

## Chapter 7

# Conclusion

In this thesis, I have presented a computationally implemented model of paradigm learning in which learners identify one form as the base form, and then learn a grammar of morphological and phonological rules that operate on the base form to derive the remaining forms in the paradigm. I have pursued two related, but logically distinct hypotheses about how this is done. The first hypothesis is that learners are limited to selecting a single form as the base, and that the base form must be a surface form from somewhere within the paradigm. Furthermore, the choice of base is global, meaning that the same part of the paradigm must serve as the base for all lexical items. Such an approach is similar to the “basic alternant” approach of some pre-generative phonologists. The second hypothesis is that learners select the base form that is maximally informative, in the sense that it preserves the most contrasts, and permits accurate and productive generation of as many forms of as many words as possible. I have argued along the way that these hypotheses could be interpreted in different ways, depending on the model of morphology that is employed. If one adopts a word-based model of morphology, in which morphological operations take whole words as their inputs, then these are hypotheses about how learners figure out which forms in the paradigm to derive from which other forms. If, on the other hand, one adopts a stem-based model of morphology, in which morphological operations combine morphemes listed in their underlying forms, then these hypotheses constitute a restriction on the set of possible underlying forms that learners are allowed to consider.

The single surface base hypothesis is far more restrictive than the standard assumptions about how speakers arrive at the underlying representations of words. It is generally assumed in the phonological literature that speakers can combine unpredictable information from anywhere in the paradigm to construct underlying forms that ideally can be used to derive all surface properties of the word (Kenstowicz and Kisseberth 1977). When inputs are restricted to a single surface form, on the other hand, it is often impossible to find a single form that unambiguously reveals the surface properties of all of the members of the paradigm.

This situation can arise in two different ways. The first is when a language has multiple neutralizing phonological processes, with some affecting some slots in the paradigm, and others affecting other slots. This was the case in Latin, discussed in chapter 4: the nominative form preserved the  $s \sim r$  contrast but suffered from various phonological neutralizations, while the oblique forms preserved most contrasts but merged stem-final  $s$  and  $r$ . This is also the case for stress-shifting nouns in Russian (section 6.1): unsuffixed forms preserve the contrast between  $a$  and  $o$  but suffer from final devoicing, while suffixed forms preserve the voicing of

final obstruents but suffer from vowel reduction. In both Latin and Russian, every slot in the paradigm suffers from at least one phonological neutralization, so no single slot reveals all properties of all words.

The other situation in which there is no surface form that preserves all contrasts arises when a neutralization is *symmetrical* (chapter 5): the language contains non-alternating segments *a* and *b*, and an alternating segment  $a \sim b$  that is neutralized to *a* in some forms and to *b* in others. This was the case in an earlier stage of Yiddish (chapter 2), since verbs with  $a \sim e$  umlaut alternations were neutralized with invariant *a* verbs in some forms, and with invariant *e* verbs in others. The same was also true for final  $a \sim e$  ablaut alternations in Lakhota (chapter 5). In all of these cases, the single surface base restriction means that the underlying form of a word cannot unambiguously encode all of the surface contrasts of the language.

If speakers cannot encode all of the unpredictable surface properties of a word in its underlying form, then some of this information must be stored elsewhere in the system. For example, if we cannot encode the contrast between root-final *s* and *r* in Latin as part of the underlying form, we need some other way to keep track of whether a particular word should end in *s* or *r* in the nominative. I proposed that in the absence of abstract underlying forms, speakers may nevertheless be able to overcome such neutralizations using a sophisticated system of morphological rules and exceptions. If we allow morphological rules to refer to details of the phonological and morphosyntactic environment—such as “after coronal obstruents”, “in neuter nouns”, and so on—then we can use these to formulate a grammar that captures intermediate and small-scale generalizations about groups of words that tend to behave alike. In the case of Latin, such a grammar includes rules specifying that masculine and feminine polysyllabic nouns ending in *-ris* in the genitive should end in *-r* in the nominative, while neuter nouns ending in *-ris* in the genitive should end in *-s* in the nominative, as should monosyllabic nouns of any gender. For words that do not conform to these generalizations, the nominative form must simply be stored as a lexical exception.

The single surface base hypothesis runs contrary to the standard use of underlying forms in phonology, but it has some advantages. First, it is computationally simpler as a learning strategy, since it greatly reduces the number of underlying representations that must be considered. It also means that there need not be a mechanism for positing abstract underlying representations that combine information from multiple forms, or that contain underspecified elements such as archiphonemes, since such elements never occur in any surface form. Moreover, I have shown throughout the course of this thesis that the single surface base hypothesis also has empirical advantages. In particular, it predicts asymmetries in which forms and which contrasts should be open to analogical change: analogy should affect only non-basic forms, and it should target only those contrasts that are not preserved in the base form. In other words, analogical changes happen only when speakers are forced to guess their way out of a neutralization.

There are several aspects of this proposal that could benefit from further inquiry. First, I have needed to suppose at various points that there is a bias towards selecting bases with higher token frequency, because they are more readily available as inputs to derive other forms. I have tried to maintain the view that this bias is not actually a formal grammar acquisition strategy, but is rather a side effect of the way that learners assess the reliability of generalizations. In particular, generalizations based on more data are more trustable, and higher frequency forms give the learner more input data from which to create generalizations. I suggested that the

confidence limit adjustment already employed by the minimal generalization learner might provide an appropriate way to model this. Further modeling, with more realistic input sets in which various forms occur in their actual proportions, is needed to better understand the trade-off between informativeness and availability, and whether confidence limit statistics are the best way to model it.

Further research is also required concerning the implications of this proposal for the use of underlying forms in phonology. In all three cases discussed here (Yiddish, Latin, Lakhota), I argued that a model that can combined unpredictable information from various parts of the paradigm cannot predict the asymmetries observed in subsequent historical changes. In all three of these cases, however, the amount of recombination that would be necessary to arrive at the “traditional” underlying forms was quite small, and the expense of having to store the extra information as lexical exceptions was not all that great. In chapter 6, I considered whether this type of analysis could be extended to a case in which significantly more recombination is necessary (Russian), suggesting that perhaps even here, an analysis restricted to using a single surface form as the UR might not require an unreasonably large number of lexical exceptions. It would be worthwhile expanding this analysis beyond the sample of words that I have considered in that discussion, to see what would be required to analyze all of Russian under this system. More generally, phonological analyses that are restricted to using a single surface form as the UR may often look quite different from traditional analyses, and have the potential to make unique predictions about what patterns speakers may notice and extend.

Finally, I have mentioned at several points that data from language change is only one of the many areas in which base effects can be seen. Two other areas that seem particularly relevant are acquisition evidence from children, and experimental evidence from psycholinguistic studies. As discussed in section 6.4, the type of acquisition evidence that would most directly support or refute this approach concerns asymmetries between forms: we can test whether overgeneralization affects some parts of the paradigm but leaves others intact, extends some patterns but not others, and so on. In psycholinguistic experiments, too, we might look for asymmetries, such as differences in how inflectionally forms affect one another in a morphological priming task, or some type of dissociation between forms that are predicted to be productively derived by the grammar and those that must be stored as lexical exceptions. Converging evidence from historical changes, child data, and psycholinguistic tasks could play a crucial role in assessing the plausibility of the current approach as a model of how speakers learn to produce and comprehend paradigms of inflectionally related forms.

