

The PAKA World

The primary contributors to this work are John Alderete, Graham Horwood, Koichi Nishitani, Alan Prince, and Bruce Tesar.

Learning Phonological Alternations

A central challenge of learning phonological systems is the determination of the mappings underlying phonological alternations. Phonotactic learning alone can (in principle) determine phonotactic distributions, the restrictions on what configurations may and may not occur in grammatical output forms. Presumably, inputs whose faithful preservation would require violation of phonotactic restrictions are not faithfully preserved, but instead mapped onto something else. But what? Phonotactic learning alone can only imply that banned structures must be mapped onto “something grammatical,” but may not in general determine which specific grammatical structure.

To gain more insight into the mapping, the learner must turn to data which motivate the learner to care: phonological alternations. If the same morpheme surfaces differently in different contexts, the simplest analysis will involve a single phonological underlying form for the morpheme and a single mapping, such that the underlying form is mapped to the (different) attested surface form in each (different) context. The learner’s task includes coming up with both of the elements of analysis just described (underlying form and mapping) in response to the surface forms indicating the alternations.

It bears emphasizing that the mapping and the underlying forms are interdependent, and both must be inferred by the learner. We will here liken this to other forms mutual dependence in language learning, such as structural ambiguity. In this case, the hidden structure that must be inferred by the learner includes the phonological underlying forms of the morphemes that collectively determine the observed form. It is again the case that if the learner knew the hidden structure (underlying forms), they could easily determine the mapping (under present assumptions, an Optimality Theoretic constraint ranking). It is further plausible that knowing the correct constraint ranking would be of great benefit to the learner when determining the underlying forms for the morphemes, although this has not yet been conclusively demonstrated. As usual, the learner starts out knowing neither the ranking nor the underlying forms, and must infer them together.

The first goal of the present work is to isolate the problem of UF/ranking interdependence from the many other challenging problems surrounding the learning of morphophonological systems. In particular, we want to abstract away from problems of segmentation and morpheme identification *in the output*. For the learning problem we wish to pose here, we will assume that the learner receives an organized set of “semi-analyzed” output forms. Each output form will include an indication of the morphemes that constitute it, and an explicit indication of which phonological elements correspond to each morpheme. The learner will have the output forms organized into paradigms, indicating a set of output forms that share a given morpheme. Under such assumptions, the goal of the learner will be to find the most restrictive grammar consistent with the data. More precisely, the learner must identify a lexicon of underlying forms, one UF for each morpheme, and a ranking, such that all the observed surface forms are correctly generated. Further, when several such grammars are possible, the learner should select from among them a most restrictive grammar.

We here define a very simple linguistic system, just about the simplest imaginable system consistent with the conditions just described, to allow us to investigate in concrete terms this learning problem. The only thing differing from language to language in this system is the assignment of stress; all other elements are fixed

across languages in this system. The only kind of phonological information that can be specified in an underlying form is an accent: a morpheme may be specified as accented, or not. The set of possible morphemes is fixed across the languages of the system. It contains markedness constraints assessing the positioning of stress, and faithfulness constraints assessing the realization of underlying accent. To learn a particular language within the system, the learner must learn the correct ranking of the constraints, and correctly assign underlying accents to the appropriate morphemes of the system.

A Simple Morphoaccentual Domain: The PAKA World

The system has a total of two roots and two suffixes.

Roots: /pa/, /ba/

Suffixes: /-ka/, /-ga/

The voiced morphemes have underlying accent, while the voiceless ones do not.

Thus, every language in this system contains a total of four words. What varies from language to language is the stressing of each of those words.

Stress is culminative: every word in every language must have exactly one stress on the surface.

Phonologically, the only hidden structure that needs to be inferred by the learner is the presence/absence of lexical accent on a morpheme.

The system contains four constraints:

MAINLEFT (ML) – main stress must be on the leftmost syllable (really, an alignment constraint)

MAINRIGHT (MR) – main stress must be on the rightmost syllable

FAITHPROMINENCE (F) – an underlyingly accented syllable must be stressed.

FAITHPROMINENCEROOT (FR) – an underlyingly accented root syllable must be stressed.

Typology (Nishitani Aug. 2001)

Factorial Typology: FAITHPROM (F), FAITHPROM_{ROOT} (FR), MAINLEFT (ML), MAINRIGHT (MR)

ID	(1)	(2)	(3)	(4)	(5)
Hierarchy	ML >> {MR, F, FR}	MR >> {ML, F, FR}	F >> ML >> {MR, FR}	F >> MR >> {ML, FR}	FR >> MR >> {ML, F}
Forms	PAka PAga BAka BAga	paKA paGA baKA baGA	PAka paGA BAka BAga	paKA paGA BAka baGA	paKA paGA BAka BAga
Possible Lexica	pa/pa' ka/ka' ba/ba' ga/ga'	pa/pa' ka/ka' ba/ba' ga/ga'	pa ka ba' ga'	pa ka ba' ga'	pa ka/ka' ba' ga/ga'
Description	Finnish		Pseudo GreekRussian		Pseudo Cupeno

Exercise: languages (1) and (2) are a “symmetric pair” if you compare their rankings; so are (3) and (4). Why doesn't the typology include a symmetric counterpart for (5)?

If we optimize the r-measure, we expect a particular hierarchy to be the “best outcome” for each language.

Language	Hierarchy	R-measure
(1)	ML » MR » {F, FR}	4
(2)	MR » ML » {F, FR}	4
(3)	F » ML » MR » FR	2
(4)	F » MR » ML » FR	2
(5)	FR » MR » ML » F	2

From the point of view of BCD:

- faith constraints go either at the top or the bottom of the hierarchy
- only one faith constraint ever need dominate markedness.

NOTE: these two observations are specific to this particular system; they do NOT hold for OT analyses in general.

The Learning Problem

Criteria for Success

The learner is presented with a paradigm of forms, indicating the morphological affiliation and the overt stresses of the syllables. The learner needs to find a constraint hierarchy and a set of underlying forms that will reproduce the paradigm.

Input to Learner: { PA-ka pa-GA BA-ka BA-ga }

Output from (successful) learner: F » ML » {MR, FR} /pa ba' -ka -ga'/

The assumption used in generating languages of the system, that (all and only) voiced morphemes are accented, implements richness of the base: there is one morpheme of each type. It is not necessary that the learner recover that lexicon in all cases, what is necessary is:

- the learner get the right surface forms, using the learner's underlying forms.
- the learner's ranking produces the same surface forms when the canonical UFs are used; this is a way (for this example) of requiring the most restrictive grammar possible.

This second criterion seems to be equivalent to the following: the learner can have a UF that differs from the target one ONLY when the learned ranking neutralizes the two UFs to a single (correct) output, for each case.

Subset Relations: The Linguist's View

The languages of the system differ in their restrictiveness: some always put stress in the same place regardless of morphological identity, while others allow accentual distinctions to surface. This creates the potential for subset relations among languages.

The base consists of four possibilities: unaccented root, accented root, unaccented suffix, accented suffix. Naively, one might expect to express subset relations in these terms: different languages preserve different subsets of this range of lexical possibilities. But things get more complicated when we consider morphologically complex forms. To see the restrictiveness problem as a subset problem, we need the right characterization of the forms and languages. In particular, we need to specify in the right way what the "sets" are that can possibly enter into subset relations. The paradigm includes the behavior of forms in different contexts, so the "things" in the "sets" must capture the behavior of single morphemes across contexts.

In the PAKA world, a morpheme's relation to its context is particularly trivial, as it is predictable from the state of the (single) context morpheme itself. If a root is stressed, its suffix is not, and vice-versa; likewise

for suffixes. Thus, the “behavior in different contexts” for a morpheme reduces to the set of context morphemes itself. There are four possible context sets for each type of morpheme. For roots, the possibilities are {ka,ga}, {KA,ga}, {ka,GA}, and {KA,GA}. For suffixes, the possibilities are {pa,ba}, {PA,ba}, {pa,BA}, and {PA,BA}.

However, to truly see the abstract inventories for what they are, we want to characterize a morpheme’s relation to its context in a way that doesn’t rely on the full identity of the individual morphemes characterizing the context. That is, we want to talk about how a *type of* root behaves in the presence of some particular *type of* suffix, rather than how /pa/ behaves in the presence of /-ka/. This requires some appropriate type theory.

The behavior of concern here is stressed/unstressed. If a morpheme is always stressed, or always unstressed, then we know how it behaves with respect to each possible context without further elaboration. The need for further elaboration arises when a morpheme alternates; then we might care about distinguishing which types of contexts cause it to surface as stressed, and which cause it to surface as unstressed. We can take a first step towards this by characterizing each root/suffix as either always stressed, always unstressed, or alternating. Such a classification is indicated in the table below.

ID	(1)	(2)	(3)	(4)	(5)
Forms	PAka PAga BAka BAga	paKA paGA baKA baGA	PAka paGA BAka BAga	paKA paGA BAka baGA	paKA paGA BAka BAga
Inventory	root-str suffix-unstr	root-unstr suffix-str	root-alt root-str suffix-unstr suffix-alt	root-unstr root-alt suffix-alt suffix-str	root-unstr root-str suffix-alt

Once the inventories are expressed in this way, it is easy to spot the subset relations among the languages:

$$(1) \subseteq (3)$$

$$(2) \subseteq (4)$$

Exercise: FAITHPROMINENCEROOT and FAITHPROMINENCE have a special/general relationship between them. Yet no subset language relations appear to be based upon substituting the general in place of the special. Why is this?

This is probably adequate for our PAKA world as it actually happens, but in principle, when a morpheme alternates, we would like to distinguish which *type of* context morpheme causes it to be stressed/unstressed. However, this creates in a tricky interdependence: we need to know what the possible types of suffixes are in order to define the possible types of roots, and vice-versa. To use this to spot subset relations between languages, we would need some sort of valid cross-language classification of all possible types of morpheme behavior.

It is important to be clear about the role of each part of the grammar. The ranking determines the inventory of possible morpheme types for a given language, while the lexicon assigns each morpheme a particular type. A morpheme’s type is assigned via its underlying form (in this case, by the presence/absence of a lexical accent). Richness of the base for underlying forms implies, in turn, richness of the base for morpheme types.

Now look back at the success criterion for learning, to the part about when a learned UF may differ from the target one. We can now express this as saying that a learned UF may differ from the target one precisely when the ranking maps both UFs to the same morpheme type. This is another way of saying that the underlying specification distinguishing the two UFs is *not contrastive*.

The idea that an underlying specification is not contrastive if UFs distinguished only by that specification neutralize is hardly new. What makes things complicated is how, exactly, we characterize neutralization: we need to track the corresponding behavior of the two UFs across contexts (including the effects on the context morphemes). In the above discussion, it is neutralization of two UFs *to the same morpheme type*.

The learner’s task, then, can be characterized as needing to find a constraint ranking which selects the correct inventory of morpheme types, and to give each morpheme a UF assigning it the correct type.

Subset Relations: The Learner’s View

In the preceding analysis of subset relations between languages, language (5) did not enter into subset relations with any of the other languages. If one is attempting to learn language (5), restrictiveness won’t be a problem. Or so it would appear. The reality isn’t so kind.

Suppose a learner is presented with the data for language (5) and comes up with the following grammar in response:

F » MR » ML » FR /pa ba’ -ka -ga/

You should be able to verify that this in fact generates the forms of language (5). But is it a correct grammar? No. In fact, the ranking is that of grammar (4). What has gone wrong? The key is that the learner is not considering the consequences of an accented suffix. Both suffixes are posited to be unaccented. But if language (5) isn’t supposed to be a subset of language (4), why isn’t the learner able to see that?

Note first that the suffix inventories themselves DO form a subset relation: (5) has only alternating suffixes, while in language (4) unaccented suffixes alternate (matching language (5) suffixes) and accented suffixes are always stressed (missing from language (5)). If we look only at suffixes, (5) is a subset of (4). Further, we can see how the learner misses this fact in the language above: because they never observe always stressed suffixes, they don’t posit an accented suffix in the lexicon.

The analysis from the linguist’s view concluded that language (5) isn’t a subset of language (4) for the following reason: (5) has an always stressed root, while (4) instead has an alternating root (both also have an unstressed root). So, our learner is being presented data from a language containing always-stressed roots, and has come up with a grammar generating alternating roots. Why can’t it see the difference? The reason is that in language (4) the accented root “alternates” in response to an accented suffix. But accented suffixes are precisely the part of the suffix base that the learner isn’t considering. If the learner were to posit an accented suffix, they would then see that their accented roots were alternating instead of always stressed, and need to change the ranking. But the learner doesn’t see this, because they have posited only unaccented suffixes.

Notice that when we compare the rankings of (4) and (5), we might independently expect a subset relation:

(4)	F » MR » ML » FR
(5)	FR » MR » ML » F

The dominant faithfulness constraints in each of the languages are in a special/general relation. The

dominant faith constraint in (5), FR, is the special one, so we might suspect that language to be a subset of the language with an otherwise-similar ranking but with the general faith constraint, F, dominant, as is the case for language (4).

The problem arises from the possibility of only observing some of a morpheme’s potential behavior. An alternating morpheme is one that appears stressed in some environments and unstressed in others. The range of behaviors of a non-alternating morpheme can thus be seen as a kind of “subset” of the range of behaviors of an alternating one. If all relevant environments are seen, we can tell that an alternating morpheme is in fact alternating. But if an alternating morpheme is seen only in contexts that cause it to surface one way, it can appear to be non-alternating. Thus, in the case of our hapless learner above, the accented root /ba’/ is alternating from the point of view of the whole system, because it is stressed before unaccented suffixes and stressed before accented suffixes. But, if we only examine it in the contexts of unaccented suffixes, we only ever see it surface as stressed, and it “appears” to be non-alternating, always-stressed.

We can now re-examine the possible subset relations among the languages in this light. Below, the table giving the inventory of morpheme types is repeated.

ID	(1)	(2)	(3)	(4)	(5)
Forms	PAka PAga BAka BAga	paKA paGA baKA baGA	PAka paGA BAka BAga	paKA paGA BAka baGA	paKA paGA BAka BAga
Inventory	root-str suffix-unstr	root-unstr suffix-str	root-alt root-str suffix-unstr suffix-alt	root-unstr root-alt suffix-alt suffix-str	root-unstr root-str suffix-alt

We had previously concluded $(1) \subseteq (3)$ and $(2) \subseteq (4)$, and we will not change those conclusions. We’ve just further concluded that $(5) \subseteq (4)$. We can now examine it under the further consideration that always stressed and always unstressed, as states, can be thought of as subsets of alternating.

Comparing (1) to (5), we see that (1) might be a subset of (5) under this view. Why? The subset relations among the root types should be clear: $\{\text{root-str}\} \subseteq \{\text{root-unstr}, \text{root-str}\}$. The suffixes also yield this conclusion under the view that $\{\text{suffix-unstr}\}$ could be a subset of $\{\text{suffix-alt}\}$. In fact, this suspicion is confirmed once we realize that the ranking of (5), $FR \gg MR \gg ML \gg F$, can generate the pattern of language (1) if we assume a lexicon in which all roots are accented: $\{\text{pa}', \text{ba}', \text{-ka}, \text{-ga}\}$. If a learner posed with the surface forms of (1) were to mistakenly select that grammar, no positive evidence (from (1)) would ever contradict them, yet their grammar would mistakenly predict the existence of words with final stress due to unaccented roots. Similar reasoning produces the conclusion that $(2) \subseteq (5)$. The language (2) results from a lexicon in which all roots are unaccented.

What about (1) and (4)? The inventories suggest a possible subset relation, and in fact one can be achieved, with the ranking of (4) and a lexicon of accented roots and unaccented suffixes. In an analogous way, $(2) \subseteq (3)$. In both cases, the superset language has high-ranking faith, so each form of the subset language can be accounted for (non-restrictively) via faithfulness.

What about (5) and (3)? The refined analysis of the morpheme type inventories suggests that $(5) \subseteq (3)$. The ranking for (3), $F \gg ML \gg MR \gg FR$, will generate (5) with a lexicon of $\{\text{pa}, \text{ba}', \text{-ka}', \text{-ga}'\}$. This lexicon

can be derived by following the reasoning of the subset relations of the inventory.

We thus end of with subset relations picture, from the point of view of the learner, of

$$[(1),(2)] \subseteq (5) \subseteq [(3),(4)]$$

Note that this is really a further articulation of the restrictiveness predicted by the r-measure. The r-measure misses the distinction between (5) and [(3),(4)] because the distinction involves the special/general relationship between the two faithfulness constraints.