Chapter 8

Morphological complexity and Input Optimization

Michael Hammond
University of Arizona

In this paper, we examine morphological complexity through the lens of Input Optimization. We take as our starting point the dimensions of complexity proposed in Anderson (2015). Input Optimization is a proposal to account for the statistical distribution of phonological properties in a constraint-based framework. Here we develop a framework for extending Input Optimization to the morphological domain and then test the morphological dimensions Anderson proposes with that framework.

The dimensions we consider and the framework we develop are both supported by empirical tests in English and in Welsh.

1 Introduction

Anderson (2015) lays out a number of dimensions of morphological complexity, ways that we might evaluate how complex different morphological systems are, e.g. number of morphemes in the system, complexity of principles governing combinations of morphemes, complexity of exponence, complexity of allomorphy, etc.

These are clearly the right kinds of dimensions for evaluating the complexity of morphological systems, and we might be inclined to use them as part of a typology of morphology. However, if we adopt these as our dimensions for calculating complexity, what follows? As I learned from Steve Anderson years ago in graduate school, typology without implications is bad typology. (For discussion, see, for example, Anderson 1999.)

In this paper, I consider the implications of these dimensions of morphological complexity for the theory of Input Optimization (Hammond 2013; 2014; 2016). This theory develops a notion of phonological complexity which languages "attempt" to minimize statistically. To the extent that different phonological representations are complex, they are under-represented statistically. This complexity shows up in the markedness of surface representations and in the complexity of input–output mappings. For example, marked segments or syllable structures are under-represented compared with their less

Michael Hammond. 2017. Morphological complexity and Input Optimization. In Claire Bowern, Laurence Horn & Raffaella Zanuttini (eds.), *On looking into words (and beyond)*, 155–171. Berlin: Language Science Press. DOI:10.5281/zenodo.495444

marked counterparts. In addition, outputs that are distinct from their inputs are under-represented with respect to outputs that are identical with their inputs.

Interestingly, there are morphological effects as well, effects that sometimes work in the opposite direction. For example, initial consonant mutation in Welsh causes a mismatch between output and input, but is over-represented. I argue that this is because phonological complexity includes morphological mapping. Specifically, to the extent that morphological distinctions are not made in the surface form, a representation is more complex. This is formalized in OT-based terms using a constraint deriving from work by Kurisu (2001).

This general approach is supported by the facts of haplology, e.g. English adjectives in *-ly* like *weekly* not getting double-marked with adverbial *-ly* and plurals like *kings* not getting double-marked with genitive *-s*. The absence of double-marking means that a morphological distinction is not made on the surface; thus these cases are under-represented as expected.

These morphological cases beg the larger question: should phonological complexity be generalized further? Should there be a more general notion of morphological complexity, built on dimensions of the sort cited above, where forms that are more complex morphologically are statistically under-represented? In this paper, I pursue just this course, formalizing a notion of morphological complexity and then testing it with cases from English and Welsh with respect to the dimensions of morphological complexity identified in Anderson (2015).

The organization of this paper is as follows. I first review some of the dimensions of complexity presented in Anderson (2015). I then outline the theory of Input Optimization and a framework for a constraint-based theory of morphology that we can assess Input Optimization with respect to. With these in hand, we then consider the predictions made by the Input Optimization framework and turn to the English and Welsh data.

2 Dimensions of complexity

Anderson (2015) discusses a number of dimensions of morphological complexity and we will not review them all here.

We explicitly set aside those systemic dimensions that cannot distinguish options within a language. For example, Anderson cites the number of elements in the morphological system as a measure of its complexity. Thus, for example, if one language has one way of marking noun plurals and another has ten ways, we might think of the first as less complex. Input Optimization makes no predictions about systemic differences like these, as we will see in the next section, so we don't consider them any further.

Anderson cites the number of morphemes in a word as a dimension of complexity.[1] This can be taken in several ways. One possibility is that one might compare across different languages, which seems to be Anderson's intent. Another possibility though

[1] This, of course, begs the question of what is a morpheme, an issue at the forefront of much of Steve's own work.

would be to compare across words in the same language and we investigate this possibility below.

Another dimension Anderson identifies is whether the morphemes present in a word depend on each other in some way. We might think of this in two ways. Some morpheme may only occur when "licensed" by some other. Gender in Spanish is an example of this. If gender is marked on a noun, then that gender marking must be present in the plural as well, e.g. *mes+a+s* 'tables' table + feminine + plural. Another example might be verbal theme vowels in Romance; person/number marking depends on the presence of the theme vowel.

The other side of this coin would be cases where the presence of some morpheme blocks another. Haplology is an example of this. For example, adverbial *-ly* in English cannot occur on an adjective that already ends in *-ly*. Thus we have *happily*, but not **weeklyly*.

Anderson also distinguishes among morpheme types in terms of complexity. Simple prefixation or suffixation is less complex than circumfixion or infixation. Presumably, non-concatenative morphology like templatic operations, ablaut, umlaut, truncation are also more complex.

Lastly, Anderson cites the complexity of allomorphy as an instance of general morphological complexity. We take this to mean that a system is more complex when there is more allomorphy. We interpret allomorphy as generously as possible to include cases where the phonology seems to be involved, say, in the different pronunciations of the English plural *-s* as [s, z, əz], but also plurals that differ on some other basis, e.g. *geese*, *criteria*, *sheep*, etc.

Anderson treats some other possibilities as well, but the ones above are quite simple and can be examined within a single language. We list them together below.

1. Number of morphemes in a word
2. Principles of morphological combination, e.g. scope, haplology, etc.
3. Complexity of exponence, e.g. circumfixes, infixes, etc.
4. Complexity of allomorphy

In the following section, we review the Input Optimization proposal and sketch out the predictions it makes for these. Our interest in Input Optimization is that it provides a mechanism by which we can assess the dimensions of morphological complexity we've just considered.

3 Input Optimization

The problem that Input Optimization addresses is that certain phonological configurations occur less often than we might otherwise expect. For example, if we look at the distribution of stress on two-syllable adjectives, we see that adjectives with final stress like *alert* [əlɹ̩́t] or *opaque* [òpʰék] are less frequent overall. Strikingly, both are even less frequent when they occur prenominally.

Hammond (2013) argues that this effect is driven by the English Rhythm Rule (Liberman & Prince 1977; Hayes 1984). Certain stress configurations in English are avoided by shifting a primary stress leftward onto a preceding secondary. Thus we have alternations like *Mìnnesóta* vs. *Mínnesòta Míke*; *thìrtéen* vs. *thírtèen mén*; etc. When an adjective with final stress occurs in prenominal position, the relevant configuration is quite likely to occur. In addition to following context, there is a restriction on preceding context. With an adjective like *òpáque*, stress shift leftward is possible because of the preceding stress, e.g. *ópàque stóry*, but with an adjective like *alért*, such a shift is impossible and the offending configuration must surface, e.g. *alért pérson.* Both kinds of cases are statistically under-represented in English. Specifically, these configurations arise significantly less frequently than we might expect based on the overall distribution of adjectives with these stress patterns.

Input Optimization is a proposal to account for statistical skewings like these that occur in the phonologies of languages. The idea is developed in Hammond (2013; 2014; 2016). The basic idea is that markedness and faithfulness violations are avoided in the phonology so as to reduce the complexity of the phonological system. Input Optimization is a generalization of the notion of Lexicon Optimization Prince & Smolensky (1993):

(1) **Lexicon Optimization**:
Suppose that several different inputs $I_1, I_2, \ldots, I_n$ when parsed by a grammar G lead to corresponding outputs $O_1, O_2, \ldots, O_n$, all of which are realized as the same phonetic form Φ—these inputs are all *phonetically equivalent* with respect to G. Now one of these outputs must be the most harmonic, by virtue of incurring the least significant violation marks: suppose this optimal one is labelled O_k. Then the learner should choose, as the underlying form for Φ, the input I_k.

The idea is that if there are multiple ways to produce an output form consistent with the facts of a language, the learner chooses the input that produces the fewest constraint violations. There are no empirical consequences to Lexicon Optimization by itself. In fact, it is defined to apply only when there are no consequences.

To refine this into something we can use, we define a notion of *Phonological complexity* that applies to individual input–output pairings, but also to entire phonological systems. (The basic logic of this is that the complexity of a phonological system is proportional to the number of asterisks in its tableaux.)

We define the output/surface forms of a language as a possibly infinite set of forms.

(2) $O = \{O_1, O_2, \ldots, O_n, \ldots\}$

Each output form has a corresponding input:

(3) $I = \{I_1, I_2, \ldots, I_n, \ldots\}$

The phonology is comprised of a finite sequence or vector of constraints:

(4) $C = \langle C_1, C_2, \ldots, C_n \rangle$

Any input–output pairing, (I_i, O_i), then defines a finite vector of violation counts, some number of violations for each constraint earned by the winning candidate for that input.

(5) $\langle n_{C_1}, n_{C_2}, \ldots, n_{C_n} \rangle$

With these notions, Phonological Complexity (PC) is defined as follows:

(6) **Phonological Complexity** (PC)
The phonological complexity of some set of forms is defined as the vector sum of the constraint violation vectors for surface forms paired with their respective optimal inputs.

To produce a *relative* measure of PC given some set of *n* surface forms, divide the PC score for those forms by *n*.

Hammond (2016) exemplifies this with a hypothetical example of nasal assimilation. Imagine we have the forms in (7) we wish to compute the PC for. Given the inputs provided in column 2, we have the constraint violations for winning candidates in columns 3 and 4.

(7)

	Input	Output	NC	IO-Faith
a.	/on pi/	om pi		*
b.	/an ba/	am ba		*
c.	/un bo/	um bo		*
d.	/en do/	en do		
e.	/on ta/	on ta		
f.	/un ti/	un ti		
g.	/an ku/	aŋ ku		*
h.	/in ga/	iŋ ga		*
i.	/on ke/	oŋ ke		*
			0	6

The relative complexity of this first system is: $\langle 0, 6\rangle/9 = \langle 0, .66\rangle$. We can compare the system in (7) with the one below in (8). Here we have a different array of output forms, but the same logic for inputs and constraint violations.

(8)

	Input	Output	NC	IO-Faith
a.	/on pi/	om pi		*
b.	/an ba/	am ba		*
c.	/en do/	en do		
d.	/on ta/	on ta		
e.	/un ti/	un ti		
f.	/in di/	in di		
g.	/an ku/	aŋ ku		*
h.	/in ga/	iŋ ga		*
			0	4

The relative complexity of the second system is: $\langle 0, 4\rangle/8 = \langle 0, 0.5\rangle$, less than the first. As argued by Hammond (2016), this notion extends obviously to weighted constraint systems. For example, in a system with strict ranking, $\langle 0.1, 0.4\rangle$ is more complex than $\langle 0, 0.5\rangle$.

The proposal then is that all phonological systems are skewed to be less complex.

(9) **Input Optimization**
All else being equal, phonological inputs are selected that minimize the phonological complexity of the system.

Note that (9) alters the frequency of input–output pairings and does not change the input–output mapping of any particular form. For example, this principle prefers (8) to (7), though both systems contain the same pairings. The difference is in the relative frequency of the pairings that occur.

Our goal in this paper is to see if it is profitable to extend this system to include morphology. In point of fact, Hammond (2016) addresses this question partially in response to statistical effects in Welsh. In particular, Welsh initial consonant mutation is statistically over-represented when, based on what we have seen so far, we might have expected the opposite.

Consonant mutation in Welsh refers to a set of phonological changes that apply to initial consonants in specific morpho-syntactic contexts. For example, the *Soft Mutation* makes the following changes:

(10)

Orthographic		Phonological	
Input	Output	Input	Output
p	b	p	b
t	d	t	d
c	g	k	g
b	f	b	v
d	dd	d	ð
g	∅	g	∅
m	f	m	v
ll	l	ɬ	l
rh	r	r̥	r

There are many contexts where this occurs, e.g. after certain prepositions, direct object of an inflected verb, after certain possessives, feminine singular nouns after the article, etc. The following figure gives some examples after the preposition *i* [i] 'to'.

(11)

pen	[pʰɛn]	'head'	i ben	[i bɛn]	'to a head'
cath	[kʰaθ]	'cat'	i gath	[i gaθ]	'to a cat'
mis	[mis]	'month'	i fis	[i vis]	'to a month'
nai	[naj]	'nephew'	i nai	[i naj]	'to a nephew'
siop	[ʃɔp]	'shop'	i siop	[i ʃɔp]	'to a shop'

The chart above also includes examples of non-mutating consonants. Note that words with these occur in mutation contexts with no change.

The Input Optimization framework would seem to predict that mutation should be under-represented. After all, mutation entails a faithfulness violation and, all else being equal, the system is less complex to the extent that such violations are avoided. This, however, is not what occurs. Instead, we get over-representation in mutation contexts. Words that begin with consonants that can mutate are over-represented in mutation contexts compared with words that begin with consonants that do not mutate.

To capture this, Hammond (2016) proposes the (revised) Realize Morpheme constraint (12). This is a slight revision of a constraint that Kurisu (2001) motivates on other (non-statistical) grounds. This constraint basically militates for the expression of morphological information.

(12) Realize Morpheme (*revised*) (RM')
Let α be a *morpheme*, β be a morphosyntactic category, and F(α) be the phonological form from which F(α+β) is used to express a morphosyntactic category β. Then RM' is satisfied with respect to β iff F(α+β) ≠ F(α) phonologically.

With this in hand, the reason why Welsh mutation is over-represented is to reduce phonological complexity by minimizing violations of RM'.

The RM' constraint is also invoked by Hammond (2016) to account for haplology in English. We've already cited the fact that forms like **weeklyly* are blocked. Similarly, we find overt marking of the genitive in English does not occur on plural forms marked with *-s*; the genitive plural of *cat* is *cats'*, not something like **catses*. Both kinds of cases are statistically under-represented in English: they are avoided to minimize violations of RM'.

While RM' (12) does what's required, it begs the question of whether a more general version of PC is appropriate. In other words, beyond the effects of RM', do we expect Input Optimization to apply to morphology?

4 Constraint-based morphology

To assess this, we need a constraint-based theory of morphology. There have been a number of proposals over the years for how to deal with morphology generally in an OT-like framework. The earliest we know of are Russell (1993; 1995); Hammond (2000), but see Aronoff & Xu (2010); Xu & Aronoff (2011) for more recent and fuller proposals. A full-on theory of this sort is well beyond the scope of this paper, but let's lay out what such a theory might look like, at least in sufficient detail so we can assess whether it makes the right predictions about Input Optimization.

Let us assume that morphology—like phonology—is a constraint-based system mapping inputs to outputs. Inputs are denuded of any morphological marking, but have sufficient featural information so that we can evaluate whether morphologically marked candidate forms satisfy relevant constraints. For example, we might imagine that plural marking in English comes about by taking a stem marked [+plural] and adding various affixes or performing other operations that do or do not express that feature. The idea is that the syntax provides a featurally complex object that the morphology can then interpret. Morphological operations like affixation, reduplication, mutation, etc. add features which do or do not match those required by the syntax. Following is a schematic partial tableau to give a sense of this.

(13)

	$/\begin{bmatrix}\text{cat}\\ \text{+pl}\end{bmatrix}/$	...
	$\begin{bmatrix}\text{cat}\\ \text{+pl}\end{bmatrix}$	
☞	$\begin{bmatrix}\text{cat}\\ \text{+pl}\end{bmatrix} + \begin{bmatrix}\text{-s}\\ \text{+pl}\end{bmatrix}$	
	$\begin{bmatrix}\text{cat}\\ \text{+pl}\end{bmatrix} + \begin{bmatrix}\text{-ed}\\ \text{+past}\end{bmatrix}$	

We would want constraints that force the correct morphological operation to take place. Presumably there would be one or more constraints that enforce a correspondence between the features required by the stem and the features offered by any affixes or other changes; to the extent that those don't match, we would have violations. For convenience, let's call this Features (Fs). The RM' constraint above, or constraints that get the same effects, should fall in this class.

We also need constraints that militate against gratuitous morphological operations. Some of this might be achieved by featural correspondence imposed by Fs, but we surely need something to account for the relative markedness of morphological operations generally. Perhaps something like this:

(14) *Ablaut ≫ *Infix ≫ *Prefix ≫ *Suffix

The basic idea is to posit constraints that militate against *any* morphological operation. These constraints are ranked with respect to each other, presumably in a universal fashion. This hierarchy would then be interleaved with the Fs constraint. For example, we might have:

(15) *Ablaut ≫ *Infix ≫ Features ≫ *Prefix ≫ *Suffix

The effect of such a ranking is that the featural needs of a stem can be met by prefixation and suffixation, but not by other operations.

This system is woefully incomplete and, in its present form, cannot do justice to the full range of effects we see in morphological systems. See, for example, Anderson (1982;

1992). It is, in some ways, quite similar to these proposals in treating affixation as an instance of more general morphological operations that interpret syntactically-motivated features. However, our goal here is not to develop a full-on constraint-based morphological theory. Rather, the point is to build enough of such a theory so that we can test Input Optimization with respect to the dimensions of morphological complexity identified above.

Let's now return to our dimensions and consider one by one what our theoretical skeleton in conjunction with Input Optimization predicts. First, we have the number of morphemes. All else being equal, the system certainly as developed militates for as few morphemes, or other morphological operations, as possible. Additional morphology entails violations of the constraints in (14) and Input Optimization predicts these should be avoided statistically.

The second dimension of complexity refers to principles of morphological combination. The system we've developed says nothing (so far) about the licensing side of this, but it does address morphological haplology. To the extent that haplology occurs, it entails violations of RM' (12) and of Fs. Previous work cited above has already established that Input Optimization applies in these cases.

The third dimension is the complexity of exponence, that certain morphological operations are intrinsically more complex than others. This is captured by the ranking, e.g. in (14). We expect morphologies to be statistically skewed against violations of the higher-ranked constraints.

The fourth dimension is complexity of allomorphy, allomorphy that is a consequence of phonology or morphophonology like English plural [s, z, əz], but also plurals that differ on some other basis, e.g. *geese*, *criteria*, *sheep*, etc. The phonological cases fall under the core Input Optimization proposal. In fact, Hammond (2013) shows statistical skewing for English plural and past allomorphy in just the expected directions. The other case cited is also accommodated by the proposal. Internal modifications like *geese* or truncation+suffixation like *criteria* violate higher-ranked constraints than simple plural suffixation; hence they should exhibit under-representation. Similarly, plurals with no change like *sheep* should violate RM' and Fs and be under-represented.

Summarizing, a constraint-based morphological theory of the sort sketched out, in conjunction with Input Optimization, makes the following predictions:

(16)
1. Words should have fewer morphemes.
2. Haplology should be avoided. (This has already been established by Hammond 2016.)
3. More marked morphological operations (per the hierarchy above) should be avoided.
4. Morphophonology should be avoided. (This has already been established by Hammond 2013.)
5. Ablaut, umlaut, truncation, etc. should be avoided. (This is essentially the same as #3 above.)
6. Zero-marking should be avoided.

We must therefore examine #1, #3/5, and #6 empirically. In the next sections, we look at all three cases with data from English and Welsh.

5 Number of morphemes

The first prediction of Input Optimization applied to our toy constraint-based theory of morphology is that a form is more complex if it has more morphemes. This is a bit tricky to test. In many cases, having fewer morphemes is not necessarily the less complex option. For example, consider the plural form *sheep* which lacks an overt plural suffix. Is this less complex than a form like *cat+s*? Probably not. The most reasonable analysis given the framework above is that the plural *sheep* surfaces with an undischarged plural feature. On that view, it is not clearly less complex than a form like *cat+s*.

We might also think of strong verb forms like *spoke*, as compared with *look+ed*. Here, however, it would be a mistake to view *spoke* as having fewer morphemes than *look+ed*. Rather, there is some operation, perhaps mostly lexical, for creating or selecting strong verb forms when available. Presumably, this would add to the complexity of *spoke*.

To find a case without these alternative analyses, we turn to Welsh plurals. Welsh has a number of ways of forming plurals. For example:

(17)

Singular			Plural	
ysgol	[ə́sgɔl]	‘school’	ysgolion	[əsgɔ́ljɔn]
cyfarfod	[kʰəvárvɔd]	‘meeting’	cyfarfodydd	[kʰəvarvɔ́dɨð]
cynllun	[kʰə́nɬɨn]	‘plan’	cynlluniau	[kʰənɬɨ́njaɯ]
problem	[pʰrɔ́blɛm]	‘problem’	problemau	[pʰrɔblɛ́maɯ]
panel	[pʰánɛl]	‘panel’	paneli	[pʰanɛ́li]
pwnc	[pʰʊ́ŋk]	‘subject’	pynciau	[pʰə́ŋkʰjaɯ]
angen	[áŋɛn]	‘need’	anghenion	[aŋhɛ́njɔn]
gorchymyn	[gɔrχə́mɨn]	‘order’	gorchmynion	[gɔrχmə́njɔn]
alarch	[álarχ]	‘swan’	elyrch	[ɛ́lɨrχ]
castell	[kʰástɛɬ]	‘castle’	cestyll	[kʰɛ́stɨɬ]

Note that there are different suffixes and stem changes.

There is another class of nouns, however, where it is the singular that is marked rather than the plural. The singular is always marked with either *-yn* in the masculine gender or *-en* in the feminine gender. For example:

(18)

Singular			Plural	
mochyn	[mɔ́χɨn]	'pig'	moch	[mɔ́χ]
blewyn	[bléwɨn]	'hair'	blew	[bléw]
morgrugyn	[mɔrgrɨ́gɨn]	'ant'	morgrug	[mɔ́rgrɨg]
marworyn	[marwɔ́rɨn]	'ember'	marwor	[márwɔr]
eginyn	[ɛgínɨn]	'sprout'	egin	[ɛ́gɪn]
mefusen	[mɛvɨ́sɛn]	'strawberry'	mefus	[mɛ́vɨs]
coeden	[kʰɔ́ɰdɛn]	'tree'	coed	[kʰɔɰd]
derwen	[dɛ́rwɛn]	'oak'	derw	[dɛ́ru]
madarchen	[madárχɛn]	'mushroom'	madarch	[mádarχ]
moronen	[mɔrɔ́nɛn]	'carrot	moron	[mɔ́rɔn]

There are some blended cases as well, where nouns marked for the singular co-occur with stem changes or take plural suffixes as well. For example:

(19)

Singular			Plural	
merlyn	[mɛ́rlɨn]	'pony'	merlod	[mɛ́rlɔd]
oedolyn	[ɔɰdɔ́lɨn]	'adult'	oedolion	[ɔɰdɔ́ljɔn]
taten	[tʰátʰɛn]	'potato'	tatws	[tʰátʰʊs]
(a)deryn	[(a)dɛ́rɨn]	'bird'	adar	[ádar]
gwreiddyn	[gwréjðɨn]	'root'	gwraidd	[gwrájð]
deilen	[déjlɛn]	'leaf'	dail	[dájl]
hwyaden	[hujádɛn]	'duck'	hwyaid	[hújajd]
cneuen	[kʰnéɰɛn]	'nut'	cnau	[kʰnáɰ]

The existence of the pairs where the singular is marked instead of the plural allows us to test the number of morphemes prediction without the problems of the English cases above.[2] On one hand, we take nouns which mark the plural with *-(i)au*, the most frequent plural suffix, and no associated stem changes. On the other, we take nouns which mark the singular with *-en* or *-yn*, and no associated stem changes or plural marking. In other words: *problem/problemau*, etc. vs. *coeden/coed*, etc. What we're interested in is whether there is a difference in the relative frequency of singular and plural forms in the two classes as a function of whether the form has an extra morpheme. Since we have both types of marking in Welsh, we can do this independent of the relative frequency of singulars and plurals in the language.

Individual lexical items have different frequencies of occurrence, so we must equalize for this. We therefore take the ratio of singular to plural as a measure of the relative frequency of singular and plural. Since this is a ratio, it abstracts away from the overall frequency of each pair.

[2] One might counter that it's possible to treat these as instances of subtractive morphology. There are two arguments against this. First, the singulatives are always marked with the suffixes *-yn* or *-en* (depending on gender). Second, there are cases where nouns end in these phonetic sequences where they are not suffixes. In these cases, normal plural formation occurs. For example: *emyn*/*emynau* 'hymn', *terfyn*/*terfynau* 'boundary', *ffenomen*/*ffenomenau* 'phenomenon', *awen*/*awenau* 'inspiration, muse', etc.

For this investigation, we use the CEG corpus (Ellis et al. 2001). This is a tagged written corpus of 1223501 words. For each word form, it also includes lemmas, so it is possible to determine singular–plural pairs fairly easily. In this corpus, we find 885 distinct pairs where the plural is marked with *-(i)au* and 41 distinct pairs where the singular is marked with *-en* or *-yn*. (As above, in both cases, we exclude pairs where stem changes are involved.)

When the plural is marked, the ratio of singulars to plurals is 11.08; when the singular is marked, the ratio is 1.26. Singulars greatly outnumber plurals that are marked, but singulars occur far less frequently when they are marked instead. This difference is significant: $t(920.287) = -8.267$, $p < .001$. This is consistent with the hypothesis that forms with more morphemes are more complex.

6 Marked morphological operations

Let's now consider the question of whether more marked morphological operations are under-represented. If they are, this would be consistent with Anderson's typology and Input Optimization.

To test this, let's look at the distribution of plurals in English using the tagged Brown corpus (Kučera & Francis 1967). The Brown corpus is a fairly old written corpus of 928181 words. The advantage of using it here is that it is tagged, so identification of singular and plural nouns is relatively easy, and it is widely used and available.

Focusing on plural nouns, we can separate them into regular plurals marked with *-s* vs. other plural forms, e.g. *men*, *stigmata*, *radii*, *oxen*, etc. When we pair these up with their respective singular forms, we get the following overall counts:

(20)

Type	Pl. tokens	Sg. tokens	Pairs/types
-s	46083	111904	4266
Irregular	2891	2517	79

Overall, there are far more regular than irregular forms, but this is, of course, to be expected by the very definition of "irregular". It is, however, also consistent with Input Optimization. The complexity of a system can be enhanced by limiting the number of forms that exhibit marked properties. It can also be enhanced by limiting the distribution of forms that do have those properties.

Is there a difference in the likelihood of a plural form given its regularity? If irregular forms are more complex, then we would expect their use to be statistically under-represented because of Input Optimization. To test this, we calculate the ratio of singular to plural tokens for each noun pair. This ratio allows us to examine the relative distribution of singular and plural forms, abstracting away from the overall frequency of any specific lexical item. The difference is plotted in Figure 1.

Strikingly, the difference goes in the wrong difference here: irregular plurals are more frequent relative to their singular forms than regular *-s* plurals with respect to their singular forms. This difference is significant: $t(90.318) = 3.151$, $p = 0.002$.

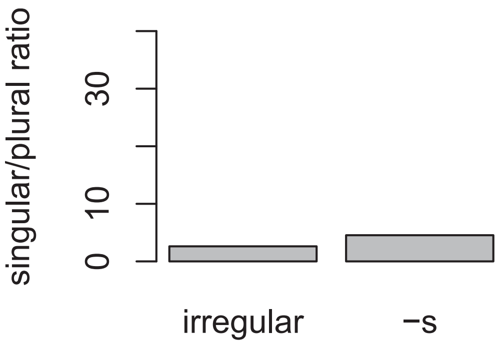

Figure 1: Singular-to-plural ratios for regular and irregular plurals in English

We conclude that the distribution of irregular plurals is ambiguous. In terms of relative frequency of singulars and plurals, the distribution goes in the wrong direction. In terms of overall distribution, however, it goes in the right direction. There are 2891 instances of irregular plurals in Brown and 46083 instances of regular plurals. If we were to assume that both types were equally likely, the difference is certainly significant: $X^2(4344, N = 48974) = 425346.797$, $p < .001$.

7 Zero marking

Let's now turn to zero marking. The claim is that zero marking is more complex and therefore the prediction is that zero marking should be under-represented.

We examine this with respect to plurals in English in the Brown corpus. Zero-marked plurals in English includes examples like: *deer*, *aircraft*, *buffalo*, etc. The difference in ratios between regular plurals in -*s* and zero plurals is shown in Figure 2.

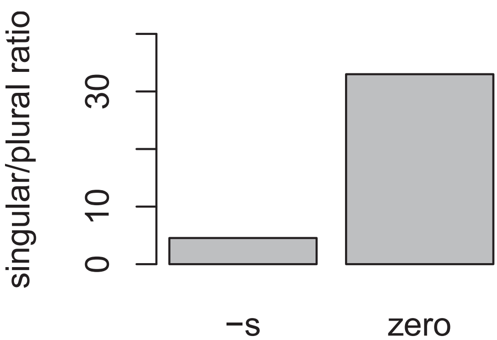

Figure 2: Singular-to-plural ratios for regular and zero plurals in English

Zero-marked plurals are far more frequent—relatively speaking—than regular plurals. Unfortunately, the variance is quite high—there is a lot of variation within each category—and though the mean difference is large, it is not significant: $t(19.002) = -1.416$, $p = 0.173$. As with the irregular plurals, however, the absolute difference is significant. There are 184 instances of zero plurals in Brown and 46083 instances of regular plurals. If we were to assume that both types were equally likely, the difference is certainly significant: $X^2(4285, N = 46267) = 311130.115$, $p < .001$. Again then, though the relative count is not significant, the absolute count goes in the right direction.

8 Summary

Our goal here has been to test the dimensions of morphological complexity proposed in Anderson (2015) with the theory of Input Optimization. As reviewed above, Input Optimization maintains that grammatical complexity, as assessed through constraint violation, is minimized at the input level of the grammar. Specifically, we should see under-representation of more marked morphological structures.

We picked out several dimensions of morphological complexity to examine, some of which have already been treated with respect to Input Optimization. The following list is repeated from Section 3 and annotated to reflect our results.

(21)
1. Words should have fewer morphemes. *This is borne out by the distribution of marked plurals and marked singulars in Welsh.*
2. Haplology should be avoided. (This has already been established by Hammond 2016.)
3. More marked morphological operations (per the hierarchy above) should be avoided. *This was tested with respect to English plurals and is borne out in an absolute comparison, but not in a relative comparison.*
4. Morphophonology should be avoided. (This has already been established by Hammond 2013.)
5. Ablaut, umlaut, truncation, etc. should be avoided. (This is essentially the same as #3 above.)
6. Zero-marking should be avoided. *This was tested with respect to English plurals and is borne out in an absolute comparison, but not in a relative comparison.*

First, all else being equal, we expect forms with more morphemes to be dispreferred to forms with fewer morphemes. We saw that this was borne out in a comparison of singular–plural pairs in Welsh where in some cases the singular has an extra morpheme and in others the plural has an extra morpheme.

Second, we predict that morphological haplology should be under-represented. This was established in previous work with respect to the English genitive plural and adjectives in *-ly*.

Third, more marked morphological operations should be under-represented with respect to less marked morphological constructions. We saw an overall effect here with English irregular noun plurals. We also saw that the relative distribution of plurals with respect to singulars went in the opposite direction.

Fourth, we predict that morphophonology should be avoided. This was established in previous work with respect to morphophonology associated with English past *-ed* and plural *-s*.

The fifth point is the same as the third.

Finally, zero-marking should be under-represented. We saw an overall effect here with English zero-marked noun plurals. We also saw that the relative distribution of plurals with respect to singulars went in the opposite direction.

We conclude that the parameters of complexity developed in Anderson (2015) tested here and in previous work are correct.

We have seen that there is some divergence in the absolute and relative representation of plural marking, but we leave investigation of that for future work.

Acknowledgements

Thanks to Nick Kloehn and Diane Ohala for useful discussion. All errors are my own.

References

Anderson, Stephen R. 1982. Where's morphology? *Linguistic Inquiry* 13(4). 571–612.

Anderson, Stephen R. 1992. *A-morphous morphology*. Cambridge: Cambridge University Press.

Anderson, Stephen R. 1999. A formalist's reading of some functionalist work in syntax. In Michael Darnell, Edith A. Moravcsik, Frederick J. Newmeyer, Michael Noonan & Kathleen Wheatley (eds.), *Functionalism and formalism in linguistics*, vol. 1, 111–135. Amsterdam: John Benjamins.

Anderson, Stephen R. 2015. Dimensions of morphological complexity. In Matthew Baerman, Dunstan Brown & Greville G. Corbett (eds.), *Understanding and measuring morphological complexity*, 11–26. Oxford: Oxford University Press.

Aronoff, Mark & Zheng Xu. 2010. A Realization Optimality-Theoretic approach to affix order. *Morphology* 20. 381–411.

Ellis, N. C., C. O'Dochartaigh, W. Hicks, M. Morgan & N. Laporte. 2001. Cronfa Electroneg o Gymraeg (CEG): A 1 million word lexical database and frequency count for Welsh. http://www.bangor.ac.uk/canolfanbedwyr/ceg.php.en.

Hammond, Michael. 2000. There is no lexicon! *Coyote Papers* 10. ROA #43 (1995), 55–77.

Hammond, Michael. 2013. Input optimization in English. *Journal of the Phonetic Society of Japan* 17. 26–38.

Hammond, Michael. 2014. Phonological complexity and input optimization. *Phonological Studies* 17. 85–94.

Hammond, Michael. 2016. Input optimization: Phonology and morphology. *Phonology* 33. 459–491.

Hayes, Bruce. 1984. The phonology of rhythm in English. *Linguistic Inquiry* 15. 33–74.

Kurisu, Kazutaka. 2001. *The phonology of morpheme realization*. University of California, Santa Cruz PhD thesis.

Kučera, Henry & W. Nelson Francis. 1967. *Computational analysis of present-day American English*. Providence: Brown University Press.

Liberman, Mark & Alan Prince. 1977. On stress and linguistic rhythm. *Linguistic Inquiry* 8(2). 249–336.

Prince, Alan & Paul Smolensky. 1993. Optimality Theory. University of Massachusetts and University of Colorado.

Russell, Kevin. 1993. *A constraint-based approach to phonology*. University of Southern California PhD thesis.

Russell, Kevin. 1995. Morphemes and candidates in Optimality Theory. ROA #44.

Xu, Zheng & Mark Aronoff. 2011. A Realization Optimality Theory approach to blocking and extended morphological exponence. *Journal of Linguistics* 47(3). 673–707.