty-of-the-Stimulus’ design, Wilson, 2006; Culbertson & Adger 2014, a.o.). In this case, the held-out data are cases in which the semantic and phonological cues conflict. During training each class is perfectly correlated with a semantic feature (e.g., animacy) *and* a phonological feature (e.g., one of two suffixes). For example, in the animacy/suffix condition, animate objects are labelled with nouns ending in *-te*, while inanimate objects are labelled with nouns ending in *-po*. The animate/*-te* stimuli always appear with one distributional marker (e.g., *kuh*), while the inanimate/*-po* stimuli always appear with the other marker (e.g., *shae*). At test, participants were presented with an object and its associated label, and were required to choose the appropriate definite article; included in the test set were items in which the semantic and phonological cues (which aligned during training) conflicted, for example an animate object labelled with a noun ending in *-po*. In this case, participants could choose *kuh* since all animate objects took this article during training, *or* they could choose *shae* since all nouns ending in *-po* took that article during training.

Table 1
Conditions in Experiment 1, ordered according to decreasing semantic (columns) and phonological (rows) salience.

Semantics				
high	medium	low		
animacy affixes	shape affixes	flexibility affixes	high	Phonology
animacy suffix	shape suffix	flexibility suffix	low	

Each participant was randomly assigned to a condition (HITs [‘Human Intelligence Tasks’, Amazon’s term for tasks] were posted simultaneously for each condition in small batches over several days). For each participant 48 pictures (24 in each semantic category for the relevant cue type) and 48 roots were randomly chosen to make up the training set. In this

set, the relevant semantic and phonological cues aligned as described above. The test set was comprised of those trained items, plus an additional 48 pictures and roots in which the cues conflict. The mapping between pictures and roots was also randomly assigned for each participant. The order of presentation was fully randomized within both training and testing for each participant.

Procedure

Participants were told that they were going to learn part of a new language which had two words for ‘the’.² They then listened to these two words in isolation. In the training phase of the experiment, participants saw a picture and heard a noun accompanied by one of the definite articles marking class. The noun was also presented orthographically below the picture. The two possible class markers were presented as choices, and participants had to click on the one they had just heard, as shown in Fig. 1A, B. If they clicked the wrong word, the trial was repeated.

Fig. 1. Example trials in Experiment 1 for animacy/suffix condition. **A**: Training trial (animate object with *-po* noun). **B**: Training trial (inanimate object with *-te* noun). **C**: Conflicting test trial (animate object with *-te* noun).

During the testing phase, participants saw a picture and heard a noun (also orthographically presented), but no word for ‘the’. They were instructed to choose the word for ‘the’ they thought the speaker would be most likely to use. No feedback was provided. An example trial is shown in Fig. 1C.

Results

Recall that during training the data were ambiguous: use of a given class marker could have been based on the semantic cue present, or the phonological cue present. Thus learners could have formed representations of the class which were based on one or the other hypothesis about class membership. Responses to disambiguating trials at test allow us to see whether

² For complete instructions, see Appendix C.

learners determine membership of conflicting items predominantly based on the phonological or semantic cue, and how this is affected by cue salience.

Fig. 2 shows the results for each of the six conditions in terms of the proportion of trials on which marker choice followed the semantic cue used in training. For aligned trials, this choice was also consistent with the matching phonological cue used in training. Performance on these trials indicates whether participants learned the classification system they were trained on. For conflicting trials, choice of the marker corresponding to the semantic cue necessarily conflicted with the marker used for the phonological cue in training.³ Performance on these trials indicates whether participants base their classification on semantics or phonology when the two conflict. The results suggest that both measures were affected by the relative salience of the cues. Mixed-effects logistic regression models with helmert contrast coding were used to analyze the data.⁴

For aligned trials, performance is relatively high across all conditions, but an effect of salience is apparent for the low salience semantic cue; when the semantic cue was flexibility, participants were less accurate at choosing the correct class marker on aligned trials, suggesting they did not learn the system as well (flexibility vs. animacy, shape: $\beta = -0.60 \pm 0.14$, $p < 0.001$). There was a marginally significant interaction driven by lower accuracy when the semantic cue was shape and the phonological cue was suffix ($\beta = -0.48 \pm 0.25$, $p = 0.05$). All other fixed effects and their interactions were not significant.

Turning to the conflicting trials, Fig. 2 illustrates clearly that both salience manipulations affect which cue learners base their choice of marker on. Independent of the phonological cue present, learners were most likely to use the semantic cue if it was animacy, and least likely if it was flexibility, with shape falling in between. This is confirmed statistically: learners were less likely to use shape than animacy ($\beta = -2.11 \pm 0.42$, $p < 0.001$),

³ We plot the individual participant outcomes on Fig. 2 because the data are bimodally distributed, making the average values misleading. In general participants are individually highly consistent, and appear to base their choice of class in conflicting trials on *either* the semantic *or* the phonological cue exclusively, rather than using some combination of both cues across trials. The likelihood with which they choose one or the other, however, is clearly affected by salience, as revealed by our statistical analysis. The tendency for participants to be consistent in their cue use may be related to the more general bias against unpredictable variation seen in some other artificial language learning studies with adult participants (e.g. Real & Griffiths, 2009; Smith & Wonnacott, 2010; Culbertson et al., 2012).

⁴ Mixed-effects logistic regression models for all experiments were run using lme4 (Bates, 2010), with by-participant, by-picture, and by-word (root) random intercepts. Random slopes for words were not included due to convergence problems (other random effects were between-subjects). For Experiment 1, helmert contrast coding was used. Helmert coding tests the hypothesis that levels of a factor are related in a step-wise fashion; the model compares each subsequent level of a factor to the mean of the previous levels (the intercept is the grand mean). For all models, the semantic cue is ordered *animacy* > *shape* > *flexibility*, meaning that helmert contrasts compare shape to animacy, and flexibility to the mean of animacy and shape, and for the phonological cue *affixes* > *suffix* (suffix is compared to affix). P-values are obtained from Wald-Z tests.

and less likely to use flexibility than shape and animacy ($\beta = -0.86 \pm 0.20$, $p < 0.001$). Independent of the semantic cue present, learners were more likely to use the semantic cue (that is *less* likely to use the phonological cue) if the phonological cue was a suffix alone rather than a prefix and suffix combination ($\beta = 1.11 \pm 0.29$, $p < 0.001$). There was also a significant interaction between semantic and phonological cue types, driven by less use of the flexibility cue in the suffix condition than predicted by these two factors alone ($\beta = -0.59 \pm 0.24$, $p = 0.01$). This reflects the larger difference between shape and flexibility when the phonological cue is suffix alone. There was no parallel interaction between shape and phonological cue ($\beta = -0.21 \pm 0.40$, $p = 0.60$).

Fig. 2. Experiment 1 results. Bars represent average proportion of trials on which marker choice followed the *semantic* cue used in training, points are individual participants. Error bars represent bootstrap 95% confidence intervals. In aligned trials (seen during training), both cues were in agreement, so higher values indicate responses consistent with the semantic and phonological cue used in training. Conflicting trials paired a given semantic cue with the conflicting phonological cue, and higher values therefore indicate a preference for using the semantic cue over the phonological cue.

Discussion

Our results reveal that under these experimental conditions, participants were more likely to choose a category marker based on the semantic cue when the absolute salience of the semantic cue is high (e.g., animacy) or when the absolute salience of the phonological cue is

low (e.g., suffix alone). These findings confirm that, as in many other domains of grammar (e.g., Slobin, 2001), cue salience matters in the acquisition of noun classification systems (Frigo & MacDonald, 1989). They also show that learners' reliance on a given type of cue depends not just on the salience of that cue relative to other cues of that type, but also relative to other cues present: when multiple types of cues are available, learners preferentially exploit the more salient cue. Interestingly, our low salience semantic cue condition led to particular poor results, even when the competing phonological cue was not a high relatively low salience one: participants were relatively unlikely to use flexibility as a cue to class, regardless of the salience of the competing phonological cue.

Is cue salience alone a likely explanation for children's over-reliance on phonology in natural language acquisition of noun classes? Gagliardi et al. (2014) report that Tsez-acquiring children used relatively low salience phonological cues (e.g., the initial segment of a noun) over high salience semantic cues (e.g., animacy) when determining the class of novel nouns; similarly, in Karmiloff-Smith (1981), French children were found to use noun endings rather than the presumably high salience natural gender cue. This observation remains mysterious given our results: phonological cues do not appear to be more salient for learners, nor is there evidence for any systematic bias against semantics in the context of noun class learning (a possibility suggested in Gagliardi et al., in press).

In Experiments 2-4 we therefore investigate a second proposed mechanism for the apparent over-reliance on phonological cues during learning. We manipulate the availability of cues at various stages of learning in order to test whether having access to a particular type of cue earlier might increase learners' likelihood of using it.

Experiment 2-4: Cue availability

In this set of experiments, we test whether having a cue available earlier in the learning process changes the extent to which learners rely on it. We do this by staging the input data; in the first stage, learners are trained on data which contains evidence for one type of cue only, in the second stage they are trained on data which provide evidence for both types of cues together. Children are likely able to learn from purely distributional and noun-internal phonological cues in the ambient language well before they have formed robust associations between nouns and their semantic referents, meaning that they should have access to phonological cues to noun class before they have access to semantic cues. Having earlier access to a particular cue may simply mean more exposure to that cue (and therefore better learning), or it may influence early formation of representations in a way that persists even when other cues become readily available. In either case, we can mirror this developmental situation experimentally by making phonological cues available before semantic cues (the 'phonology first' condition), or we can reverse this 'natural' order by making semantic cues available before phonological cues (the 'semantics first' condition).

Experiment 2

In Experiment 2 we use the two high salience cues from Experiment 1—animacy and affixes—and compare participants' behavior on conflicting test trials depending on whether the semantic or phonological cue was present in the first stage.

Participants

Participants were 64 English-speaking adults recruited via Amazon Mechanical Turk (32 per condition). Participants were paid \$2, the experimental session lasted approximately 15 minutes.

Stimuli/materials

Stimuli were identical to those used in the animacy/affixes condition of Experiment 1.

Design

Cue salience was held constant by using both a high salience semantic and phonological cue. However, participants were randomly assigned to one of two conditions. In the phonology first condition, they were first trained on stimuli including only a noun and a corresponding class marker. In the semantics first condition, participants were first trained on stimuli including only a picture and a corresponding class marker. These conditions are illustrated in Figure 3.

Fig. 3. Example trials for staged conditions in Experiment 2. A: Phonology first condition. In the first stage of training, participants are exposed to the determiner and noun form without seeing the associated referent. B: Semantics first condition. In the first stage of training, participants hear the determiner and see the referent, without being exposed to the phonological form of the noun.

Procedure

As in Experiment 1, participants were told that they were going to learn part of a new language which had two words for 'the'. They then listened to these two words in isolation. As shown in Fig. 3, in the semantics first condition participants then saw trials in which a picture of an object appeared, and they heard one of the two words for 'the'. They were instructed to click on the one they heard. In the phonology first condition, participants saw trials in which a noun appeared, and they heard that noun followed by one of the two words for 'the'. They were instructed to click on the one they heard. For both conditions, as in Experiment 1, clicks on the wrong word resulted in the trial being repeated. Upon completion

of this first stage of the experiment, participants moved on to a second training phase with the same items, but introducing the second cue. This stage which was identical across participants, and the same as the training phase of Experiment 1.⁵ In other words in the second stage of training, both a semantic and a phonological cue were present for all learners, and both were equally predictive of class. The testing phase was also identical across participants and the same as the testing phase of Experiment 1; participants in Experiment 2 therefore received more training in total than participants in Experiment 1.

Results and discussion

Our hypothesis is that if a cue is available earlier in learning, learners should be more likely to rely on it to determine class membership relative to other cues present at later stages. This predicts that when the two competing cues are both high salience, the earlier available cue will be used to determine class. This is clearly borne out by the data, shown in Fig. 4: in conflicting test trials, learners in the phonology first condition were less likely to base their choices on the semantic cue (i.e. were more likely to use the phonological cue) than learners in the semantics first condition. This was confirmed by a mixed-effects logistic regression model⁶ predicting use of the semantic cue on conflicting trials from staging ($\beta = 2.38 \pm 0.68$, $p < 0.001$).

Fig. 4. Experiment 2 results: effect of staging on high salience competing cues. Error bars represent bootstrap 95% confidence intervals.

⁵ For complete instructions, see Appendix C.

⁶ Mixed-effects logistic regression models for Experiments 2-4 used helmert contrast coding comparing the phonology first condition to the semantics first condition unless otherwise noted.

This result confirms that staging can influence the extent to which learners use high salience cues. In Experiment 3 we investigate how early availability affects a weaker semantic cue depending on the salience of the competing phonological cue.

Experiment 3

In this experiment we use the mid-salience semantic cue from Experiment 1, shape, and compare the effect of staging when the phonological cue is either weak (suffix alone, Experiment 3a) or strong (affixes, Experiment 3b).

Participants

Participants were 128 English-speaking adults recruited via Amazon Mechanical Turk (32 per condition). Participants were paid \$2, the experimental session lasted approximately 15 minutes.

Stimuli/materials

Stimuli were identical to those used in the shape/affixes and shape/suffix condition of Experiment 1.

Design

The semantic cue salience was held constant by using shape across conditions. However, phonological cue and staging were crossed; participants were either trained with a suffix only or affixes, and were trained either on semantics first or phonology first.

Procedure

The procedure was identical to Experiment 2.

Fig. 5. Experiment 3a results: effect of staging on shape/suffix. The first two sets of bars show use of the semantic cue when either that cue or the phonological cue was available first. The third reproduces the results of the shape/suffix condition in Experiment 1 (cues available at the same time). Error bars represent bootstrap 95% confidence intervals.

Fig. 6. Experiment 3b results: effect of staging on shape/affixes. The first two sets of bars show use of the semantic cue when either that cue or the phonological cue was available first. The third reproduces the results of the shape/affixes condition in Experiment 1 (cues available at the same time). Error bars represent bootstrap 95% confidence intervals.

Results and discussion

Fig. 5 and 6 show shape/suffix and shape/affix respectively, with both types of staging and unstaged (data from Experiment 1). Looking at the two staging conditions, the results of this experiment suggest that the effect of staging on use of a mid-salience semantic cue is quite strong, regardless of the phonological cue salience. In both the suffix and affixes condition, having the phonological cue available first leads to lower rates of use of the semantic cue compared to having the semantic cue first. This is confirmed by a mixed-effects model on the combined data of Experiments 3a and 3b, revealing a significant effect of staging ($\beta = 0.95 \pm 0.26$, $p = 0.002$), a marginal effect of phonological cue salience ($\beta = -0.52 \pm 0.28$, $p = 0.06$), and no interaction between staging and phonological cue salience ($\beta = 0.15 \pm 0.31$, $p = 0.64$).

We can also compare this data to the original results reported in Experiment 1. However, one caveat is in order here: recall that participants in Experiment 3 receive more training data in total for one of the cue types (48 vs. 96 trials). This mirrors our hypothesis that children have access to some cues earlier than others, which leads to that cue being available earlier *and* to a greater total volume of exposure to that cue. Despite the difference in overall volume of training data, there was no significant difference between the staged and not staged conditions in aligned trials (not staged vs. semantics first: $\beta = 0.08 \pm 0.43$, $p = 0.86$; not staged vs. phonology first: $\beta = 0.49 \pm 0.42$, $p = 0.25$). However, conflicting trials reveal an

interesting asymmetry. Specifically, in the suffix conditions, the effect of putting phonology first is stronger than the effect of putting semantics first; compared to Experiment 1, the phonology first suffix condition results in much lower use of the semantic cue. By contrast, in the semantics first suffix condition, learners do not appear more likely to use the shape cue relative to Experiment 1. This is confirmed by a mixed-effects model, revealing a significant difference from the no staging condition only for the phonology first condition ($\beta = -3.72 \pm 1.07$, $p < 0.001$), not for the semantics first condition ($\beta = -0.61 \pm 1.02$, $p = 0.55$). For affixes (the stronger phonological cue), this is reversed; having the phonology first does not appear to have a strong effect compared to Experiment 1, but having the semantics first leads to higher levels of reliance on the shape cue. This is again confirmed by a mixed-effects model, here revealing a significant difference from the no staging condition only for the semantics first condition ($\beta = 2.66 \pm 1.08$, $p = 0.01$), not for the phonology first condition ($\beta = -0.61 \pm 1.01$, $p = 0.57$).⁷ This suggests that staging is most likely to have an effect on relatively weak cues.

The natural language acquisition findings suggest that even weak, low salience phonological cues may also be strengthened relative to higher salience semantic cues if they are available earlier. We test this explicitly in Experiment 4a. We also test the inverse prediction: that a weak *semantic* cue may be strengthened relative to a strong phonological cue (Experiment 4b).

Experiment 4a

In Experiment 4a we ask whether a relatively weak phonological cue—namely a suffix alone—can be strengthened by early availability, even when the competing semantic cue is highly salient. We do this by comparing unstaged data from the animacy/suffix condition of Experiment 1 with a new phonology first condition. In this condition, the suffix cue is available first, before the high salience animacy cue is introduced.

Participants

Participants were 32 English-speaking adults recruited via Amazon Mechanical Turk. Participants were paid \$2, the experimental session lasted approximately 15 minutes.

⁷ Mixed-effects models comparing no staging to semantics first and phonology first are reported above with treatment coded contrasts (using no staging as a baseline). However, this results in relatively high correlations of fixed effects. Therefore, both models were also run with sum coded contrasts (comparing semantics and phonology first conditions to the grand mean). These models replicated the basic finding reported in the text. For aligned trials, no significant differences are found. For conflicting trials, phonology first is significantly different from the grand mean for shape/suffix ($\beta = -2.28 \pm 0.56$, $p < 0.001$), while semantics first is significantly different from the grand mean for shape/affix ($\beta = 1.98 \pm 0.60$, $p < 0.001$). In the latter case, phonology first is also different from the grand mean ($\beta = -1.29 \pm 0.62$, $p = 0.04$).

Stimulus, design and procedure

Stimuli were identical to those used in the animacy/suffix condition of Experiment 1. The design and procedure follow the phonology first condition in Experiment 2.

Results and discussion

Fig. 7 shows learners' use of the semantic cue with and without staging. The unstaged results here are those reported for the flexibility/affixes condition in Experiment 1, when both the phonological and semantic cues are simultaneously available from the start of training. For aligned trials, a model comparing use of the semantic cue (animacy) across the staged and unstaged experiments reveals no significant difference ($\beta = -0.48 \pm 0.48$, $p=0.31$). For conflicting trials however, a model comparing staged to unstaged confirms that making the weak phonological cue available first significantly decreases learners' use of the semantic cue on conflicting trials ($\beta = -4.4 \pm 1.5$, $p<0.001$). In other words, they use the phonological cue more. The results of this experiment provide clear evidence that earlier availability of a weak phonological cues can lead to increased use of this cue even when competing semantic cues are highly salient. This experiment brings us the closest to replicating the situation that both Tsez and French child learners are in.

Fig. 7. Experiment 4a results: effect of staging on weak phonological cue. The first set of bars shows use of the semantic cue when the phonological cue was available first. The second set of bars reproduces the results of the animacy/suffix condition in Experiment 1 (cues available at the same time). Error bars represent bootstrap 95% confidence intervals.

Experiment 4b

In our final experiment, we ask whether a similar effect of staging is found for weak *semantic* cues, by pitting the weakest semantic cue (flexibility) against the strongest phonological cue (affixes).

Participants

Participants were 32 English-speaking adults recruited via Amazon Mechanical Turk. Participants were paid \$2, the experimental session lasted approximately 15 minutes.

Stimulus, design and procedure

Stimuli were identical to those used in the flexibility/affixes condition of Experiment 1. The design and procedure follow the semantics first condition in Experiment 2.

Results

Fig. 8 shows learners' use of the semantic cue with and without staging. The unstaged results here are those reported for the flexibility/affixes condition in Experiment 1, when both the phonological and semantic cues are simultaneously available from the start of training. A comparison of these results indicates at first glance that there is no effect of making the weak semantic cue available first. Indeed, a model comparing use of the semantic cue (flexibility) across the staged and unstaged experiments confirms that making the weak semantic cue available first does not have a significant effect on learners' likelihood of using the semantic cue on conflicting trials ($\beta = -0.23 \pm 0.44$, $p = 0.60$). However, (unlike in Experiments 2-3 and 4a) in this case there is a significant effect on staging on the aligned trials; participants are more successful at learning the categorization system if the semantic cue is available first ($\beta = 0.89 \pm 0.40$, $p = 0.03$). Although this is a relatively subtle effect, it suggests that staged presentation is useful for this difficult-to-learn cue. The benefit could derive from having the low salience cue present first, or having it present in isolation (without the distraction of a 'competing' phonological cue), or simply of having additional data for that cue.

Fig. 8. Experiment 4b results: effect of staging on weak semantic cue. The first set of bars shows use of the semantic cue when that cue was available first. The second set of bars reproduces the results of the flexibility/affixes condition in Experiment 1 (cues available at the same time). Error bars represent bootstrap 95% confidence intervals.

Discussion

Experiment 4 investigated the effect of staging on weak phonological and semantic cues in order to test whether these cues might be strengthened if they are available to learners earlier. In Experiment 4a we found that weak phonological cues are indeed strengthened by early availability—confirming our hypothesis based on previous studies of natural language acquisition of noun classes. Experiment 4b found a small but reliable effect of staging on learning the classification system, but failed to find an effect of staging on the use of the weak semantic cue when it conflicts with a high salience phonological cue. There are (at least) two possible explanations for the difference between these weak phonological and semantic cues. It may be that this reflects some inherent difference in the way these cue types are learned. However, we believe it is more likely due to the extremely low salience of flexibility, also reflected in the rarity of this feature as a primary cue to classification cross-linguistically (Adams & Conklin, 1973). In other words, it may be our low salience semantic cue is less salient than our low salience phonological cue. This is in fact suggested by the data from Experiment 1, in which flexibility is unlikely to be used to determine class membership regardless of the salience of the competing phonological cue.

General Discussion

Studies of the acquisition of noun class systems in languages including French, German, Spanish, Russian, Sesotho and Tsez all suggest that, counterintuitively, child learners rely on phonological cues to class even in the face of highly salient, reliable semantic cues. For example, in French, natural gender, likely a highly salient cue, is strongly predictive of grammatical gender. However, when presented with novel nouns, French children appear to ignore natural gender, and assign grammatical gender based on phonological cues. Findings like this have led researchers to posit a number of potential mechanisms underlying this behavior (e.g., Mulford, 1985; Polinsky & Jackson, 1999; Demuth, 2000; Gagliardi & Lidz, 2014; Gagliardi et al., in press; Culbertson & Wilson, 2013). One hypothesis is that learners are biased against using external cues like semantics in forming grammatical categories. However, previous studies using artificial language learning to investigate noun class acquisition have suggested that (at least when no phonological cues are present), adult learners readily make use of semantic cues (Culbertson & Wilson, 2013). It could be that a bias against semantics is found only in child learners, or that the mere presence of noun-internal phonological cues triggers a bias against noun-external cues. Alternatively, perhaps no inherent bias against semantics exists, and rather, preferential use of phonological information is the result of differences in salience and availability of cues. In the experiments reported above, we investigated these two hypotheses experimentally, using artificial language learning with adult participants.

In Experiment 1, we manipulated the relative salience of competing phonological and semantic cues to class, holding constant their reliability. Participants were trained on a system in which both a semantic *and* phonological cue were equally predictive of class membership. During testing, they were given trials in which the two types of cues conflicted. We found that the salience of cues of a given type (e.g., phonological or semantic), *and* the relative salience of competing phonological and semantic cues affected which cue participants used.

When cues conflicted, participants were more likely to assign class based on high salience semantic or phonological cues (e.g., animacy, or a prefix+suffix combination). Moreover, when conflicting cues differed in salience, learners went with the more salient cue—in general, regardless of whether it was the semantic or phonological cue.

In natural languages, both cue salience and cue reliability typically differ within and across cue types. The results of Experiment 1 suggest that differences along both dimensions may combine to drive patterns of early development in the acquisition of classification systems. However, this also seems unlikely to be the full story; in some reported studies of natural language acquisition, children appear to ignore even very high salience (and high reliability) semantic cues like natural gender or animacy (e.g., Karmiloff-Smith, 1981; Gagliardi & Lidz, 2014). We therefore investigated a further hypothesized mechanism, namely the early availability of phonological cues. In natural language acquisition, there is reason to believe that learners may build representations of categories first based on phonology, before they have acquired the relevant word-to-meaning mappings. While this may eventually be overridden by semantic information (in cases where such information is more reliable, and equally salient), there may be a stage during which phonological cues thus take priority.

In Experiments 2-4 we staged the availability of phonological and semantic cues such that participants were first exposed to input with only one or the other type of cue. For example, in our phonology first conditions, participants were first trained only on nouns with accompanying distributional class markers. Semantic referents were subsequently added in, and thus semantic cues were accessible only later in learning. We found that (in most cases) this had a clear impact on which cue type participants used when the cues conflicted. When cue types were of equal salience, the earlier available cue was more likely to be used. When a weak (low salience) phonological cue was available earlier, the extent to which participants based their classification choices on that cue increased.

It remains to be seen whether the boost we see from early cue availability is a consequence of early availability itself, or whether this is simply a reflection of greater exposure to the early available cue: in Experiments 2-4, learners receive twice as many training trials featuring the early available cue than the late-available cue. It could be that simple volume of data (without staging) would produce the same result, or that early availability contributes beyond simply increasing total volume of exposure. Since our staging experiments were motivated by theories regarding cue availability for real-world learners, where early availability and overall volume of exposure are confounded in exactly the same way, we do not regard this as a defect in our experimental method; nonetheless, our methods could straightforwardly be adapted to disentangle these two possible components of early availability. In any case, there is reason to believe that older children/adults do eventually base their classification decisions on semantic cues, when these are more reliable (e.g., Gagliardi, 2012), suggesting that the initial advantage of phonological cues is eventually balanced out.

Despite the primacy of phonological cues during natural language acquisition, particularly in early acquisition, semantic cues persist in the noun classification system of the world—phonological cues, despite their early dominance, do not take over. Moreover, in many cases the relevant semantic cues continue to be highly reliable predictors of class. One

possible explanation for this is that noun classification systems often start out as purely semantic (Aikhenvald, 2000). Languages may therefore still retain a reflex of these early systems. Alternatively, the maintenance of semantic features as cues to class may be due to the high salience of the features that tend to be involved in noun classification systems in natural languages. The results reported here are consistent with the idea that the cross-linguistically most frequent semantic cues, like animacy and shape, are those which are most easily learned. This suggests an interesting possibility, that later available cues may actually need to be *more* salient and *more* reliable than earlier available cues. Whether this is indeed the case cross-linguistically will require quantitative typological data, as many languages make use of unique, less reliable semantic cues in addition to higher reliability, higher salience cue (e.g., words for trajectories tend to be feminine in French, and words for types of cars tend to be feminine in Italian). Whether we can find evidence for this hypothesis experimentally, we leave for future work.

We have argued here for the role of a particular mechanism—early availability—in explaining why children use phonological cues over semantic cues to noun class. However, it is worth noting again the possibility that children might show an inherent bias against using semantic cues, even if adults do not. The current study cannot rule out this possibility. Along the same lines, it may be that children find particular semantic or phonological cues more or less salient than adults do. Though there is reason to believe that properties like shape and animacy are highly salient for children just as they are for adults (Rakison et al. 2001, Landau et al. 1988), we also know that the role these features play in language undergoes developmental changes (e.g., Smith et al. 2002). Therefore, the next step in this work is to replicate these findings with child learners.

Conclusions

Noun class systems exploit a mix of phonological and semantic cues to class membership, but semantic cues seem to be privileged, both within individual languages (semantic cues tend to be more reliable than phonological cues) and cross-linguistically (a small set of semantic features are used by many languages to determine noun class, e.g., animacy, and no attested languages entirely eschew semantic cues to class). Surprisingly, this favoring of semantic cues is not seen in child learners, who seem to preferentially rely on phonological cues, and select phonological over semantic cues when the two are in conflict.

In this paper we have explored two possible mechanisms which might begin to account for the reliance on phonological cues during acquisition: we explored whether phonological cues are more salient in the context of noun class learning, and whether learners will prefer to use cues which they have early access to (as will be the case if children learn word forms before they learn word meanings). Our results show that cue salience matters: absolute and relative salience of competing phonological and semantic cues affects the likelihood that learners will use them in determining noun class. However, we found no evidence that phonological cues are less salient in general. By contrast, cue availability was revealed as a likely mechanism for explaining children's reliance on phonology. We found that if a cue is available earlier, then it is more often used by learners to determine

categorization. Our results therefore confirm that both salience and availability are likely to play a role in explaining how learners exploit cues to class during acquisition.

What does this work say about the mismatch between the facts from acquisition (phonological cues are preferred) and typology (semantic cues are privileged)? A sensible working hypothesis is that the typological facts reflect biases in acquisition, with more learnable linguistic systems being cross-linguistically more frequent. However, our data show that learnability is not a simple reflection of cue salience (or, by extension, cue reliability): the extent to which learners are able to exploit a given cue depends on its salience, but other factors (such as availability) play a role. This means that the cross-linguistic distribution of noun class systems likely reflects the complex interplay between multiple factors contributing to the ability of learners to exploit phonological and semantic cues during noun class learning.

Acknowledgments

We thank members of the Centre for Language Evolution at the University of Edinburgh as well as audience members at the Linguistic Society of America Annual Meeting, 2015 for their comments and feedback. Funding for this work was provided by the school of Philosophy, Psychology and Language Sciences at the University of Edinburgh.

References

- Adams, K., & Conklin, N. (1973). Toward a theory of natural classification. In C. Corum, T. C. Smith-Stark, & A. Weiser (Eds.), *Proceedings of the Ninth Regional Meeting of the Chicago Linguistics Society* (pp. 1–10). Chicago, IL: Chicago Linguistic Society.
- Aikhenvald, A. (2000). *Classifiers: A typology of noun categorization devices*. New York, NY: Oxford University Press.
- Arnon, I., & Ramscar, M. (2012). Granularity and the acquisition of grammatical gender: How order-of-acquisition affects what gets learned. *Cognition*, 122, 292–305.
- Bates, D. (2010). lme4: Mixed-effects modeling with R. URL <http://lme4.r-forge.r-project.org/book>.
- Bates, E., & MacWhinney, B. (1989). Functionalism and the competition model. In B. MacWhinney, & E. Bates (Eds.), *The crosslinguistic study of sentence processing* (pp. 3–76). New York: Cambridge University Press.
- Becker, M. (2009). The role of NP animacy and expletives in verb learning. *Language Acquisition*, 16, 283–296.
- Braine, M. (1987). What is learned in acquiring word classes: A step toward an acquisition theory. In B. MacWhinney (Ed.), *Mechanisms of language acquisition* (pp. 65–87). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Braine, M. D. S., Brody, R. E., Brooks, P. J., Sudhalter, V., Ross, J. A., Catalano, L., & Fisch, S. M. (1990). Exploring language acquisition in children with a miniature artificial language: Effects of item and pattern frequency, arbitrary subclasses, and correction. *Journal of Memory and Language*, 29, 591–610.
- Brooks, P. J., Braine, M. D., Catalano, L., Brody, R. E., & Sudhalter, V. (1993). Acquisition of gender-like noun subclasses in an artificial language: The contribution of phonological markers to learning. *Journal of Memory and Language*, 32, 76 – 95.

- Bruening, P., Brooks, P., Alfieri, L., Kempe, V., & Dabasinskiene, I. (2012). Children's tolerance of word-form variation. *Child Development Research*, Article ID 401680.
- Bunger, A., & Lidz, J. (2006). Constrained flexibility in the acquisition of causative verbs. In D. Bamman, T. Magnitskaia, & C. Zaller (Eds.), *Proceedings of the 30th Annual Boston University Conference on Language Development* (pp. 60–71). Cambridge, MA: Cascadia Press.
- Carroll, S. E. (1999). Input and SLA: Adults' sensitivity to different sorts of cues to French gender. *Language Learning*, 49, 37.
- Comrie, B. (1989). *Language universals and linguistic typology: Syntax and morphology*. Chicago, IL: University of Chicago press.
- Culbertson, J., & Adger, D. (2014). Language learners privilege structured meaning over surface frequency. *Proceedings of the National Academy of Sciences*, 111, 5842–5847.
- Culbertson, J., Smolensky, P., & Legendre, G. (2012). Learning biases predict a word order universal. *Cognition*, 122, 306–329.
- Culbertson, J., & Wilson, C. (2013). Artificial grammar learning of shape-based noun classification. In M. Knauff, M. Pauen, N. Sebanz, & I. Wachsmuth (Eds.), *Proceedings of the 35th Annual Meeting of the Cognitive Science Society* (pp. 2118–2123). Austin, TX: Cognitive Science Society.
- Culbertson, J., & Wilson, C. (Unpublished results). Rapid learning of semantic noun classification in an artificial grammar. Talk given at the Linguistic Society of America Annual Meeting 88, Minneapolis, Jan. 2–5.
- Demuth, K. (2000). Bantu noun class systems: loanword and acquisition evidence of semantic productivity. *Systems of Nominal Classification*. Cambridge University Press, Cambridge, UK, (pp. 270–292).
- Demuth, K., & Ellis, D. (2008). *Revisiting the acquisition of Sesotho noun class prefixes*. Mahwah, NJ: Lawrence Erlbaum.
- Denny, P. J. (1976). What are noun classifiers good for? In S. S. Mufwene, C. A. Walker, & S. B. Steever (Eds.), *Proceedings of the Twelfth Regional Meeting of the Chicago Linguistics Society* (pp. 122–132). Chicago, IL: Chicago Linguistic Society.
- Dixon, R. (1986). Noun classes and noun classification in typological perspective. In C. G. Craig (Ed.), *Noun classes and categorization* (pp. 105–112). Philadelphia, PA: John Benjamins.
- Frigo, L., & McDonald, J. (1998). Properties of phonological markers that affect the acquisition of gender-like subclasses. *Journal of Memory and Language*, 39, 218–245.
- Gagliardi, A. (2012). *Input and Intake in Language Acquisition*. PhD dissertation, University of Maryland.
- Gagliardi, A., Feldman, N., & Lidz, J. (in press). Modeling statistical insensitivity: Sources of suboptimal behavior. *Cognitive Science*.
- Gagliardi, A., & Lidz, J. (2014). Statistical insensitivity in the acquisition of Tsez noun classes. *Language*, 90, 58–89.
- Gerken, L., & Knight, S. (2015). Infants generalize from just (the right) four words. *Cognition*, 143, 187–192.

- Goldschneider, J. M., & DeKeyser, R. M. (2001). Explaining the “natural order of L2 morpheme acquisition” in English: A meta-analysis of multiple determinants. *Language Learning*, 51, 1–50.
- Gómez, R., & Maye, J. (2005). The developmental trajectory of nonadjacent dependency learning. *Infancy*, 7, 183–206.
- Hawkins, J. A., & Cutler, A. (1988). Psycholinguistic factors in morphological asymmetry. In J. A. Hawkins (Ed.), *Explaining language universals* (pp. 280–317). Oxford: Blackwell.
- Hupp, J. M., Sloutsky, V. M., & Culicover, P. W. (2009). Evidence for a domain-general mechanism underlying the suffixation preference in language. *Language and Cognitive Processes*, 24, 876–909.
- Karmiloff-Smith, A. (1981). *A functional approach to child language: A study of determiners and reference*. New York, NY: Cambridge University Press.
- Kempe, V., & Brooks, P. J. (2001). The role of diminutives in the acquisition of Russian gender: Can elements of child-directed speech aid in learning morphology? *Language Learning*, 51, 221–256.
- Lakoff, G. (1987). *Women, Fire and Dangerous Things—What Categories Reveal about the Mind*. Chicago, IL: University of Chicago Press.
- Landau, B., Smith, L. B., & Jones, S. S. (1988). The importance of shape in early lexical learning. *Cognitive Development*, 3, 299–321.
- Lee, M. (1987). The cognitive basis for classifier systems. In J. Aske, N. Michaelis, & H. Filip (Eds.), *Proceedings of the Thirteenth Annual Meeting of the Berkeley Linguistics Society* (pp. 395–407). Berkeley, CA: Berkeley Linguistics Society.
- Levy, Y. (1988). On the early learning of formal grammatical systems: Evidence from studies of the acquisition of gender and countability. *Journal of Child Language*, 15, 179–187.
- Lyster, R. (2006). Predictability in French gender attribution: A corpus analysis. *Journal of French Language Studies*, 16, 69–92.
- MacWhinney, B. (1978). The acquisition of morphophonology. *Monographs of the society for research in child development*, 43, 1–123.
- Mariscal, S. (2009). Early acquisition of gender agreement in the Spanish noun phrase: starting small. *Journal of Child Language*, 36, 143–171.
- Mills, A. E. (1985). The acquisition of German. In D. Slobin (Ed.), *The crosslinguistic study of language acquisition*, volume 1. Hillsdale, NJ: Lawrence Erlbaum.
- Mulford, R. (1985). Comprehension of Icelandic pronoun gender: semantic versus formal factors. *Journal of Child Language*, 12, 443–453.
- Nelson, D. (2005). French gender assignment revisited. *Word*, 56, 19–38.
- Newport, E. L., & Aslin, R. N. (2004). Learning at a distance I. Statistical learning of non-adjacent dependencies. *Cognitive Psychology*, 48, 127–162.
- Perez-Pereira, M. (1991). The acquisition of gender: What Spanish children tell us. *Journal of Child Language*, 18, 571–590.
- Polinsky, M., & Jackson, D. (1999). Noun classes: Language change and learning. In B. A. Fox, D. Jurafsky, & L. Michaelis (Eds.), *Cognition and Function in Language* (pp. 29–50). Stanford, CA: CSLI.

- Rakison, D. H. and Poulin-Dubois, D. (2001). Developmental origin of the animate–inanimate distinction. *Psychological Bulletin*, 127(2):209.
- Real, F., & Griffiths, T. L. (2009). The evolution of frequency distributions: Relating regularization to inductive biases through iterated learning. *Cognition*, 111, 317–328.
- Rodina, Y., & Westergaard, M. (2012). A cue-based approach to the acquisition of grammatical gender in Russian. *Journal of Child Language*, 39, 1077–1106.
- Senft, G. (2000). *Systems of nominal classification*. Cambridge: Cambridge University Press.
- Slobin, D. I. (2001). Form-function relations: how do children find out what they are? In M. Bowerman, & S. C. Levinson (Eds.), *Language acquisition and conceptual development*, volume 3 (pp. 406–449). Cambridge University Press.
- Smith, L. B., Jones, S. S., Landau, B., Gershkoff-Stowe, L., & Samuelson, L. (2002). Object name learning provides on-the-job training for attention. *Psychological Science*, 13(1):13–19.
- Smith, K., & Wonnacott, E. (2010). Eliminating unpredictable variation through iterated learning. *Cognition*, 116, 44–449.
- Surridge, M. E. (1986). Genre grammatical et derivation lexicale en français. *Canadian Journal of Linguistics*, 31, 267–283.
- Taraban, R. (2004). Drawing learners’ attention to syntactic context aids gender-like category induction. *Journal of Memory and Language*, 51, 202–216.
- Vitevitch, M. S., & Luce, P. A. (2004). A web-based interface to calculate phonotactic probability for words and nonwords in English. *Behavior Research Methods, Instruments, & Computers*, 36, 481–487.
- Williams, J. N. (2004). Implicit learning of form-meaning connections. In J. Williams, B. VanPatten, S. Rott, & M. Overstreet (Eds.), *Form-meaning connections in second language acquisition* (pp. 203–218). Mahwah, NJ: Lawrence Erlbaum Associates.
- Wilson, C. (2006). An experimental and computational study of velar palatalization. *Cognitive Science*, 30, 945–982.

Appendix A

Example Pictures from each semantic category.

Animacy cue.

Animate	Inanimate

Shape cue.

Narrow

Flat

Flexibility cue.

Flexible

Rigid

Appendix B

Words used (Orthographic form and IPA).

Ortho	IPA	Ortho	IPA	Ortho	IPA	Ortho	IPA
barm	bɑːm	foose	fus	proaf	pɹɔf	terk	tɹɪk
blore	blɔɪ	forch	fɔɪtʃ	quire	kwaɪɪ	tief	tɪf
cep	sɛp	forn	fɔɪn	roaf	ɹɔf	tirth	tɹɪθ
cerp	sɹɪp	frag	fɹæɡ	scode	skod	todge	tɔdʒ
chale	tʃeɪl	glait	glɹɪt	shope	ʃɔp	tork	tɔɹk
chount	tʃaunt	glane	gleɪn	skern	skɹɪn	treak	tɹɪk
clade	klɹd	glate	glɛɪt	skun	skɹn	treast	tɹɪst
cleap	klɪp	gleaf	glɪf	sleaf	slɪf	trog	tɹɪɡ
cleem	klɪm	glout	glɹɹt	slock	slɹk	truff	tɹɹf
cleep	klɪp	glurn	glɹɹn	slome	slɔm	vade	vɛɪd
clep	klɛp	greach	ɡɹɪtʃ	slork	slɔɹk	vark	vɹɹk
clom	klɹm	gream	ɡɹɪm	snart	snɹt	veen	vɪn
cloon	klun	groach	ɡɹɔtʃ	sneach	snɪtʃ	verk	vɹɹk
clound	klɹund	grome	ɡɹɔm	snirk	snɹɹk	vide	vɪd
clount	klɹunt	gume	ɡum	snoke	snok	voast	vɔst
clum	klɹm	jeal	dʒɪl	snork	snɔɹk	voke	vok
clurt	klɹɹt	lige	lɪdʒ	spag	spæɡ	voose	vus
clut	klɹt	lour	lɹɹɹ	speat	spɪt	vop	vɔp
crade	kɹɛɪd	mer	mɹɹɹ	splake	spleɹk	vork	vɔɹk
crail	kɹɹɹl	moint	mɔɪnt	splane	spleɪn	vun	vɹn
crɪd	kɹɪd	moke	mok	splear	spleɹ	vune	vun
crin	kɹɪn	mudge	mɹdʒ	splock	splɹk	vurn	vɹɹn
croad	kɹɔd	noaf	nof	splurn	splɹɹn	vurnd	vɹɹnd
crod	kɹɹd	noast	nɔst	spone	spon	vut	vɹt
crom	kɹɹm	nork	nɔɹk	sprock	spɹɹk	yode	yod
cruit	kɹɹt	nuff	nɹf	spuff	spɹf	yome	yom
cust	kɹst	nust	nɹst	sput	spɹt	yort	yɔɹt
dape	deɪp	pam	pɹɹm	starm	stɹɹm	yote	yot
dase	deɪs	pedge	pɛdʒ	steaf	stɪf	zait	zeɪt
deet	dɪt	plice	plɪs	sterk	stɹɹk	zast	zæst
firk	fɹɹk	plird	plɹɹd	stird	stɹɹd	zear	zɪɹ

flaid	flaid	ploak	plok	stirth	stɹɪθ	zipe	zɹɪp
flarp	flɹɹp	plock	plak	stoon	stun	zoat	zot
flear	fliɹ	plound	plaund	stope	stop	zodge	zadz
floud	flaud	podge	padʒ	stug	stʌg	zote	zot
flum	flʌm	porce	pɔɹs	tarce	tɹɹs	zun	zʌn

Appendix C

Exact instructions for the experimental tasks.

Experiment 1.

Introduction: “In this experiment, you will be learning part of a new language. The language has two different ways of saying the word “the”. The word for “the” in this language is either [marker 1] or [marker 2]. Click play to hear them.”

Training: “Now you’ll hear a speaker of the language use the different words for “the”. Look at the picture, and then click the word for “the” that matches what you heard the speaker say. Can you find a pattern for when [marker 1] is used and when [marker 2] is used?”

Testing: “In the next part of the experiment, you will show what you have learned about this new language. Look at the picture, listen to the word, and click on the choice of “the” you think a speaker of the language WOULD BE MOST LIKELY TO SAY. Try to do as well as you can, remembering what you learned in the first part of the experiment, but don’t worry if once in a while you have to guess.”

Experiments 2-4.

Introduction (both conditions): “In this experiment, you will be learning part of a new language. The language has two different ways of saying the word “the”. The word for “the” in this language is either [marker 1] or [marker 2]. Click play to hear them.”

Training 1 (semantics first): “Now you’ll hear a speaker of the language use the different words for “the”. Look at the picture, and then click the word for “the” that matches what you heard the speaker say. Can you find a pattern for when [marker 1] is used and when [marker 2] is used?”

Training 2 (semantics first): “Now you’ll see a picture appear along with the phrase. Look at the picture, listen to the word, and then click the choice of “the” that matches what you heard the speaker say.

Training 1 (phonology first): “Now you’ll hear a speaker of the language use the different words for “the”. Listen to the phrase, and then click the word for “the” that matches what you heard the speaker say. Can you find a pattern for when [marker 1] is used and when [marker 2] is used?”

Training 2 (phonology first): “Now you’ll hear a phrase along with the picture. Look at the picture, listen to the word, and then click the choice of “the” that matches what you heard the speaker say.

Testing (both conditions): “In the next part of the experiment, you will show what you have learned about this new language. Look at the picture, listen to the word, and click on the choice of “the” you think a speaker of the language WOULD BE MOST LIKELY TO SAY. Try to do as well as you can, remembering what you learned in the first part of the experiment, but don’t worry if once in a while you have to guess.”