

Logical Complexity in Morphological Learning: effects of structure and null/overt affixation on learning paradigms

Katya Pertsova *

University of North Carolina at Chapel Hill

1. Introduction

Language learning is to a large extent learning how to classify objects based on their properties (features). For instance, learning a morphological paradigm can be viewed as learning conditions for affix insertion, where the distribution of each affix is determined by morpho-syntactic features. Sometimes the distribution of an affix can be described simply in terms of a single conjunction of features (e.g., use */z/* in [3rd person singular] contexts), but sometimes, due to syncretism, an affix has a heterogeneous distribution that is impossible to state as a simple conjunction of features. (For example, the form *are* of the verb “to be” occurs either in a plural context or in a 2p. singular context.) A natural hypothesis is that paradigms with simpler affix distributions should be acquired faster, with less errors, and, therefore, be less prone to historical change. However, what is the relevant metric of simplicity (or complexity) for the human learners? We address this question by way of artificial grammar learning experiments that provide a controlled setting for studying what factors affect pattern complexity. Similar types of experiments have a long history in psychology. In particular, there is an extensive literature on learning of artificial categories defined by visual features such as shape, color, size, and so on. The robust findings in this literature can serve as a starting point for

*I am grateful to Elliott Moreton for many suggestions related to this work. This project was in part supported by the University of North Carolina’s Research Council Small Grant Program.

identifying what linguistic patterns are more complex than others. Namely, we can test whether the results found for non-linguistic patterns extend to the linguistic domain. This line of inquiry not only lets us investigate complexity of linguistic pattern learning but also provides a way of addressing the hotly debated question of language-specialization: are there different learning mechanisms for learning of comparable linguistic vs. non-linguistic patterns?

This paper considers a morphological analogue of a psychological experiment on learning categories with the structure of Boolean connectives AND, OR, and XOR (exclusive OR). This experiment will test the effects of logical structure and the presence of null/overt affixes on paradigm complexity. The results are partially consistent with the previous psychological studies, however, they also differ in several ways from analogous non-linguistic experiments. At this point, it is unclear exactly what factors are responsible for these differences.

2. Complexity of Boolean categories

Some of the simplest types of categories are those that can be described with two binary features. There are four such non-isomorphic categories that have the logical structure of *affirmation* (AFF), a category that is defined by a presence or absence of a single feature-value, *conjunction* (AND), defined by a conjunction of two feature-values, *disjunction* (OR), defined by a disjunction of two feature-values, and *exclusive disjunction* (XOR), a category that is defined by a disjunction of two conjunctions. These four types are illustrated below in the order of their learning difficulty as found in numerous psychological studies on acquisition of artificial categories (Bruner *et al.* 1956, Neisser and Weene 1962, Gottwald 1971).

Table 1: Categories over two binary features: shape (triangle/circle) and color (black/white). Boxes pick out the examples that belong to the category.

(a) AFF circle	(b) AND circle AND black	(c) OR triangle OR white	(d) XOR/ \leftrightarrow (black AND triangle) OR (white AND circle)

Evidence from linguistic typology (see below) is consistent with the above complexity order suggesting that this order might be universal across domains. Some such evidence comes from recent typological studies into the syncretism of person/number marking (Cysouw 2003, Baerman and Brown 2005, Baerman *et al.* 2002). According to the data from the World Atlas of Language Structures on verbal subject agreement markers, 80 out of 140 languages in the sample that mark agreement have no syncretism (Baerman and Brown 2005). That is, in *at least 57%*

Logical Complexity in Morphological Learning

of languages, the distribution of every subject agreement affix can be described as a conjunction of morphosyntactic features. This number is just a lower bound since languages with syncretism can also have paradigms in which every affix is describable with a conjunction of features. This would happen if the syncretic cells formed a natural class. To get an estimate of how often that happens, Pertsova (2007) examined 88 subject agreement paradigms from 30 languages with syncretic verbal agreement markers. She found that 51% of these paradigms contained only those instances of syncretism that could be described either as a conjunction of several features or an affirmation (i.e., the syncretic cells formed a natural class). She also showed that the number of paradigms that would be expected to have at least one non-conjunctive affix by chance grows very fast with the size of a paradigm. For example, there are 4,140 possible partitions of an 8-cell paradigm over 3 binary features. Only 3% of these partitions are expected to involve exclusively conjunctive affixes. So, taken together these facts suggest that paradigms in which every affix is conjunctive (or, alternatively, has a 1-1 correspondence between form and meaning) occur more often than expected by chance: more than half of verbal agreement paradigms are 1-1.

Cysouw (2003) also looked at frequencies of different types of person/number syncretisms in personal pronouns and in verbal agreement markers. His sample was not fixed ahead of time and was intended to include as many different patterns of syncretism as possible (and, as a result, it was skewed towards the rarer patterns). The rarest types of syncretism in this sample were cases in which the syncretic cells did not form a natural class (1st sg.=2nd pl; 3rd sg.=2nd pl; 2nd sg.=1st pl), and among the most frequent types of syncretism were cases in which all singular or all non-singular persons were syncretic and cases in which either 1st and 2nd persons (participants in the discourse) or 2nd and 3rd persons (non speakers) were syncretic. This again suggests that syncretisms that can be described in terms of a simple conjunction of feature values are relatively more frequent than other types of syncretism.¹

Similar findings have been reported with respect to phonological patterns: phonologically active classes typically involve a set of segments that can be described as a conjunction of feature values and rarer as a disjunction of a small number of such conjunctions (Mielke 2004). Moreton analyzed patterns from Mielke's database (P-base1.93, Mielke (2008)) to estimate exactly how many classes of phonological segments fall into one of the four patterns in table 1 (Moreton 2012). His results based on 2034 vowel classes and 5682 consonant classes (from 627 languages) support the ranking of the four categories reported above. Namely, conjunctions and

¹ These results of course depend on the kinds of features one adopts, and linguists commonly choose feature systems that make the description of linguistic patterns, including syncretism, simple. Nevertheless, the choice of features is also motivated by other factors such as semantics, facts about agreement, allomorphy, and so on. If a host of different factors all point to the same grouping of morpho-syntactic categories across different languages, we can be more sure that the simplicity effects we find in typological data are real.

affirmations are significantly overrepresented in the sample compared to chance, while XOR and OR are either underrepresented or similar to what one would expect by chance.

Based on the facts reported in this section, there are reasons to expect that learning artificial morphological paradigms in the lab would produce the same complexity ranking found in non-linguistic experiments and in linguistic typology. This result would not only strengthen our belief in the domain-independent nature of the ranking, but also suggest that the learning strategies subjects use in the lab are not vastly different from the strategies used in natural language acquisition (that is, strategies that affect typology).

3. Learning categories vs. learning paradigms

Notice, that in table 1 categories AND and OR are complements of each other and, therefore, are part of the same pattern. (In the rest of the paper, the term *pattern* refers to a partition of the feature space into several categories). Indeed, in terms of morphological paradigms there is a single AND/OR pattern: a paradigm in which one morph corresponds to the conjunctive category, and the other – to the disjunctive category. The morphological analogues of the three logically possible patterns are shown below.

Table 2: Syncretism in paradigms with two binary features

Hebrew V pres. subj. agreement			English past forms of <i>do</i>			Somali def. articles		
affirmation			AND/OR			XOR/↔		
	sg.	pl.		sg	pl		sg	pl
–part fem.	-et	-ot	–part	does	do	m.	-ta	-ka
+part fem.	-et	-ot	+part	do	do	f.	-ka	-ta

A few psychological studies (Peters and Denny 1971, Gottwald 1971) have investigated whether the difficulty of learning a pattern is affected by whether the same task is presented as (a) learning a distinction between examples that belong or do not belong to a category, or (b) learning the same pattern a distinction between two complementary categories A and B (similar to learning the distribution of two distinct affixes in a paradigm). Most relevantly for this paper, Gottwald (1971) considered learning of the four categories in table 1 in two conditions: *biased* labeling, in which instances were labeled as described in (a) above, and *neutral* labeling as described in (b) above. The stimuli were schematic bugs, with two relevant features being body shape and spot color. The measure of learning difficulty was the number of trials it took a subject to correctly categorize 16 instances (pictures of bugs) in a row.

Logical Complexity in Morphological Learning

Interestingly, the type of labeling showed an interaction with logical structure so that neutral labeling condition was harder for the XOR type, but of intermediate difficulty for the AND/OR type (the effect of labeling was not significant for AFF). Notice that there are two biased versions of the AND/OR pattern: one corresponding to the AND category in which instances that belong to the category have the distribution shown in table 1 (b), and one corresponding to the OR category in which instances that belong to the category have the distribution in 1 (c). Gottwald found the following order of learning difficulty (*b* stands for biased labeling and *n* for neutral labeling).

AFF_n, AFF_b > AND_b > AND/OR_n > XOR_b, OR_b > XOR_n

This order is consistent with the well-established rankings: AND_b > OR_b and AND_b > XOR_b (as well as AND/OR_n > XOR_n), but it does not include the expected ranking OR_b > XOR_b. It also shows that the neutral category AND/OR, which differs from its biased counterparts in that all instances are positively labeled, is harder to learn than AND_b but easier to learn than OR_b. Gottwald speculated that the “use of positive/negative labels focuses attention on the positive subset of stimuli and this would facilitate learning characteristics of these stimuli.” (*Ibid.* p.32). However, it is not clear why focusing attention of the positive subset of stimuli should help, and, in particular, why a biased OR_b category should be harder to learn than the neutral AND/OR_n category given that in the neutral category all stimuli are effectively positive and, hence, there are two different categories to learn. One possible explanation for the observed order AND/OR_n > OR_b is that in the neutral mode, where all labels are positive and equally frequent, subjects would quickly learn the conjunctive part of the pattern and treat the disjunctive part as the “elsewhere” case (the complement of the conjunctive part). On the other hand, in the biased OR mode subjects would not be (at least initially) inclined to analyze the disjunctive category as the complement of the conjunctive category (i.e., the negative instances) because they would not attend to negatively labeled stimuli as much as to positively labeled stimuli. For a concrete articulation of this proposal in terms of the learning mechanism that derives Gottwald’s results see the proposal in Pertsova (2012).

4. Morphological experiment

In this section we discuss a morphological analogue of Gottwald’s experiment that involved learning phrases in an artificial language rather than visual shapes. In this experiment, the difference between the neutral and biased labeling corresponded to the difference between realizing all morphological contrasts overtly (neutral labeling) vs. having no overt realization for one of the affixes (biased labeling). The rationale behind treating overt/null marking as an analogue of the type of labeling was as follows: it is possible that when learning a paradigm in which one of the

two affixes is null the learner will focus attention on learning the contexts for the overt affix assuming that no change (i.e., zero affix) occurs by default in all other contexts.

4.1. Stimuli and experimental conditions

The subjects had to learn the distribution of two suffixes in phrases of the form Prep + Noun-Suff as a function of two features: the type of the preposition (“above” / “beside” / “behind”) and the type of the noun (“bird” / “fish” / “insect”). Two random values for each of the features were chosen as relevant for each subject, so that each feature was binary. The suffix can be seen as marking a kind of a locative case with allomorphy conditioned by a semantically determined inflectional class.² Additionally, the choice of features was constrained by the following three criteria: (1) both features had to be morphologically irrelevant in the subjects’ native language (English) (2) the feature values for each feature had to be equally “marked”; that is, no feature value could be seen as somehow more basic, prior, or more frequent. The reason for the second constraint is the correlation between conceptual and morphological markedness (Battistella 1990). For instance, plural is considered to be more marked than singular and, correspondingly, there are many languages that only mark plural overtly, but almost no languages that only mark singular overtly and have null plural inflection throughout the lexicon (Croft 2003). Since we are interested in testing how presence of null morphs (the biased mode) would interact with the logical structure of a pattern, we want to abstract away from other factors that may favor null marking.

Two values for each feature were randomly chosen for each subject. There were five experimental conditions resulting from crossing two factors: the logical type of the paradigm (AND/OR vs. XOR) and the type of labeling (biased/neutral). These five conditions are summarized using schematic paradigms below.

Table 3: Experimental conditions: example paradigms

	neutral			biased (ANDb)			biased (ORb)		
	F1\F2	bird	fish	F1\F2	bird	fish	F1\F2	bird	fish
AND/OR	above	-B	-A	above	-A	-∅	above	-∅	-A
	beside	-A	-A	beside	-∅	-∅	beside	-A	-A
XOR	neutral			biased					
	F1\F2	bird	fish	F1\F2	bird	fish			
	above	-B	-A	above	-∅	-A			
	beside	-A	-B	beside	-A	-∅			

The artificial language (nonce roots, prepositions, and suffixes) were randomly

² There are many languages in which inflectional classes are at least partly determined by semantic factors such as animacy, mass/count distinction, animal/plant distinction and so on (e.g., Bantu languages).

Logical Complexity in Morphological Learning

generated for each subject. All roots had the shape CVCV, and all prepositions and suffixes had the form CV, where C was one of the following consonants {p, t, k, b, d, g, f, s, v, z}, and V was one of the following vowels {i, e, a, o, u}. Languages were designed not to contain any phrases with repetitions of identical syllables next to each other.

4.2. Procedure

The subjects were seated in front of a computer. They were told that they have to learn how to say simple phrases in another language, and that to do well in this task it is important to pay attention to the prepositions and the types of nouns they will see. They were told that they will encounter two of three possible prepositions “above,” “beside,” or “behind,” and that nouns could refer to fish, insects, or birds. That is, the features and their values were explicitly introduced to the subjects.³ Subjects then proceeded to a training phase which lasted for 64 trials. On each trial, they first saw a picture of an animal with an English label and its “translation” into the artificial language, which they were asked to read out loud. The next screen showed an English phrase using this noun (e.g., “beside the goldfish”), an illustrating picture (a goldfish with a black circle beside it), and two possible translations of this phrase that differed only in the choice of the suffix. The subject had to pick the translation they thought was correct and were given feedback in the form of a correct or incorrect beep. While this training procedure differs drastically from natural language learning, it is similar to most psychological studies of categorization that established the complexity ranking in question. The training phase was followed by a testing phase which consisted of 12 novel trials with no feedback.

4.3. Subjects

Sixty subjects (12 per condition) participated in the experiment. The subjects were recruited from the population of students and staff at the University of North Carolina at Chapel Hill. They were paid \$7 for participation.

5. Results

The recorded measurements for each subject were the number of errors during the testing phase, and the number of trials it took to reach the learning criterion. The learning criterion was set to 19 correct responses on 20 consecutive trials. For subjects who did not reach the learning criterion, this number was set to the maximum number of trials, 64. For others, the number of trials to criterion marks the number of trials to the last of the 20 consecutive responses. The descriptive statistics for the collected data are summarized in the table below.

From this table, we can already see that the ANDb condition was the easiest (as expected). The hardest pattern to learn was the biased XORb while in Gottwald’s experiment it was the neutral XORn.

³ In a pilot experiment which did not introduce the features, subjects could not learn the patterns.

The data was analyzed using logistic mixed-effects regression with two factors, type of pattern (AND/OR vs. XOR) and type of labeling and a random intercept for each participant. The labeling factor had three levels: neutral, up, and down. The “up” and “down” conditions differ with each other with respect to which of the two subsets of the pattern was overtly labeled. Note that for the XOR condition, both subsets are of equal size, so the difference between “up” and “down” conditions is not really meaningful and was not expected to be significant. For the AND/OR condition, the “up” version corresponds to ANDb and the “down” version corresponds to ORb. The results of the regression run on correct/incorrect responses during the testing phase are reported in table 5. The model can be interpreted as follows. With the change from ORb (the intercept) to XORb, the log odds of a correct response go down significantly ($p = 0.02$). The same is true with the change from ORb to the neutral mode, AND/OR ($p = 0.06$, marginally significant). There is a slight and non-significant improvement in performance with change from ORb to ANDb, and a marginally-significant interaction between XOR and the neutral mode ($p = 0.09$), suggesting that the neutral XOR was easier than expected based on the other cells. This interaction is also noticeable from the descriptive statistics: in the AND/OR category, subjects did numerically better on the biased conditions, while in the XOR category subjects did numerically better on the neutral condition.

The pairwise comparisons of the five conditions using Mann-Whitney U test with the Bonferroni adjusted alpha level of 0.005 per test showed the following significant differences summarized in the table below. (The data for the two biased XOR conditions was pooled together since the two groups are structurally identical.)

Overall, the biased conditions ANDb and ORb were significantly easier to learn than all other conditions. This is consistent with the well-established ranking $AND > XOR$, and with Gottwald’s result $ANDb > AND/OR$. However, there was no significant difference between ANDb and ORb, and the order $AND/ORn > XORb > XORn$ observed in Gottwald’s experiment was not replicated. Particularly surprising is the subjects’ performance on the XORn condition which was better than

Table 4: Descriptive statistics

	Num. of people to reach crit. (out of 12)	Mean num. of errors on test trials (out of 12)	Mean trials to criterion
AND/OR			
AND/ORn	6	2.2	42.5
ANDb	10	0.7	29
ORb	8	1.3	28.2
XOR			
XORn	8	1.4	35.7
XORb	4	2.3	46.7

Logical Complexity in Morphological Learning

Table 5: Coefficients of logistic regression

	Coeff.	Std. error	p
(Intercept)	2.0519	0.286	8.03e-13
XOR	-1.096	0.481	0.02
neutral	-0.732	0.398	0.06
up	0.101	0.405	0.79
XOR × neutral	1.042	0.619	0.09
XOR × up	0.809	0.691	0.24
	Variance		
Random eff. - subj	0.823		
Log-likelihood	-2093.000		
Deviance	4184.001		
AIC	4199.001		
BIC	4244.897		
N	4560		
Groups	60		

Table 6: Pairwise comparisons of all groups (> indicates that the category in the column was better than the category in the row)

	ANDb	ORb	AND/OR	XORb
ORb	n.s.	–	–	–
AND/OR	>	>	–	–
XORb	>	>	n.s.	–
XORn	>	>	n.s.	n.s.

expected as the results of the logistic regression suggest. This condition was virtually unlearnable in Gottwald’s experiment: 6 out of his 8 subjects failed to learn this pattern after 256 trials. On the other hand, in the present study 8 out of 12 subjects learned this pattern after less than 64 trials. (See table 2 for a morphological example of this pattern involving definite articles in Somali.)

6. Discussion

The findings of this study indicate that paradigms with the logical structure of AND/OR in which one of the affixes is null are easier to learn than an isomorphic paradigm in which all affixes are overt (AND/ORn) or paradigms with the logical structure of XOR. These results are only partially consistent with the previous non-linguistic findings. One of the differences is that the current study found ORb to be easier than XORb while Gottwald’s non-linguistic results show no significant differences between these two conditions. Another difference is that Gottwald established the ranking AND/ORn > XORb > XORn, while the current study found

no differences among these three conditions. Finally, the most significant difference between this linguistic study and its non-linguistic analogue is the difference involving OR_b and AND/OR_n. In learning artificial paradigms, OR_b was easier than AND/OR_n, while Gottwald found the reverse (AND/OR_n > OR_b) in his study. In the rest of this section we discuss possible reasons for these divergences in the experimental results.

This study was concerned with two factors affecting learning difficulty, the logical structure of a pattern (AND/OR vs. XOR) and type of labeling (biased vs. neutral). We hypothesized that the presence of null marking in a paradigm would have a similar effect to biased labeling by way of focusing subjects' attention on the overtly marked contrasts. However, it is possible that this is not the case. In other words, it could be that while subjects pay less attention to the negatively labeled stimuli than positively labeled stimuli, they do not pay less attention to the zero affix compared to an overtly affix (especially in the absence of semantic or other reasons to treat one of the feature values as marked). This would explain the lack of difference between the two biased conditions, AND_b and OR_b. If equal attention is paid to null and overt affixes, then our biased conditions should not be that much different from each other and from the neutral condition. It might still be somewhat easier to learn a biased pattern given that it places less burden on memory (one less new phonological string has to be memorized). This type of an explanation can be invoked for the found ranking AND_b, OR_b > AND/OR. However, if this explanation is on the right track, we would also expect the biased XOR_b to be easier than the neutral XOR_n. Yet, the two were not statistically different, and, in fact, showed a trend towards the opposite order (XOR_b was the hardest condition to learn).

Another possibility is that the null/overt marking has a similar effect as biased/neutral labeling after all, and the differences between the findings of two experiments are due to other differences in how these experiments were conducted. In particular, the experiments differed in two crucial respects that can significantly affect learning difficulty: the number of irrelevant features, known to affect pattern learning (Kepros and Bourne 1966), and the probability distribution of the input data. In the present experiment there were no other features that were systematically varied across the stimuli, while in Gottwald's experiment there were two additional irrelevant features (so the morphological subjects were doing just "rule learning" while Gottwald's subjects were doing both "rule learning" and "feature identification"). Also, Gottwald used a response-balanced distribution of data, where each response label was correct on half the trials. We did not use this type of distribution because we found that it often led subjects to adopt an incorrect one-feature hypothesis when learning AND/OR paradigms. This is because when the stimuli is balanced to have equal number of each label (i.e., each affix), a one-feature hypothesis is correct on 5/6 of all AND/OR stimuli. Therefore, we used a distribution in which each cell in the paradigm had an equal probability of occurrence in the data. This meant that in the AND/OR conditions, one of the affixes was three times more frequent than the other. These details may have big consequences for the results.

Logical Complexity in Morphological Learning

Such consequences can only be estimated with respect to a concrete learning strategy. Different learning models are differently affected by the variations in the probability distribution of the data. It would be useful to simulate experimental conditions and compare different models' predictions to the empirical results obtained under such conditions. Taken together, the empirical data from artificial category learning experiments and predictions from learning models can help us to zero in on what factors affect pattern complexity, and whether these factors are different for linguistic vs. non-linguistic categories. Pertsova (2012) is a recent attempt to provide a learning model that predicts the differences between learning in the neutral vs. the biased mode. This model is largely consistent with Gottwald's results. Additionally, it makes somewhat different predictions for learning AND/OR and XOR patterns depending on the distribution of instances in the data. It predicts that OR_b should be learned faster than XOR_b when truth-table cells are equally frequent (as in our experiment), but not when each response label is equally frequent. This prediction is consistent with the differences observed between the Gottwald's experiment and the present experiment. However, the same algorithm predicts the difficulty order AND/OR_n > OR_b which was observed in Gottwald's experiment but the reverse was observed in the present experiment, as well as the order XOR_b > XOR_n which, on the other hand, is similar to what was found in the present experiment but the opposite from the order observed in Gottwald's experiment.

In short, further work, both on the empirical and modeling front, is required to better understand the interplay of different factors contributing to pattern-learning complexity. The results reported here make the first steps in that direction. These results are promising because they replicate several biases found in learning of non-linguistic categories and in typological data. As discussed in this section, the differences between our results and previous findings may be due to additional factors that affect complexity (e.g., the probability distribution of the data). These factors may be particularly relevant for learning of linguistic categories where it is not uncommon to find unequal frequency distributions among different affixes.

7. Conclusion

We showed that learning morphological paradigms is subject to some of the same learning biases as learning of non-linguistic patterns. In particular, we replicated the well established complexity ranking AND, OR > XOR with artificial morphology data. (This ranking is also predicted by majority of categorization models.) That is, paradigms with the logical structure of XOR (cf. Somali in table 2) exhibiting the so-called "polarity effects" are more difficult to learn than paradigms in which the overt affix occupies a natural class of cells (corresponding to a conjunctive category), or a set of cells that can be described as a complement of some natural class. The prevalence of conjunctive patterns is also evident in typological data for both morphological and phonological categories (see section 2). This fact lends support to the hypothesis that the biases discovered in the lab are similar to

those that operate in natural language learning. We have also found that paradigms with null affixes are easier to learn than paradigms in which all affixes are overt if the paradigm has the AND/OR (but not the XOR) structure. The reasons for this finding are less clear and demand further investigation. Overall, effects of different learning factors interact in complex ways. Artificial grammar learning experiments and computational modeling provide tools for disentangling these interactions.

References

- Baerman, Matthew and Brown, Dunstan, 2005. *Syncretism in verbal person/number marking*. In M. Haspelmath, M. S. Dryer, D. Gil, and B. Comrie, eds., *The world atlas of language structures*. London: Oxford University Press.
- Baerman, Matthew, Brown, Dunstan, and Corbett, Greville, 2002. *The Surrey Person Syncretisms Database*. <http://www.smg.surrey.ac.uk/syncretism/index.aspx>. Accessed in 2007.
- Battistella, Edwin, 1990. *Markedness: the evaluative superstructure of language*. State University of New York Press.
- Bruner, Jerome S., Goodnow, Jacqueline J., and Austin, George A., 1956. *A study of thinking*. New York: John Wiley and Sons.
- Croft, William, 2003. *Typology and universals*. Cambridge University Press, 2nd edition.
- Cysouw, Michael, 2003. *The paradigmatic structure of person marking*. Oxford studies in typology and linguistic theory, Oxford: Oxford University Press.
- Gottwald, Richard L., 1971. *Effects of response labels in concept attainment*. *Journal of Experimental Psychology* 91(1):30–33.
- Kepros, Peter C. and Bourne, Lyle E., 1966. *Identification of biconditional concepts: effect of number of relevant and irrelevant dimensions*. *Canadian Journal of Psychology/Revue Canadienne de Psychologie* 20(2):198–207.
- Mielke, Jeff, 2004. *The emergence of distinctive features*. Ph.D. thesis, Ohio State University.
- Morton, Elliott, 2012. *Work note: Complexity in P-Base classes*. Manuscript, Department of Linguistics, UNC-Chapel Hill, April 15, 2012.
- Mielke, Jeff, 2008. *The emergence of distinctive features*. Oxford University Press.
- Neisser, Ulrich and Weene, Paul, 1962. *Hierarchies in concept attainment*. *Journal of Experimental Psychology* 64(6):640–645.
- Pertsova, Katya, 2007. *Learning Form-Meaning Mappings in the Presence of Homonymy*. Ph.D. thesis, UCLA.
- Pertsova, Katya, 2012. *Linguistic categorization and complexity*. In *12th Meeting of the ACL Special Interest Group on Computational Phonology and Morphology*.
- Peters, Kenneth G. and Denny, J. Peter, 1971. *Labeling and memory effects on categorizing and hypothesizing behavior for biconditional and conditional rules*. *Journal of Experimental Psychology* 87(2):229–233.

Logical Complexity in Morphological Learning

Katya Pertsova
Linguistics Department, CB #3155
University of North Carolina at Chapel Hill
Chapel Hill, NC 27599

pertsova@unc.edu