

Constraints, Exceptions and Representations

T. Mark Ellison

Centre for Cognitive Science, University of Edinburgh
2 Buccleuch Pl., Edinburgh EH8 9LW, U.K.
marke@cogsci.ed.ac.uk

Abstract

This paper shows that default-based phonologies have the potential to capture morphophonological generalisations which cannot be captured by non-default theories. In achieving this result, I offer a characterisation of Underspecification Theory and Optimality Theory in terms of their methods for ordering defaults. The result means that machine learning techniques for building declarative analyses may not provide an adequate basis for morphophonological analysis.

Introduction

In other work, I have shown (Ellison 1992, forthcoming) that interesting phonological constraints can be learned despite the presence of exceptions. Each of these constraints imposes a limit the set of possible words at a common level of representation. In this paper, I consider possible limits to the usefulness of these constraints in representing morphemes and finding concise representations of lexical entries.

In order to compare a strictly declarative formalism with other constraint formalisms, a common formal environment must be established. Using model theory to establish the relationship between description and object, and then a modal formalism to define the structures to which constraints apply, we can compare the different effects of strict constraints and defaults. In particular, a strict declarative approach can be compared with other constraint frameworks such as Underspecification Theory (UT) (Archangeli, 1984) and Optimality Theory (OT) (Prince & Smolensky, 1993). This discussion is followed in the latter part of the paper by consideration of the possibility of using machine learning to constraint systems that use defaults.

Morphophonology

To structure the discussion, I offer four desiderata for morphophonology. The first is that the morphophonology must allow concise lexical representations. Where information is predictable, it should not have to be specified in the lexicon. This desideratum is not a matter of empirical accuracy, rather one of scientific aesthetics. For example, English contains no front rounded vowels,

so a vowel which is marked as front in the lexicon need not be marked as unrounded.

The second desideratum is that the morphophonology should allow generalisations to be made over phonologically conditioned allomorphs. For example, a representation of the Turkish plural affixes *-lar*, *-ler*, that uses the feature $[\pm\text{front}]$ is superior to a segmental representation because a single representation for the two allomorphs can be achieved by not specifying the value for this feature in the representation of the morph.

The third desideratum requires that the specific allomorphs be recoverable from the generalisations. If *-lar* and *-ler* are generalised in a single representation, such as *-lAr*, then the morphophonology should make the recovery of the allomorphs in the correct environments possible.

The final desideratum is, like the first, a matter of scientific aesthetics: a priori abstractions should not be used in an analysis any more than is necessary. For example, the feature $[\pm\text{front}]$ should not be used in the analysis of a language unless it is motivated by structures in the language itself. This desideratum may conflict with the first: a priori features may result in a more concise representation.

These four desiderata provide a framework for evaluating the relative merits of monostratal systems of phonological constraints with other current theories such as Underspecification Theory and Optimality Theory.

Model Theory and Modal Logic

A fundamental distinction in any formal account is the distinction between description and object. Failure to make the distinction can lead, at best, to confusion, and, at worst, to paradoxes, such as Russell's Paradox. Because this theory is talking *about* theories, it must make the distinction explicitly by formalising the relationship between description and object. This distinction is pursued in below and developed into a formalism for complex structures in the following section.

Model theory

In model theory, the meaning of a statement in a formal language is provided by means of an INTERPRETATION

FUNCTION which maps the statement onto the set of objects for which the statement is true. If L is a language and W is a set of objects, and $\mathcal{P}(W)$ is the set of all subsets of W , then the interpretation function I maps L onto $\mathcal{P}(W)$:

$$I : L \rightarrow \mathcal{P}(W).$$

As an example, suppose \clubsuit is a horse, \diamond is a ferret and \heartsuit is a large stone, and that these are the objects in our world. We might define a language L_0 containing the terms *big*, *animate*, *slow* and *human*, and assign these terms the interpretations in (1).

Term T	Interpretation $I_0(T)$
<i>big</i>	$\{\clubsuit, \heartsuit\}$
<i>animate</i>	$\{\clubsuit, \diamond\}$
<i>slow</i>	$\{\diamond, \heartsuit\}$
<i>human</i>	$\{\}$

This language can be expanded to include the logical operations of conjunction, disjunction and negation. These are provided a semantics by combining the semantics of the terms they apply to.

Term	Interpretation
$l \in L_0$	$I_0(l)$
$X \wedge Y$	$I(X) \cap I(Y)$
$X \vee Y$	$I(X) \cup I(Y)$
$\neg X$	$W \setminus I(X)$

With this interpretation function, we can determine that *big* \wedge *animate* \wedge *slow* is a CONTRADICTION having a null interpretation in W , while *big* \vee *slow* is a TAUTOLOGY as $I(\text{big} \vee \text{slow})$ is the same as $I(\text{big}) \cup I(\text{slow})$ which equals W .

The term PREDICATE will be used to describe a statement in a language which has a well-defined interpretation.

Modal logics

Model theory as defined in section applies only to domains with atomic, unstructured objects. More complex structures can be captured by extending the theory of models to refer to different worlds and the relationships between them. Such a complex of worlds and relations is called a MODAL logic.

A modal theory consists of a universe U is a set of worlds $W_{j,j \in \mathcal{W}}$, called TYPES, together with a set of relations $R_{k,k \in \mathcal{R}} : W_{\text{dom}(j)} \rightarrow W_{\text{cod}(k)}$ from one world to another. Types may contain other types, and whenever a type is so contained, it defines a characteristic relation which selects elements of that subtype from the larger type. A language for this universe is more complex as well, needing a function $\omega : L \rightarrow I$ to indicate the type $W_{\omega(l)}$ in which any given expression l is to be interpreted. A MODAL OPERATOR r_k is a special symbol in the language which is interpreted as the relation R_k .

Modal operators can combine with predicates to construct new predicates. If ϕ is a predicate, r_k is a

modal operator and $\omega(\phi) = \text{cod}(k)$ then we can define an interpretation, $I(r_k \phi) \subseteq W_{\text{dom}(k)}$, for $r_k \phi$, namely $R_k^{-1}[I(\phi)]$. Furthermore, we define the type of the expression to be the domain of the functor: $\omega(r_k \phi) = \text{dom}(k)$. The interpretation of any well-formed sentence in this language is a subset of the corresponding world $I(\phi) \subseteq W_{\omega(\phi)}$.

From here on, we will assume that the $R_{k,k \in \mathcal{R}}$ are functions, and call the corresponding operators of the language FUNCTORS. Functors simplify the interpretation of predicates: inverses of functions preserve intersection, so functors distribute over conjunction as well as disjunction.

A path equation defines a predicate which selects entities that have the same result when passed through two different sequences of functions. Suppose that p and q are two sequences of functors with the same first domain and last codomain, and that the composition of the corresponding sequences of functions are P and Q respectively. Then the interpretation of $p = q$ is the set of entities x in the common domain such that $P(x) = Q(x)$.

Suppose the universe U consists of seven worlds, a , b , c , *alphabet*, *nullstring*, *nonnullstring* and *string*. Some of these worlds are built from others: *alphabet* is the disjoint union of a , b and c , while *string* is the disjoint union of *nullstring* and *nonnullstring*. Linking these types are the three functors shown in (3).

(3) <i>right</i>	:	<i>nonnullstring</i>	\rightarrow	<i>string</i>
<i>left</i>	:	<i>nonnullstring</i>	\rightarrow	<i>string</i>
<i>head</i>	:	<i>nonnullstring</i>	\rightarrow	<i>alphabet</i>

We subject these definitions to the path equation that *right left x* and *left right x* equal x for all non-null strings x .

A predicate in the corresponding modal language, using only the characteristic predicates of the types and the functors, might be: *head a* meaning the set of non-null strings whose first letter is a , *left head a* \wedge *right head c* to specify the context a_c , or *head c* \wedge *right(head a* \wedge *right(head b* \wedge *right null))*.

By the use of functors, we can move from one type to another, or from one item in a type to another item in the same type. Metaphorically, we will call the types joined by functors LOCATIONS, particularly when the type instances are only distinguished by functorial relationships with other types.

In a complex structure, like a string, the functors provide a method for interrogating nearby parts of the the structure within a predicate applying at a given position. By an appropriate choice of types and functors, complex feature structures and/or non-linear representations can be defined. For the sake of simplicity, the discussion in the remainder of this paper will be restricted to strings constructed using the types and functors defined above.

Constraints in a modal theory

In model-theoretic terms, a constraint is any well-formed expression in the language to which an interpretation is attached. Phonologists also use the term, usually intending universal application. It will be used here for a single predicate applying at a particular location in structure.

As an example of a constraint, consider front vowel harmony in Turkish¹. Informally, we can write this constraint as *if the last vowel was front, so is the current one*. In the format of a phonological rule, this might be written as $[+front]C^*V \Rightarrow [+front]$, where C^* stands for zero or more consonants. F is used to represent the disjunction of all of the front vowels.

$$(4) \text{ Left} = (\text{left head } C \wedge \text{left Left}) \vee \text{left head } F$$

$$\text{Constraint} = \text{head } F \vee \neg \text{Left}$$

In (4) the left context is abstracted into a named predicate called *Left*. This is because the left context iterates over consonants. This iteration appears in the definition of *Left* as the recursive call: if the immediate left segment is a consonant, move left and check again. *Left* succeeds immediately if the immediate left segment is a front vowel.

Note the the predicate defined here imposes no restrictions at all on where it applies except that it be a non-null string. On the other hand, it only applies at the current location in structure. The relationship between constraints and locations is the topic of the next section; first in the discussion of features, and then in the prioritisation of default feature assignment.

Features, Underspecification and Defaults

The question arises as to what basic predicates should be used in defining the lexical specification of phonological items. Lexical specifications in phonology are traditionally built from binary features. While the feature values usually correspond to a priori predicates, there is no reason why a feature cannot be defined for an arbitrary predicate: ϕ defining the feature $[+\phi]$ everywhere that ϕ is true and $[-\phi]$ everywhere that ϕ is false. This section includes discussion of two kinds of feature system here: A PRIORI and EXCEPTION-MARKING.

A priori features

Traditionally, the choice of features is made a priori (an A Priori Feature System — APFS). This does not mean that phonologists do not select their feature sets to suit their problems, rather that they do not approve of doing so. Instead, acoustic or articulatory grounds

¹Turkish has eight vowels, **a**, **e**, **i** the back version of **i**, **o** and its front correlate **ö**, and **u** and the corresponding front vowel **ü**.

are sought for a universal set of features which will serve for all analyses.


Furthermore, features in traditional systems are context free. The predicates defining the features do not make reference to neighbouring structures, such as the segment to the right or the left, in order to determine the feature polarity in a given position. Feature values depend only on the segment at that position in the string.

Continuing to draw our examples from Turkish vowels, *front* can be thought of as the predicate *head* ($e \vee i \vee \ddot{o} \vee \ddot{u}$). This predicate is context-free: there are no uses of the functors *left* and *right* in the definition. We can define the feature values $[+front]$ and $[-front]$ as holding at each non-null position in the string where *front* is true and false respectively.


Exception-marking features

A more adventurous feature system brings context together with the local segmental value to define its features. The question arises as to which predicates from this wider range should be chosen. The principle of Epicurus (Asmis, 1984) suggests that no choice should be made until direct evidence is adduced. In this domain the evidence comes in the form of a constraint on phonological structure. So, if it appears that ϕ is an interesting constraint on phonological structure, then $[\pm\phi]$ should be used as a feature. This choice is less ad hoc than introducing new predicates a priori.

As an example of this kind of feature assignment, consider the constraint (4) applied to the word *seviyorum* *I like (cts)*, which has the structure shown in (5).



The features assigned by the constraint are shown in (6). For clarity, the segments and *head* functors are not shown. To make the clearer, the positive and negative feature marks are shown as ticks and crosses respectively.



In only one case does this feature assign a negative value, ie. there is only one exception to the constraint in this word. This exception is the occurrence of the back vowel **o** after the front vowel **i**.

The segments themselves provide non-arbitrary context-free predicates which can be used as features. For example, we could define a feature $[\pm a]$ which is true if and only if *head a* is true.

These kind of feature systems are called EXCEPTION-MARKING FEATURE SYSTEMS (EMFSs) because it is exceptions to identified constraints which define all but the most basic features.

Underspecification

In EMFSs the number of features is likely to be much larger than in traditional systems. On the other hand, each of the features correspond to either a segment or a phonological constraint or a segment, so the system as a whole is ontologically simpler than a APFS. Nevertheless, unless some method of compression is used, EMFSs will demand verbose lexical forms. Two types of compression are familiar to, though seldom distinguished by, phonologists: redundancy and defaults². In terms of model theory the distinction is clear. Redundancy rules have no effect on the interpretation function I , while defaults modify it. This section discusses underspecification that eliminates redundancy. The next section discusses defaults.

A predicate ϕ is FULLY SPECIFIED FOR another predicate ψ if either ψ is more specific than ϕ , that is, $I(\phi) = I(\phi) \cap I(\psi)$, or ψ contradicts ϕ , $I(\phi) \cap I(\psi) = \emptyset$. A FULLY SPECIFIED predicate is one which is fully specified for all other predicates.

Intuitively, a fully specified predicate is one which is indivisible. There is no extra restriction which can be imposed which will make it more specific; it can only be contradicted. If ϕ is a fully specified predicate, then there is no point in adding further information to it.

If the interpretation function I is computable, then each feature value at each position in a fully-specified structure can be calculated. If the conjunction of the feature predicate with the structure has a null interpretation, then the feature is false, otherwise it is true. Consequently, so long as a predicate remains fully specified, any feature specifications which are removed from it can be recovered.

In APFSs, the constraints associated with features will not be very interesting. When the features are contextual constraints, however, regaining the full specification amounts to a process of phonological derivation albeit one of simultaneous application of constraints.

Let us utilise the Turkish vowel set for another example. Suppose each vowel is assigned a feature, and so is the vowel harmony constraint, (4). For each vowel, \checkmark marks the presence of the vowel, \times its absence. The same symbols mean the satisfaction of a constraint or its failure. Table (7) shows redundant feature specifications with a box around them. The example word is *severim* *I like*. Features for the consonants are not shown for the sake of brevity.

(7)

	s	e	v	e	r	i	m
Constraint (4)	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark
a	\times	\times	\times	$\boxed{\times}$	\times	$\boxed{\times}$	\times
e	\times	\checkmark	\times	\checkmark	\times	\times	\times
ı	\times	\times	\times	$\boxed{\times}$	\times	$\boxed{\times}$	\times
i	\times	\times	\times	\times	\times	\checkmark	\times
o	\times	\times	\times	$\boxed{\times}$	\times	$\boxed{\times}$	\times
ö	\times	\times	\times	\times	\times	\times	\times
u	\times	\times	\times	$\boxed{\times}$	\times	$\boxed{\times}$	\times
ü	\times	\times	\times	\times	\times	\times	\times

Note that this is not the only possible selection of redundant specifications. If the vowel feature specifications are regarded as primary and non-redundant, then the constraint feature values can all be regarded as redundant.

At this point we can define the declarative phonological formalism we are evaluating. It is an EMFS with redundant features removed, called Exception Theory (ET).

Defaults

Identifying fully specified predicates allows us to compress representations by removing predictable specifications from predicates. This compression method can be enhanced by modifying the interpretation function so that more predicates are fully specified.

A DEFAULT is defined in terms of a special predicate which will not need to be specified in individual representations. A representation will be conjoined with the default predicate unless it is already fully specified for it.

There may be a number of default predicates in a default system. For this reason the formal definition of the effect of defaults on the interpretation function has the recursive structure shown in (8).

$$(8) \quad I_{\Delta, \delta}(\phi) = \begin{array}{l} I_{\Delta}(\phi) \\ \text{if } \phi \text{ is fully specified for } \delta \text{ wrt } I_{\Delta}, \\ \text{or } I_{\Delta}(\phi) \cap I_{\Delta}(\delta) \text{ otherwise.} \end{array}$$

Each default predicate specifies its action at only one position in the structure. If the default is to apply at many positions in a structure, more default predicates must be added to cover each position in the structure.

For example, take the default predicate δ to be the feature [-front] equivalent to the predicate *head* ($a \vee ı \vee o \vee u$). Let ϕ be the partial specification for *kızları* (*her*) *girls* in which each vowel is underspecified for the feature front. Then the interpretation $I_{[-\text{front}]}(\phi)$ of ϕ subject to the default δ applied at the location *a* contains only the four forms *kızları*, *kızları*, *kızları* and *kızları*. Forms such as *kızları* are ruled out by the default at the *a* position.

To make the same default restriction at the other vowels in the word, we would need to other defaults such as *left left* δ and *right right* δ .

²Calder & Bird (1991) make this distinction using the (IPSG-like terms feature-cooccurrence restrictions (FCRs) and feature-specification defaults (FSDs).

Default ordering

Applying defaults is not necessarily commutative. One default may preclude the action of another. Consider the case where two feature values [-front] and [+front] are imposed as defaults to the completely unspecified predicate *true*. Because *true* is not fully specified for either [-front] or [+front], these defaults add specifications to the predicate: $I_{[-front]}(true)$ is $I([-front])$ while $I_{[+front]}(true)$ is $I([+front])$. But [-front] is fully specified for [+front] (and vice-versa), so adding [-front] (or [-front]) as a default will have no effect on the interpretation. Thus the two orderings of the defaults produce conflicting interpretations.

$$\begin{aligned}
 (9) \quad & I_{[+front][-front]}(true) \\
 &= I_{[-front]}([+front]) \\
 &= I([+front]) \\
 &\neq I([-front]) \\
 &= I_{[+front]}([-front]) \\
 &= I_{[-front][+front]}(true)
 \end{aligned}$$

Since the two orderings produce different results, a decision about the ordering of defaults must be made.

Default Ordering Schemes

Ordering

Defaults need to be ordered. There are a number of ways that the ordering of groups of defaults can be specified. Three of these are presented here.

Ordering by feature

One method for ordering defaults is to order the features they instantiate. We begin with an ordering on the features, so that, for example, feature [+F] has higher priority than feature [+G], in symbols $[+F] > [+G]$. This ordering on features, can then be extended to an ordering on defaults specified with those features.

Suppose p and q are paths in string structure, composed of sequences of *left* and *right* functors. Then for any defaults filling in predicates $\delta = p[+F]$ and $\epsilon = q[+G]$, δ is ordered before ϵ if and only if [+F] has higher priority than [+G].

Suppose a language is analysed as imposing a higher priority default that front vowels cannot occur after round vowels. Assume that the defaults insert the features [+front] and [+round] in all positions. Given a form $kVtV$ where V represents the completely uninstantiated vowel, there are two different instantiations depending on the ordering of the two features. If the [+front] default applies first, then the resulting form will be $k[+front]t \begin{bmatrix} + \text{ front} \\ + \text{ round} \end{bmatrix}$. If, on the other hand, the [+round] default applies first, the derived form will be $k[+round]t \begin{bmatrix} + \text{ front} \\ + \text{ round} \end{bmatrix}$.

Ordering by failure count

Another approach orders defaults instantiating the same feature in different positions. The preferred default minimises the number of contradictions to the default feature value.

Suppose the default feature value to be ordered is [+F]. The failure count default ordering mechanism uses a default predicate for each possible number of exceptions. The predicates, δ_i , are defined in (10).

$$\begin{aligned}
 (10) \quad & \delta_i = \bigvee_{j+k=i} \text{right } \delta_j \wedge \delta_k \\
 & \delta_0 = \text{left } (null \vee \delta_0 \wedge [+F]) \\
 & \delta_i = \text{left } (null \vee \delta_i \wedge [+F] \vee \delta_{i-1} \wedge [-F]) \\
 & \delta_0 = \text{right } (null \vee \delta_0 \wedge [+F]) \\
 & \delta_i = \text{right } (null \vee \delta_i \wedge [+F] \vee \delta_{i-1} \wedge [-F])
 \end{aligned}$$

If δ_i is compatible with a predicate ϕ , then there is a fully-specified restriction on ϕ which has no more than i occurrences of [-F]. The ordering on the defaults is imposed by requiring that for any feature [+F_{*i*}], with the corresponding predicate δ_i , δ_i has priority over δ_j iff $i < j$.

Suppose we already have a number of higher priority constraints on stress: that it can only be assigned once and in only one position within a syllable, and that consecutive syllables cannot be stressed. Collapsing the representation of syllables into a single symbol σ for convenience, table (11) gives the assignment of stress to a number of partially specified representations. The default feature is [+Stress], and this is applied to minimise the number of failures.

ϕ [\pm Stress]									
After defaults	+	-	+	-	+	-	+	-	-
or	-	+	-	+	-	+	-	+	-
Location	σ	σ	σ	σ	σ	σ	σ	σ	σ
ϕ [\pm Stress]									
After defaults	+	-	+	-	+	-	+	-	+
Location	σ	σ	σ	σ	σ	σ	σ	σ	σ
ϕ [\pm Stress]									
After defaults	-	+	-	+	-	+	-	+	-
Location	σ	σ	σ	σ	σ	σ	σ	σ	σ

Ordering by position

Another possibility is to order defaults by how far away from the starting position they specify their features. There are two simple ways of relating distance to priority: closer means higher priority, or further away means higher priority.

The formal definitions for this kind of default ordering are straightforward. Suppose, once again, that [+F] is the feature value to be filled in by the defaults. Now, δ_i will denote the specification of a default value at a distance of i functors to the left, or i to the right of the starting position.

$$\begin{aligned}
 (12) \quad & \delta_i = \text{right } \delta_i \wedge \delta_i \\
 & \delta_0 = [+F] = \delta_0 \\
 & \delta_{i+1} = \text{left } \delta_i \vee \text{null} \\
 & \delta_{i+1} = \text{right } \delta_i \vee \text{null}
 \end{aligned}$$

To prefer near defaults, prefer δ_i over δ_j when $i < j$. For far defaults, do the reverse.

Directional default preferences mimic the application of phonological rules in a left-to-right or right-to-left direction. Using this ordering, directional defaults can restrict some structures which the counting defaults cannot. Consider once again the stress assignments by defaults in table (11). Instead of simply trying to maximise the number of stresses, assume that the starting position is the left end of the word, and that near stresses are given priority. Under this system of defaults, the first of the three underspecified representations is rendered more specific, while the other two make the same restriction. These results are shown in table (13).

(13)

ϕ [\pm Stress]	+	-	+	-	+	-	+	-	-
After defaults	σ	σ	σ	σ	σ	σ	σ	σ	σ
Location									
ϕ [\pm Stress]		-							
After defaults	+	-	+	-	+	-	+	-	+
Location	σ	σ	σ	σ	σ	σ	σ	σ	σ
ϕ [\pm Stress]		+							
After defaults	-	+	-	+	-	+	-	+	-
Location	σ	σ	σ	σ	σ	σ	σ	σ	σ

Three Theories

Underspecification Theory

Within the framework given above, it is possible to define a form of Underspecification Theory. What is described here is not precisely the Underspecification Theory of Archangeli (1984), differing in that the structures described are linear and segmental. This is, however, not a necessary limitation of the framework, and the definition of underspecification theory presented here could be applied to autosegmental representations if suitable types and functors were defined for them.

In UT, lexical specifications are made in terms of an a priori fixed set of features. For example, Archangeli & Pulleyblank (1989) use the four features [\pm high], [\pm low], [\pm back] and [\pm ATR] to describe seven Yoruba vowels. All lexical specifications of vowel quality are assumed to involve specifications for some subset of these features.

In the lexical specifications, redundant information is left unmarked. The Yoruba vowel *a* does not need to be marked for any feature other than [+low], because there is only one vowel which is [+low]. Consequently, the feature values [+back], [-high] and [-ATR] are all redundant.

In UT, redundant features are filled by rule. Special constraints, such as the Redundancy Rule Ordering Constraint (Archangeli, 1984:85) ensure that redundancy rules apply before the features they instantiate are referred to. Furthermore, these constraints apply as often as necessary (Archangeli & Pulleyblank, 1989:209-210). This has the same effect as the automatic specification of redundant feature values in the current framework.

Only one type of feature value is ever lexically specified in UT. Opposite feature values are filled in by default rules. This allows the feature specifications for some segments to be subspecifications of those for other segments.

Apart from the context-free features used in lexical specifications, there are also context-sensitive constraints which are regarded in UT as fully-fledged phonological rules. For example, the Yoruba vowel harmony rule can be summarised as *a vowel on the left of a [-ATR] vowel will also be [-ATR]*. Regularity to this constraint in one position may conflict with regularity in another position. In UT, the defaults associated with such constraints are ordered by position: Yoruba vowel harmony applies right-to-left in the sense that constraint applications further from the beginning of the word have higher priority.

This directionality is not the only ordering of defaults. As it happens, there are no [+high] vowels in Yoruba which are also [-ATR]. Consequently, the default rule marking vowels as [+high] can conflict with the default that spreads [-ATR]. In the analysis of Archangeli & Pulleyblank the [+high] default is ordered first. All defaults constructed from the one feature have priority over all defaults built on the other.

The general structure of UT, therefore, is to have an a priori limited set of features for lexical specification and a set of defaults for these features and for constraints. The defaults associated with each feature or constraint are ordered by position.

Optimality Theory

Optimality Theory (Prince & Smolensky, 1993) is apparently a very different theory, but, when classified in terms of its use of defaults, is actually quite similar.

In contrast to UT, OT is deliberately vague about underlying representations. Instead of discussing the manipulation of representations directly, OT refers to their interpretations, terming them CANDIDATE SETS.

Constraints in OT apply exactly like defaults. If they can be imposed without resulting in a contradiction (empty candidate set), then they are. Each constraint imposes a set of defaults, and these are primarily ordered by an extrinsic ordering placed on the constraints. If any two defaults pertaining to two constraints conflict, the default of the higher order constraint is preferred.

As with UT, there is the possibility that the imposition of the the same constraint at different locations will conflict. Rather than ordering these defaults by position, they are ordered by the number of exceptions to the constraint that they allow. If there is a candidate form with a certain number of exceptions, all candidates with more exceptions will be eliminated by the default. This ordering on defaults is the ORDERING BY FAILURE COUNT described earlier.

Exception Theory

In contrast to the other two, more standard, phonological theories, Exception Theory does not use defaults. In ET, each lexical form is fully specified, and any feature in it may be removed so long as this property is preserved.

The set of features includes a feature for each segment type, and a feature for each constraint. While this results in a large set of features, underspecification of redundant features means that many feature specifications may be eliminated. Nevertheless, there will be more feature specifications needed in ET than in, for example, UT, because of the absence of default values.

On the other hand, because ET uses no defaults, there is no need for any form of constraint or rule ordering. All features have an immediate interpretation through the interpretation function, and so a minimum of computation is needed to identify the denotation of a representation.

Summary

Table (14) summarises the attributes of the three theories. UT and OT are primarily distinguished by the use of different methods to order defaults built from constraints. ET differs in that it does not use defaults at all.

(14)

	UT	OT	ET
A priori features	✓	×	×
Defaults	✓	✓	×
By Feature	primary	primary	×
By Failure Count	×	secondary	×
By Position	secondary	×	×

Discussion

Early in this paper, four desiderata for morphophonological theories were introduced. This section considers whether using defaults is advantageous with respect to these desiderata.

Conciseness

The first desideratum sought concise lexical representations for morphemes. Since default-based theories can also exploit underspecification of redundant feature values, they are at least as concise as non-default theories. If there are ever contrastive feature specifications, then they are more concise, allowing one side of the contrast to be left as a default value to be instantiated.

Note that the concept of conciseness which is being used here is feature-counting, not an information-theoretic measure. In a direct application of information theory, contrasting a [+F] feature value with whitespace carries as much information as contrasting it with [-F]³.

Abstracting and recovering morphemes

Defaults also provide advantages in abstracting morpheme representations from which allomorphs can be

³It may be possible, nevertheless, to provide an information theoretic basis for the feature-counting notion by couching the feature specifications in a suitable descriptive language.

recovered. As well as making representations more concise, using defaults allows more allomorphs to be brought together within a single phonological representation. As there are no feature changing rules in the framework, all feature values in the abstract representation must survive to the surface in each allomorph. Conversely, the abstract representation can only contain feature specifications common to all of the allomorphs. So the upper bound on feature specifications for the abstract morpheme is the intersection of the featural specifications for all of the allomorphs of the morpheme.

As an example, consider four allomorphs of the Turkish second person plural possessive suffix: *-ınız*, *-iniz*, *-unuz* and *-ünüz*. If the vowels are specified with the three features [±front], [±round] and [±high], then the intersection of the specifications of the four allomorphs is the sequence [+high]n[+high]z.

While it is always possible to form abstract representations by intersecting feature values (the second desideratum), there is no guarantee that the allomorphs will be readily recoverable (third desideratum). If they are not recoverable, then there is no single featural generalisation which captures the phonological structure of the morphemes.

One important question is whether defaults allow recoverable generalisations about a greater range of morphemes than non-default representations. The answer is yes. If the morphological alternations is one-dimensional, then there is no difference between having defaults and not. Suppose δ is a default predicate, and, equally, an exception feature. If all allomorphs are specified [+ δ] then the abstraction will share this feature, and so the default does not need to apply. Similarly if all allomorphs are specified [- δ], so will the abstract forms be, and the default cannot apply. If the allomorphs vary in their specification for [± δ], then the abstraction will not have include a specification for this feature. Consequently, the default will specify [+ δ] when the correct value is [- δ], and so not fail to produce the correct result. In the non-default interpretation, the representation is never fully specified.

On the other hand, if the morphological alternations form a two-dimensional paradigm, then it is possible that the paradigm might be decomposable into morphemes only with the use of defaults. Suppose, once again, that δ is a default predicate and exception feature. The default feature value is [+ δ]. Suppose further, that there is a paradigm with the feature specification for [± δ] shown in (15).

(15)

	[- δ]	[0 δ]
[- δ]	[- δ]	[- δ]
[0 δ]	[- δ]	[+ δ]

The margins show the 'morphemes' extracted by intersecting the feature values. The conjunction of the two [0 δ] specifications is not fully specified for δ , and so its direct interpretation does not recover the corresponding

component of the paradigm. If, however, the default [+δ] is applied, the full specification of the paradigm is recovered.

So it is possible to have paradigms where the morphological components cannot be assigned common phonological representations without the use of defaults⁴.

A priori specifications

The final desideratum is the avoidance of a priori information in a model. UT makes use of an a priori set of features for lexical specification. As other generalisations in the formalism are only visible insofar as they affect the values of these features, this limits the possible constraints which can be identified. This is the reason why vowel harmonies such as that of Nez Perce are so problematic for phonologists⁵: the sets of vowels used in the harmony do not have a neat definition in terms of traditional features.

Greater claims about a priori features are made in OT. Prince & Smolensky (1993:3) state that *constraints are essentially universal and of very general formulation ... interlinguistic differences arise from the permutation of constraint-ranking*. In other words, all of the predicates which define features in OT are prior to the analysis of an individual language.

In ET, very little is assumed a priori. Any constraint which captures interesting phonological generalisations about the phonology defines a feature which can be used to specify structure. Because ET does not use defaults, it need not be concerned with ordering constraints, only with finding them. Consequently, interlinguistic differences can only result from distinct sets of constraints.

Conclusion

In this paper I have presented a rigorous framework for characterising theories that use defaults with phonological structure. The framework provides a straightforward characterisation of Underspecification Theory and Optimality Theory in terms of the action of defaults.

Using this framework, I have shown that non-default theories cannot be sure of capturing all of the generalisations which are available to default theories. For this reason, the non-default constraints learnt by programs such as those described by Ellison (1992, forthcoming), are not as powerful for morphophonological analysis as default-based theories. Furthermore, defaults lead to more concise, and consequently preferable, lexical representations.

⁴If general predicates are permitted for specifying morphemes, rather than just featural specifications, the distinction between default and non-default systems disappears. If the entries in the paradigm are ξ_{ij} , define α_i to be $\bigvee_j \xi_{ij}$ and β_j to be $\bigwedge_i (\xi_{ij} \vee \neg \alpha_i)$. Then, so long as the α_i are distinct (which will be the case if the ξ_{ij} are all distinct), then the paradigm will be fully recoverable without defaults.

⁵Anderson & Durand (1988) discuss some of this literature.

The question, therefore, is how to enhance the learning algorithms to involve the use of defaults. The introduction of defaults means that constraints must be ordered; so learning must not only discover the right constraint, it must assign it a priority relative to other constraints. This makes the learning task considerable more complicated. However difficult a solution for this problem is to find, it will be necessary before machine-generated analyses can be sure of competing successfully with man-made analyses.

Acknowledgements

This research was funded by the U.K. Science and Engineering Research Council, under grant GR/G-22084 *Computational Phonology: A Constraint-Based Approach*. I am grateful to Richard Sproat and Michael Gasser for their comments on an earlier version of this paper.

References

- Anderson, J. & Durand, J. (1988). Vowel harmony and non-specification in Nez Perce. In H. van der Hulst & N. Smith (Eds.), *Features, Segmental Structure and Harmony Process (Part II)* (pp. 1-17). Foris.
- Archangeli, D. (1984). *Underspecification in Yawelmani Phonology and Morphology*. PhD thesis, Massachusetts Institute of Technology.
- Archangeli, D. & Pulleyblank, D. (1989). Yoruba vowel harmony. *Linguistic Inquiry*, 20, 173-217.
- Asmis, E. (1984). *Epicurus' Scientific Method*. Ithaca, NY: Cornell University Press.
- Calder, J. & Bird, S. (1991). Defaults in underspecification phonology. In S. Bird (Ed.), *Declarative Perspectives on Phonology* (pp. 107-125). University of Edinburgh.
- Ellison, T. M. (1992). *The Machine Learning of Phonological Structure*. PhD thesis, University of Western Australia, Perth.
- Ellison, T. M. (1994). The iterative learning of phonological rules. Technical report (forthcoming), Cognitive Science, University of Edinburgh.
- Prince, A. S. & Smolensky, P. (1993). *Optimality Theory: Constraint Interaction in Generative Grammar*. Technical Report 2, Center for Cognitive Science, Rutgers University.