

Implicative organization facilitates morphological learning

SCOTT SEYFARTH,¹ FARRELL ACKERMAN,¹ & ROBERT MALOUF²

¹*University of California, San Diego*; ²*San Diego State University**

1 Introduction

An essential challenge that faces language-learners is linking wordforms to meanings. In many languages, this requires them to make generalizations across related wordforms that convey different morphosyntactic properties. For example, all of the Russian forms below mean ‘factories’, but a learner must infer the relation between each specific form and its morphosyntactic content.

NOM.	ACC.	GEN.	DAT.	LOC.	INST.	
<i>zavodi</i>	<i>zavodi</i>	<i>zavodov</i>	<i>zavodam</i>	<i>zavodax</i>	<i>zavodam’i</i>	‘factories’

One dimension of this task is segmentation. For example, how can learners separate the stems from the affixes that signal a particular morphosyntactic property? If this information was all that learners had, they might hypothesize that *zavod* has the lexical meaning ‘factory’, and the *-ov* suffix signals genitive plural. A large body of experimental research addresses this syntagmatic, structural challenge of identifying recurrent partial forms (e.g., Saffran et al. 1996; Finley and Newport 2011; Aslin and Newport 2012).

However, there is also a paradigmatic aspect to the problem. For example, the table below shows some alternative possibilities for how a plural form might be realized with different cases in Russian (Baerman et al. 2009).

NOM.	ACC.	GEN.	DAT.	LOC.	INST.	
<i>zavodi</i>	<i>zavodi</i>	<i>zavodov</i>	<i>zavodam</i>	<i>zavodax</i>	<i>zavodam’i</i>	‘factories’
<i>dela</i>	<i>dela</i>	<i>del</i>	<i>delam</i>	<i>delax</i>	<i>delam’i</i>	‘things’
<i>strani</i>	<i>strani</i>	<i>stran</i>	<i>stranam</i>	<i>stranax</i>	<i>stranam’i</i>	‘countries’
<i>kost’i</i>	<i>kost’i</i>	<i>kost’ej</i>	<i>kost’am</i>	<i>kost’ax</i>	<i>kost’am’i</i>	‘bones’

Languages often use different markers to signal the same morphosyntactic distinction, such as a particular case and number. Furthermore, a single marker can be used to signal different properties for different lexical items. If learners focus only on segmentation, they will be faced with sets of affixes that each correspond to multiple meanings, and sets of meanings that correspond to multiple affixes. The task thus involves learning not only segmentation, but also paradigm organization: the word- or subclass-specific correspondences between markers and meanings.

*Thanks to Jeremy Boyd, Colin Bannard, Olivier Bonami, Anne Canter, Melanie Emr, Vsevolod Kapatsinski, Jeanne Monahan, the morphology group at UCSD, and the audience at BLS 40. This project was supported by an NSF Graduate Research Fellowship.

Given this task, a crucial component is to identify the formal pattern that each lexical item uses to signal different properties. For example, in the paradigm above, a language-user must determine if a new word should be realized according to the inflectional patterns of *zavodi*, *dela*, *strani*, *kost'i*, or another pattern entirely.

In identifying a word's case-marking pattern, some forms of the word will be more useful than others. For example, the form *dela* is very useful. If a learner encounters an accusative plural form ending in *-a*, they can be completely certain about the other plural forms of this word: they must follow the *dela* pattern in the second row above. All other accusative plural forms end with *-i*, which means that all other patterns can be eliminated as possibilities if the accusative plural is an *-a* form. An accusative plural *dela* form is thus a perfectly diagnostic form, because this form alone can be used to predict the other forms for this lexical item. There is an implicative relation between this form and other forms of a word: accusative plural *dela* implies genitive plural *del* and nominative plural *dela*.

On the other hand, if a learner encounters one of the locative plural forms ending in *-ax*, that form is unhelpful in identifying the other case-marked forms that a word should take. Every pattern uses *-ax* to signal locative plural. Consequently, none of the locative plurals are diagnostic forms for their patterns.

Implicative relations provide a useful source of information for learners who must identify the formal patterns of new words. Furthermore, the analysis of systems on the basis of such relations is a key component of word-and-paradigm morphology (see Blevins (to appear); Stump and Finkel 2013). Recent work quantifies the extent to which implicative relations reduce the uncertainty associated with predicting a form of a word given another form of the same word, and explores the consequences of this uncertainty for inflectional paradigms (Ackerman et al. 2009; Bonami et al. 2011; Sims 2011). For example, Ackerman and Malouf (2013) show that, in a cross-linguistic variety of complex morphological systems, any given form is much more diagnostic of its correct pattern on-average than one might expect if existing markers had been assigned to their patterns at random. Implicative relations thus significantly reduce the apparent complexity of these systems, and would improve the ability of language-users to learn and generalize morphological paradigms.

1.1 The current study

This study evaluates whether learners use intra-paradigm implicative relations to identify the other forms that a word takes. If a learner knows the relevant paradigm and encounters a diagnostic form of a word, they should be able to infer the other forms that it implies (see also Brooks et al. 1993; Frigo and MacDonald 1998).

This hypothesis is tested using two artificial grammar experiments, in which experimental participants first learned an artificial number-marking paradigm through trial-and-error. After exposure to the the full paradigm, participants were presented with new wordforms that they had not previously encountered, and were asked to generate related forms of the same word. The prediction is that if participants are first presented with a diagnostic form, they should be able to correctly produce a related form that is implied by the diagnostic form. On the other hand, if the form that they first encounter is not diagnostic, they should be unable to reliably produce a correct related form.

2 Experiment 1

2.1 Methods

2.1.1 Participants

39 individuals were recruited from the UC San Diego community, and received partial undergraduate course credit in exchange for participation. Participants were between 18–33 years old (mean 21). 35 reported having excellent, native, or native-like knowledge of English; and 23 reported having excellent, native, or native-like knowledge of a non-English language (including 7 Spanish and 4 Korean speakers).

2.1.2 Procedure

The experiment had a training phase and a testing phase. In the training phase, participants were asked to learn how to mark a variety of familiar household objects for number (singular, dual, and plural) in an alien language. In the testing phase, participants were asked to produce number-marked forms of household objects that they had not previously been trained on. Participants were seated at a computer in a quiet room. All instructions and trials were shown on the screen using the PsychoPy experiment presentation software (Peirce 2007), and participants typed free written responses with the keyboard.

2.1.3 Number paradigm

Items Word stems were the English singular forms of common household objects, such as *table*, *chair*, and *bed*. Five nonce syllables from the ARC Nonword Database (Rastle et al. 2002) were assigned to be inflectional markers for singular, dual, and plural wordforms in two subclasses. For example, the singular marker for Subclass 1 might be *yez*. A stem was inflected for number by concatenating the appropriate marker to the end. For example, the inflected singular form of *table* might be *table-yez*, or the inflected dual form of *chair* might be *chair-guk*.

As mentioned, the number-marking paradigm had two subclasses. For two numbers, the subclass of the stem determined which of two markers (allomorphs) should be used. The third number was always indicated with the same marker, regardless of the stem’s subclass. Table 1 shows a sample realization of the number-marking paradigm used in the experiment. In this version, singular and dual forms have different markers depending on the subclass, whereas the plural form is marked with *-lem* in both subclasses.

The number-marking paradigm was designed to have a structure so that only some markers reliably predicted a referent’s subclass membership. In other words, only some number markers could be used to predict the other markers that a given stem would be inflected with. The hypothesis is that when participants encounter a new word that has a number-marker that predicts subclass membership, they should be able to produce the correct number-markers for that word in its other forms. When participants encounter a new word with a number-marker that does not predict its subclass, they should have only a 50% chance of correctly guessing another marker.

	SINGULAR	DUAL	PLURAL
<i>Subclass 1</i>	table-yez	table-cav	table-lem
<i>Subclass 2</i>	chair-taf	chair-guk	chair-lem

Table 1: Sample number-marking paradigm that participants saw in Experiment 1. Objects with English names were randomly assigned to each subclass of markers.

Randomization There were 30 possible objects that were used in the experiment. For each participant, these objects were randomly divided into two subclasses which comprised 15 stems each. During the training phase, each participant learned the correct inflections for a random subset of six of the 30 stems, with three per subclass.

For each participant, the five number-markers were randomly assigned to the cells, although the structure of the paradigm was always the same. For example, one participant might see *-yez* and *-taf* as the singular markers, with *-lem* as the plural marker, as in the sample paradigm. Another participant might see *-lem* and *-taf* as the singular markers, and *-yez* as the plural marker. Each participant thus saw a different realization of the paradigm.

As noted above, two numbers had unique markers for each subclass, and one number had the same marker for both subclasses. The number that had the same inflectional marker in both subclasses was randomly selected for each participant. For example, unlike the sample realization in table 1, a participant might encounter a version in which all of the stems take the same singular marker, but the plural form has a different marker depending on the subclass that the stem belongs to.

Implicative organization Some markers were fully diagnostic of subclass membership once a participant had learned the paradigm structure. For example, in the sample realization, if a participant learns that *table* is written as *table-yez* when *table* is singular, they can predict that it should be written as *table-cav* when *table* is dual. The *-yez* marker only appears in Subclass 1, so the participant can infer that if a stem takes the *-yez* marker, it must belong to Subclass 1. Therefore, it should take the other markers that belong to Subclass 1. A marker that can be used to predict subclass membership of a stem is thus a diagnostic form.

Note that, to make this generalization, it is not necessary to infer the existence of subclasses per se and memorize their markers. If a participant knows that all of the stems that are marked with *-yez* in the singular are also marked with *-cav* in the dual, then they can reasonably infer that a new word that takes *-yez* in the singular should also take *-cav* in the dual. There is thus a strong implicative relation where $-yez \rightarrow -cav$.

On the other hand, the marker that appears in both subclasses cannot be used to diagnose subclass membership. If a participant learns that *shoe* is written as *shoe-lem* in the plural, they cannot use that fact to determine whether *shoe* belongs to Subclass 1 or 2, since *-lem* is used to mark plural number in both subclasses. A participant will have observed that every stem takes the *-lem* marker in the plural, and therefore there is no implicative relation in which *-lem* predicts either of the singular or dual markers.

2.1.4 Instructions

Participants were asked to learn, by trial-and-error, how a group of aliens referred to different numbers of objects in their language. They were given an example about English *balloons* and *geese* to sensitize them to the possibility of different number-marking patterns, and were told that two balloons might be referred to as *a pair of balloons* (in contrast with just *balloons*) to prepare them for the dual-plural distinction that is not inflected in English.

2.1.5 Training phase

In the training phase, participants learned how to add markers to stems to inflect them for singular, dual, or plural number.

In each trial of the training phase, participants saw one, two, or many black-and-white line drawings of one of the six training objects. They were asked to tell the aliens what they saw on the screen. For example, they might see two images of a chair on the screen. With the paradigm realization in table 1, a correct response was ‘chair-guk’, and an incorrect response was anything else. Participants were only able to type lower-case letters and hyphens, and pressed the enter key to indicate that they had completed their response.

After giving a response, participants were immediately told whether they were correct or not. If they answered correctly, they also saw a smiling cartoon face for two seconds. Participants then saw the image again along with the correct answer, which appeared for six seconds if they answered incorrectly, or for only three seconds if they answered correctly. They also received 100 points for each correct answer, and the total number of points earned was displayed in the top right throughout the training phase.

Each training object appeared with each number (singular, dual, or plural) in one trial per training block. The order of trials was randomized within each block. There were five training blocks, which comprised 90 training trials in total (6 training stems * 3 numbers * 5 blocks). Participants were given a break between each block for as long as they wanted, although most chose to continue the experiment immediately.

2.1.6 Testing phase

Once they had completed the training, participants saw a second set of instructions. They were told that the aliens wanted to see how they could do on new objects that they’d never seen before. In each trial of the testing phase, participants saw one of the remaining 24 objects that they had not seen during training. On the left side of the screen, they saw one, two, or many line-drawings of the object (e.g., a phone). They were shown the correct number-marked form for that stem (e.g., ‘phone-guk’), and told that this form (the GIVEN FORM) is what the aliens would say to refer to the object(s) on the left side of the screen.

After a three second pause, a different number of line-drawings of the same object appeared on the right side of the screen. Participants were asked to tell the aliens what they saw on the right side (the TARGET FORM). They were not explicitly told to use the given form to predict the target form. However, in the instructions before the testing phase, they were encouraged to think about what they had learned if they were not sure of the correct response.




Figure 1: Mean accuracies at guessing the correct target form after being presented with a diagnostic or non-diagnostic given form in each testing trial.

Each of the 6 possible relations (e.g., given form is singular, target form is dual; given form is singular, target form is plural; etc.) was tested for both subclasses. Tests appeared in a random order within each block, and there were two testing blocks. Participants saw each of the 24 testing objects on exactly one testing trial (6 relations * 2 subclasses * 2 blocks).

2.2 Analysis

A mixed-effects logistic regression was used to model whether or not participants produced a correct target form after seeing a particular given form during the testing phase. The variable of interest was whether or not the given form implied the target form (IMPLICATIVE-RELATION). The prediction was that if the given form implied the target form, participants should produce more correct targets than if the given form did not imply the target form. For example, in table 1, participants should be able to guess that *lamp-cav* is the dual form if they are given *lamp-yez* as the singular form. They should not be able to guess that *lamp-cav* is the dual form if they are given *lamp-lem* as the plural form, since *-lem* does not allow participants to identify the subclass that *lamp* belongs to.

It is trivially true that each of the singular and dual forms in table 1 imply the plural form, since *-lem* must always be the plural form. Therefore, trials in which there was only one possible target marker (because it was used by both subclasses) were excluded from the analysis. In fact, participants produced the correct form in over 98% of such trials. A chart showing the relative accuracy of participants under each condition appears in figure 1.

	SD	Corr.
SUBJECT (INTERCEPT)	0.24	—
IMPL.-RELATION	2.41	0.19

Table 2: Random effects summary for Experiment 1.

2.3 Results

As predicted, IMPLICATIVE-RELATION was highly significant, and had a large effect on the accuracy of responses ($\beta = 2.20$, $z = 4.77$, $p < 0.0001$). Participants guessed target forms significantly more often when they first encountered a wordform that could be used to predict subclass membership. A summary of the random effects appears in table 2.

Additional analyses indicated that participants performed equally well on dual forms as on singular and plural (even though dual is not distinctly inflected in English), and that the regression was not significantly improved by adding random or fixed parameters for stems, markers, or forms. Participants who reported native, native-like, or excellent proficiency in a language with regular morphological subclasses (such as noun gender) did not perform significantly better or worse in the testing phase than other participants.

3 Experiment 2

The first experiment had a very simple structure. There were two symmetrical subclasses, and implicative relations always went in both directions: if $-yez$ predicted $-cav$, then $-cav$ also predicted $-yez$. The structure was as simple as possible for the design.

Natural-language paradigms often have much more complex structures. A second experiment was conducted with two asymmetrical paradigms in order to evaluate whether participants would be still able to take advantage of implicative relations when the structure is more complex.

3.1 Methods

The methods of Experiment 2 were the same as Experiment 1, except as described here.

3.1.1 Participants

60 individuals were recruited from the UC San Diego community, and received partial undergraduate course credit in exchange for participation. Participants were between 18–25 years old (mean 21). All participants reported having excellent, native, or native-like knowledge of English; and 42 reported having excellent, native, or native-like knowledge of a non-English language (including 13 Spanish, 10 Mandarin, and 5 Vietnamese speakers). 32 participants were randomly assigned to learn the SINGLE PRINCIPAL PART paradigm, and 28 were assigned to the ASYMMETRICAL paradigm, described below.

3.1.2 Procedure

In this experiment, there were 27 common English household objects that were used as stems. Nine objects were randomly assigned to each of three subclasses. During training, participants learned the correct inflections for three stems per subclass. The remaining 18 stems were held in reserve for the testing phase. In the testing phase, each of the 6 possible relations was tested for the three subclasses in one trial each, using each of the stems that the participant had not seen during training.

3.1.3 Number paradigms

Two paradigm structures were designed to evaluate whether it was easier for participants to identify and take advantage of diagnostic forms when the markers for one number could consistently be used to predict other forms. The SINGLE PRINCIPAL PART paradigm was designed so that, regardless of subclass, the forms for one number (e.g., singular) were always diagnostic of other forms and of subclass membership. In the ASYMMETRICAL paradigm, this was not the case.

A sample realization of the Single Principal Part paradigm structure appears in 3a, and a sample realization with the Asymmetrical paradigm appears in 3b. In 3a, the singular number-marker in each subclass is a reliable predictor of other forms. In 3b, there is no number where each marker is reliably diagnostic in every subclass.

Randomization of stems, markers, and numbers was done as described in section 2.1.3. Thus, each participant saw a different realization of one of the two paradigms, but as before, the paradigm structure was always the same.

There were two features of these paradigms that could help participants choose the appropriate marker for a specific number and stem: *marker frequency* and *implicative relations*.

Marker frequency In these paradigms, some markers are more frequent than others. For example, *-cav* marks dual in two of the three subclasses (in both paradigms), whereas *-guk* marks dual in only one subclass.

Implicative organization As before, some markers implied other markers. For example, if a participant who has learned the Single Principal Part sample paradigm knows that a wordform in the singular is marked with *-taf*, they can be sure that the same wordform is marked with *-guk* in the dual and *-lem* in the plural. The marker *-taf* only occurs in Subclass 1, so knowing that a word takes *-taf* is enough to determine that that word belongs to Subclass 1 and therefore must take the other markers that are used in Subclass 1.

On the other hand, if they only know that a wordform in the plural is marked with *-lem*, they cannot diagnose that word's subclass membership. They do know that *-lem* is not the plural marker in Subclass 3 (that is *-nup*), but they do not know whether a wordform marked with *-lem* belongs to Subclass 1 or 2, because *-lem* occurs in both subclasses. Thus, some markers are diagnostic of subclass membership and other markers, and some are not.

	SINGULAR	DUAL	PLURAL		SINGULAR	DUAL	PLURAL
<i>Subclass 1</i>	chair-taf	chair-guk	chair-lem	<i>Subclass 1</i>	chair-taf	chair-guk	chair-lem
<i>Subclass 2</i>	bed-yez	bed-cav	bed-lem	<i>Subclass 2</i>	bed-yez	bed-cav	bed-lem
<i>Subclass 3</i>	table-seb	table-cav	table-nup	<i>Subclass 3</i>	table-yez	table-cav	table-nup

(a) Sample Single Principal Part paradigm.

(b) Sample Asymmetrical paradigm.

Table 3: Participants saw a paradigm realization with either the Single Principal Part or the Asymmetrical structure in Experiment 2.

3.2 Analysis

A mixed-effects logistic regression was used to model whether or not participants produced a correct target form after seeing a particular given form of a word during the testing phase. As before, the variable of interest was whether or not the given form implied the target form. The primary hypothesis is that participants should guess the correct target form more often during the testing phase when the given form implies the target form.

For this analysis, an additional variable was included in the regression: whether the target form was a frequent marker (such as *-cav* in the sample paradigm) or whether it was not. The expectation is that participants should be biased toward guessing more frequent markers, all else equal, and so should guess correctly more often when the target form is a more frequent marker.

The regression also included a variable for paradigm condition (SINGLE PRINCIPAL PART or ASYMMETRICAL) and pairwise interactions (plus the three-way interaction) between each of the variables. Finally, the model included random per-subject intercepts and slopes for each fixed effect.

3.3 Results

Paradigm condition, and all of the interactions except for the IMPLICATIVE-RELATION x FREQUENCY interaction, were found to be non-significant. Therefore, these fixed effects and their corresponding random effects were removed from the model. This did not affect the significance of the remaining variables.

The final fixed effects summary appears in table 4, and a summary of the random effects appears in table 5. As before, IMPLICATIVE-RELATION had a significant positive effect on the probability of correctly guessing a target form. The effect size was much smaller than in the simpler paradigm in Experiment 1.

Two other effects were significant. First, participants were much more likely to guess frequent markers. Second, there was a significant negative interaction for having a diagnostic form when the target form was a frequent marker. The negative effect size indicates that when the given form implies the target form but the target form is frequent, participants failed to guess the target form as well as they should have.

Thus, participants did significantly better when the given form implies the target form, but only when the given form predicts one of the infrequent markers. To verify this interpretation, two alternative models were fit only to trials in which the target was frequent or




Figure 2: Mean accuracies at guessing the correct target form after being presented with a diagnostic or non-diagnostic given form in each testing trial.

	Estimate	SE	z	p
IMPLICATIVE-RELATION	0.87	0.32	2.72	0.0066
FREQUENCY	2.11	0.38	5.55	< 0.0001
IMPLICATIVE-RELATION X FREQUENCY	-1.42	0.50	-2.86	0.0043

Table 4: Fixed effects summary for Experiment 2.

infrequent. These models confirmed that the effect of IMPLICATIVE-RELATION was significant when the target form was an infrequent marker ($\beta = 0.84$, $z = 2.75$, $p < 0.01$) but was non-significant when the target form was a frequent marker ($\beta = -0.55$, $z = -1.81$, $p > 0.05$). The chart in figure 2 shows the relative accuracy of participants at guessing the correct target form under different conditions.

Finally, to verify that this apparent effect would not have arisen from a simpler strategy by the participants (Lignos 2013), 100 simulations of the experiment were conducted under the assumption that participants were probability-matching (i.e., selecting the more frequent of the two possible number-markers 2/3 of the time, and the less frequent one 1/3 of the time) rather than using diagnostic forms. Under this assumption, IMPLICATIVE-RELATION and its interaction with frequency were not significant more often than expected by chance, which suggests that the significant effect shown in figure 2 would not have occurred under this simpler strategy.

	SD	Corr.
SUBJECT (INTERCEPT)	0.88	—
IMPL.-RELATION	1.78	-0.49
FREQUENCY	2.31	-0.99
IMPL.-R. x FREQ.	3.00	0.86

Table 5: Random effects summary for Experiment 2.

4 Discussion

An implicative relation between wordforms exists when one form of a word allows the language-user to determine another form of the word that is marked for a different morphosyntactic property. The main result of the experiments presented here is that artificial language-learners were able to identify and take advantage of implicative relations between number-markers.

In Experiment 1, participants were more successful at guessing a second form of a word if they first encountered a diagnostic form that implied the target form, as compared to when they encountered a non-diagnostic form. In Experiment 2, this was also true, but only when a given diagnostic form implied that the correct form was one of the infrequent forms.

4.1 Accounting for the differences between frequent and infrequent targets

As noted in section 3.1.3, participants had two sources of information in Experiment 2 to help them choose the correct marker. In some cases, knowledge of one form provides enough information to predict the required form. For example, if the participant knows that a word takes the *-taf* marker in the singular, then that implies that the word must take the *-guk* marker in the plural.

However, even when not given a diagnostic form, learners might make use of their knowledge of relative marker frequency. As a default strategy, if a learner knows nothing else about the word, it is reasonable to guess that the appropriate marker is the more frequent of the two possibilities. For example, if a learner had to guess a dual marker for a word that they had never seen before, they should guess that the word takes *-cav*, just because it is a more common form.

One might expect learners to follow the first strategy—the use of predictive relations—when it is available to them, and to only fall back on relative marker frequencies when they are not given a helpful diagnostic form. However, the participants’ strategy was not optimal in this respect.

The results of Experiment 2 indicate that when provided with a diagnostic form that predicts one of the infrequent markers (like *-taf* or *-guk*), participants did choose the infrequent target significantly more often than they would otherwise. In this case, they did take advantage of the useful information that a given wordform provides about another form of that word. They did not perform as well as they could have, but this strategy did significantly improve their accuracy.

On the other hand, when given a diagnostic form that predicts a frequent marker, the participants did not behave in the same way. In this case—when the diagnostic form implies a frequent marker—participants were just as likely to choose the frequent marker as if they had no implicative information at all. They did not take advantage of the helpful relation if it implied a frequent marker that they were already likely to guess anyway.

In this experiment, the participants may have learned a strong prior preference for the frequent marker—their default guess is the more common marker. When they receive extra information from a diagnostic form that indicates that they should actually select the infrequent marker, they pay attention to that information. On the other hand, if they get extra information that confirms their expectation for the frequent marker, that does not make a large difference in their preferences.

In fact, this pattern can be seen in Experiment 1 as well. If learners use only the implicative relations and ignore other sources of information, they should be perfectly accurate when they encounter a diagnostic form. However, their accuracy was actually closer to 75% on such trials (see figure 1). One way of describing this behavior is that the participants learned to have an equal prior expectation for each number-marker, since the markers were equally frequent. When they get a predictive form, they adjust their expectation by some amount to favor the predicted marker, but not as much as they would if they were putting complete faith in the reliable implicative relations.

As a second possible account for the results of Experiment 2, Yurovsky et al. (2013) demonstrate that adults suffer from a *dilution effect* when they are presented with two sources of evidence that point toward the same outcome, but one piece of evidence is much stronger than the other. In this scenario, infants behave as though they have only the single stronger piece of evidence. However, the same situation actually weakens adults' confidence in the outcome: they average the strength of the two pieces of evidence, and behave as though the total evidence is not as strong as the stronger piece of evidence alone.

This phenomenon might describe the results of Experiment 2. When confronted with a stronger and a weaker source of evidence (implication and frequency, although it is not necessarily clear which is stronger) that both point toward the same (frequent) target, participants are less confident in that target than if they had only one source of evidence. Since infants do not suffer from the dilution effect, the results of Experiment 2 might not be the same if it were possible for infants to take the experiment.

Finally, a third possible account might be that participants were learning a different kind of relation. In Experiment 2, they may have learned to pair suffixes together (e.g., *-taf* goes with *-guk*). They might have avoided suffix pairs if one of the pair members also occurs with another suffix (e.g., *-cav* goes with *-lem*, but also with *-nup*). In other words, they only acquired bi-directional implicative relations. Future experiments should test this hypothesis.

4.2 Relation to earlier work on artificial subclass acquisition

Previous experimental work has argued that learners are unable to use diagnostic forms to predict subclass membership, unless the learned stems have redundant phonological or semantic cues to subclass membership (Brooks et al. 1993; Frigo and MacDonald 1998; Gerken et al. 2009). The results of the current study conflict with this claim. In both experiments, participants demonstrated knowledge of formal paradigmatic relationships, despite experi-

mental randomization that was designed to avoid redundant information that would signal subclass membership.

Some researchers have argued that enriching word classes with alternative kinds of redundancies, such as an association between each subclass and the semantics of the items in that subclass (Braine 1987; Ouyang et al. 2012), can allow learners to acquire morphological subclasses or relations within subclasses. There is also evidence that learners successfully make paradigmatic inferences in artificial syntactic paradigms that lack redundant cues (Mintz 2002; Reeder et al. 2013). Together with the present results, these findings support a view in which (short-term) paradigm learning is enabled simply whenever memory demands are sufficiently low (Frank and Gibson 2011). Previous work (e.g., Frigo and MacDonald 1998; Ouyang et al. 2012) lowered memory demands through additional phonological or semantic coherence in subclasses.

In the current experiments, a few methodological factors kept the memory demands low. First, the learning target was very narrow—only the new suffixes and their relations had to be learned, and there were relatively few items per subclass during training. Second, the task involved inferring relations among actual referents, rather than purely among abstract linguistic labels. These referents were presented prior to their labels, and this presentation style is known to enhance association learning (Ramscar et al. 2010). Furthermore, participants were required to actively produce a free response to every item, from the very start of the task. Participants received two kinds of affective feedback following successful responses. Many participants reported that they enjoyed the task and the challenge of learning a new language.

4.3 Summary

This study investigated whether learners use implicative relations to identify the other forms that a word takes. Evidence from two experiments suggest that artificial paradigm organization is salient to morphological learners, and that they do use intra-paradigm relations to correctly infer related wordforms. This supports typological research that highlights the importance of implicative relations in the formal structure of morphological systems. Future work should investigate the interaction of this kind of structure with the relative frequencies of markers and subclasses, as well as the different strategies that L1 and L2 language-learners use to acquire and represent implicative relations.

References

- Ackerman, F., Blevins, J., and Malouf, R. (2009). Parts and wholes: Implicative patterns in inflectional paradigms. In Blevins, J. and Blevins, J., editors, *Analogy in Grammar: Form and Acquisition*, pages 54–82.
- Ackerman, F. and Malouf, R. (2013). Morphological organization: The low conditional entropy conjecture. *Language*, 89(3):429–464.
- Aslin, R. N. and Newport, E. L. (2012). Statistical learning: From acquiring specific items to forming general rules. *Current Directions in Psychological Science*, 21(3):170–176.

- Baerman, M., Brown, D., and Corbett, G. G. (2009). Morphological complexity: A typological perspective. In *European Science Foundation Exploratory Workshop on 'Words in Action'*, Pisa. Istituto di Linguistica Computazionale.
- Blevins, J. P. (to appear). *Word and Paradigm Morphology*. Oxford University Press, Clarendon.
- Bonami, O., Boyé, G., and Henri, F. (2011). Measuring inflectional complexity: French and Mauritian. In *Workshop on Quantitative Measures in Morphology and Morphological Development*, San Diego.
- Braine, M. (1987). What is learned in acquiring word classes: a step toward an acquisition theory. In *Mechanisms of Language Acquisition*, pages 65–87. Lawrence Erlbaum, Hillsdale, NJ.
- Brooks, P., Braine, M., Catalano, L., and Brody, R. (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.
- Finley, S. and Newport, E. L. (2011). Morpheme segmentation in school-aged children. *University of Rochester Working Papers in the Language Sciences*.
- Frank, M. C. and Gibson, E. (2011). Overcoming memory limitations in rule learning. *Language Learning and Development*, 7(2):130–148.
- Frigo, L. and MacDonald, J. (1998). Properties of phonological markers that affect the acquisition of gender-like subclasses. *Journal of Memory and Language*, 39:218–245.
- Gerken, L., Wilson, R., Gómez, R. L., and Nurmsoo, E. (2009). The relation between linguistic analogies and lexical categories. In *Analogy in Grammar: Form and Acquisition*. Oxford University Press.
- Lignos, C. (2013). You can't get there from here: On interpreting learning experiments. *University of Pennsylvania Working Papers in Linguistics*, 19(1):97–106.
- Mintz, T. H. (2002). Category induction from distributional cues in an artificial language. *Memory & Cognition*, 30(5):678–686.
- Ouyang, L., Boroditsky, L., and Frank, M. C. (2012). Semantic coherence facilitates distributional learning of word meanings. In *Proceedings of the 34th Annual Conference of the Cognitive Science Society*.
- Peirce, J. W. (2007). PsychoPy—Psychophysics software in Python. *Journal of Neuroscience Methods*, 162(1-2):8–13.
- Ramscar, M., Yarlett, D., Dye, M., Denny, K., and Thorpe, K. (2010). The effects of feature-label-order and their implications for symbolic learning. *Cognitive Science*, 34(6):909–957.
- Rastle, K., Harrington, J., and Coltheart, M. (2002). 358,534 nonwords: The ARC nonword database. *Quarterly Journal of Experimental Psychology*, 55A:1339–1362.

- Reeder, P. A., Newport, E. L., and Aslin, R. N. (2013). From shared contexts to syntactic categories: The role of distributional information in learning linguistic form-classes. *Cognitive Psychology*, 66(1):30–54.
- Saffran, J. R., Aslin, R. N., and Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, 274(5294):1926–1928.
- Sims, A. D. (2011). Information theory and paradigmatic morphology. In *Information-theoretic Approaches to Linguistics Workshop*, Boulder, Colorado.
- Stump, G. and Finkel, R. (2013). *Morphological Typology: From Word to Paradigm*. Cambridge University Press.
- Yurovsky, D., Boyer, T. W., Smith, L. B., and Yu, C. (2013). Probabilistic cue combination: less is more. *Developmental Science*, 16(2):149–158.