

Modeling Productivity with the Gradual Learning Algorithm: The Problem of Accidentally Exceptionless Generalizations

Adam Albright
Massachusetts Institute of Technology

Bruce Hayes
University of California,
Los Angeles

1. Introduction¹

Many cases of gradient intuitions reflect conflicting patterns in the data that the child receives during language acquisition. An area in which learners frequently face conflicting data is inflectional morphology, where different words often follow different patterns. Thus, for English past tenses, we have *wing* ~ *winged* (the most common pattern in the language) *wring* ~ *wrung* (a widespread [ɪ] ~ [ʌ] pattern), and *sing* ~ *sang* (a less common [ɪ] ~ [æ] pattern). In cases where all of these patterns could apply, such as the novel verb *spling*, the conflict between them leads English speakers to entertain multiple possibilities, with competing outcomes falling along a gradient scale of intermediate well-formedness (Bybee and Moder 1983; Prasada and Pinker 1993; Albright and Hayes 2003).

In order to get a more precise means of investigating this kind of gradience, we have over the past few years developed an implemented formal model for the acquisition of inflectional paradigms. An earlier version of our model is described in Albright and Hayes (2002), and its application to various empirical problems is laid out in Albright, Andrade, and Hayes (2001), Albright (2002), and Albright and Hayes (2003). Our model abstracts morphological and phonological generalizations from representative learning data and uses them to construct a stochastic grammar that can generate multiple forms for novel stems like *spling*. The model is tested by comparing its “intuitions,” which are usually gradient, against human judgments for the same forms.

In modeling gradient productivity of morphological processes, we have focused on the **reliability** of the generalizations: how much of the input data do they cover, and how many exceptions do they involve? In general, greater productivity is correlated with greater reliability, while generalizations covering few forms or entailing many exceptions are relatively unproductive. For English past tenses, most generalizations have exceptions, so finding the productive patterns requires finding the generalizations with the fewest exceptions.

¹ For helpful comments and advice we would like to thank Paul Boersma, Junko Ito, Armin Mester, Jaye Padgett, Hubert Truckenbrodt, the editors, and our two reviewers, absolving them for any shortcomings.

Intermediate degrees of well-formedness arise when the generalizations covering different patterns suffer from different numbers of exceptions.

The phenomenon of gradient well-formedness shows that speakers do not require rules or constraints to be exceptionless; when the evidence conflicts, they are willing to use less than perfect generalizations. One would expect, however, that when gradience is observed, more reliable generalizations should be favored over less reliable ones. In this article, we show that, surprisingly, this is not always the case. In particular, we find that there may exist generalizations that are exceptionless and well-instantiated, but are nonetheless either completely invalid, or are valued below other, less reliable generalizations.

The existence of exceptionless, but unproductive patterns is a challenge for current approaches to gradient productivity, which generally attempt to extend patterns in proportion to their strength in the lexicon. We offer a solution for one class of these problems, based on the Optimality-theoretic principle of constraint conflict and employing the Gradual Learning Algorithm (Boersma 1997, Boersma and Hayes 2001). In the final section of the paper we return to our earlier work on gradience and discuss the implications of our present findings.

2. Navajo Sibilant Harmony

The problem of exceptionless but unproductive generalizations arose in our efforts to extend our model to learn non-local rule environments. The first example we discuss comes from Sibilant Harmony in Navajo, a process described in Sapir and Hoijer (1967).

Sibilant harmony can be illustrated by examining the allomorphs of the *s*-perfective prefix. This prefix is realized as shown in (1) (examples from Sapir and Hoijer):²

- (1) a. [ši-] if the *first segment* of the stem is a [–anterior] sibilant ([č, č', č^h, š, ž]), for example in [ši-č^hit] ‘he is stooping over’
- b. [ši-] or [si-] if *somewhere later in the stem* is a [–anterior] sibilant, as in [ši-t^hé:ž] ~ [si-t^hé:ž] ‘they two are lying’ (free variation)
- c. [si-] otherwise, as in [si-t^hí] ‘he is lying’³

A fully realistic simulation of the acquisition of Navajo sibilant harmony would require a large corpus of Navajo verb stems, along with their *s*-perfectives. Lacking such a corpus, we performed idealized simulations using an artificial language based on Navajo: we selected whole Navajo words (rather than stems) at random from the electronic version of Young,

² We have rendered all transcriptions (including Sapir and Hoijer’s) in near-IPA, except that we use [č č^h č' š ž] for [tʃ tʃ^h tʃ' ʃ ʒ] in order to depict the class of nonanterior sibilants more saliently.

³ Sapir and Hoijer specifically say (pp. 14-15): “Assimilation nearly always occurs when the two consonants are close together (e.g. ši-čà:ʔ, from si-čà:ʔ ‘a mass lies’; ... but it occurs less often when the two consonants are at a greater distance.”

Morgan, and Midgette's dictionary (1992), and constructed *s*-perfective forms for them by attaching [ši-] or [si-] according to the pattern described in (1).

3. The Learning Model

Our learning system employs some basic assumptions about representations and rule schemata. We assume that words are represented as sequences of phonemes, each consisting of a bundle of features, as in Chomsky and Halle (1968). Rules and constraints employ feature matrices that describe natural classes, as well as variables permitting the expression of non-local environments: ([+F]) designates a single skippable segment of the type [+F], while ([+F])* designates any number of skippable [+F] segments. Thus, the environment in (2):

(2) / ___ ([+seg])* [-anterior]

can be read “where a non-anterior segment follows somewhere later in the word” ([+seg] denotes the entire class of segments).

The model is given a list of pairs, consisting of bases and inflected forms. For our synthetic version of Navajo, such a list would be partially represented by (3):

- (3) a. [pà:ʔ] [si-pà:ʔ]
 b. [č'it̚] [ši-č'it̚]
 c. [č^hò:jin] [ši-č^hò:jin]
 d. [kàn] [si-kàn]
 e. [k'áz] [si-k'áz]
 f. [k^héškã:] [ši-k^héškã:], [si-k^héškã:]
 g. [sí:t̚] [si-sí:t̚]
 h. [t^hãš] [ši-t^hãš], [si-t^hãš]
 i. [t^hó:ʔ] [si-t^hó:ʔ]
 j. [t̚t̚é:ž] [ši-t̚t̚é:ž], [si-t̚t̚é:ž]

Where free variation occurs, the learner is provided with one copy of each variant; thus, for (3f) both [k^héškã:] ~ [ši-k^héškã:] and [k^héškã:] ~ [si-k^héškã:] are provided.

The goal of learning is to determine which environments require [si-], which require [ši-], and which allow both. Learning involves generalizing bottom-up from the lexicon, using a procedure described below. Generalization creates a large number of candidate environments; an evaluation metric is later employed to determine how these environments should be employed in the final grammar.

Learning begins by parsing forms into their component morphemes and grouping them by the morphological change they involve. The data in (3) exhibit two changes, as shown in (4); the box surrounds cases of free variation.

(4)	I. Prefix [si-]	II. Prefix [ši-]
a.	[pà:ʔ] [si-pà:ʔ]	a. [č ^h ò:jìn] [ši-č ^h ò:jìn]
b.	[kàn] [si-kàn]	b. [č'it̚] [ši-č'it̚]
c.	[k ^h éškã:] [si-k ^h éškã:]	c. [k ^h éškã:] [ši-k ^h éškã:]
d.	[t̚é:ž] [si-t̚é:ž]	d. [t̚é:ž] [ši-t̚é:ž]
e.	[t ^h ãš] [si-t ^h ãš]	e. [t ^h ãš] [ši-t ^h ãš]
f.	[k'áz] [si-k'áz]	
g.	[sí:t̚] [si-sí:t̚]	
h.	[t ^h ó:ʔ] [si-t ^h ó:ʔ]	

For each change, the system creates hypotheses about which elements in the environment crucially condition the change. It begins by treating each pair as a “word-specific rule,” separating out the changing part from the invariant part. Thus, the first three [ši-] forms in (4) would be construed as in (5):

- (5) a. $\emptyset \rightarrow \text{ši} / [_ \text{č}^{\text{h}}\text{ò:jìn}]$
 b. $\emptyset \rightarrow \text{ši} / [_ \text{č}'\text{it̚}]$
 c. $\emptyset \rightarrow \text{ši} / [_ \text{k}^{\text{h}}\text{éškã:}]$

Next, the system compares pairs of rules that have the same change (e.g., both attach [ši-]), and extracts what their environments have in common to form a generalized rule. Thus, given the word-specific rules in (6a), the system collapses them together using features, as in (6b).

- (6)a. $\emptyset \rightarrow \text{ši} / [_ \text{t}^{\text{h}}\text{ãš}]$
 $\emptyset \rightarrow \text{ši} / [_ \text{t̚é:ž}]$
- b. $\emptyset \rightarrow \text{ši} / [_ \quad \text{t}^{\text{h}} \quad \quad \quad \text{ã} \quad \quad \quad \text{š} \quad]$
 + $\emptyset \rightarrow \text{ši} / [_ \quad \quad \quad \text{t̚} \quad \quad \quad \text{é:} \quad \quad \quad \text{ž} \quad]$
-
- = $\emptyset \rightarrow \text{ši} / [_ \quad \begin{bmatrix} -\text{sonorant} \\ -\text{continuant} \\ +\text{anterior} \end{bmatrix} \quad \begin{bmatrix} +\text{syllabic} \\ -\text{high} \\ -\text{round} \end{bmatrix} \quad \begin{bmatrix} -\text{sonorant} \\ +\text{continuant} \\ -\text{anterior} \\ +\text{strident} \end{bmatrix} \quad]$

In this particular case, the forms being compared are quite similar, so determining which segment should be compared with which is unproblematic. But for forms of different lengths, such as [č^hò:jìn] and [č'it̚] above, this is a harder question.⁴ We adopt an approach that lines up

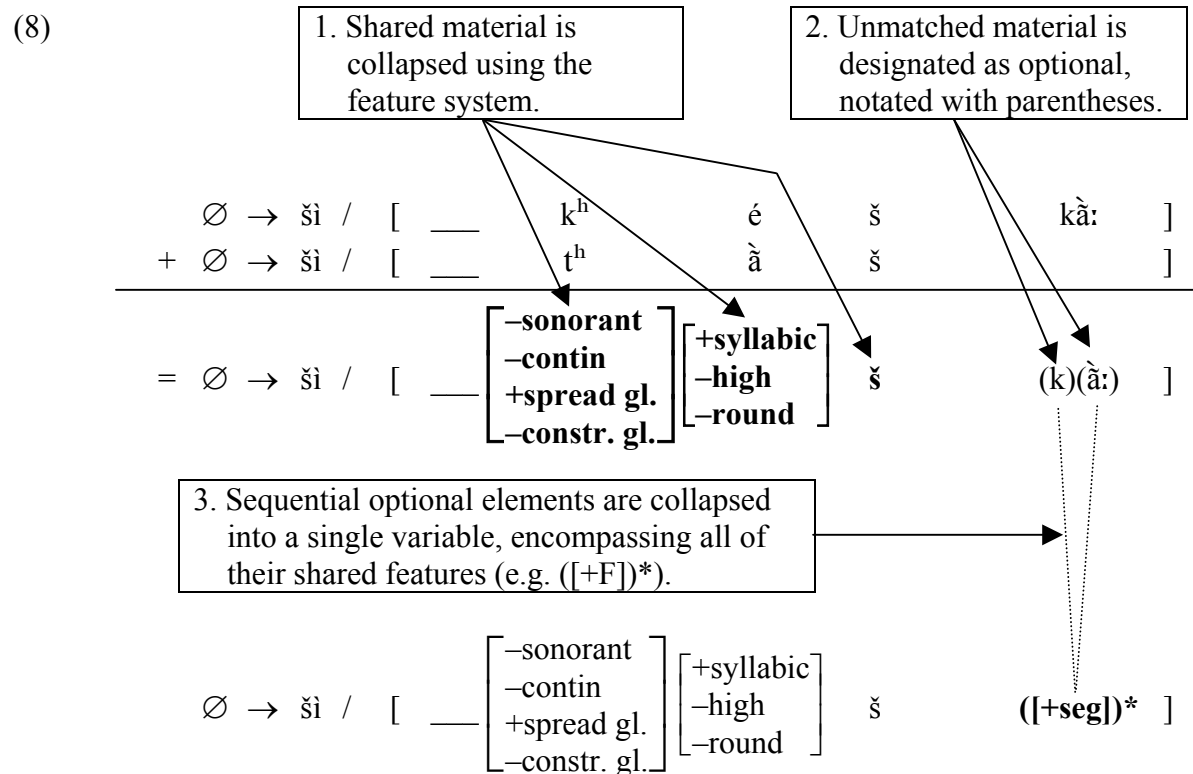
⁴ The issue did not arise in an earlier version of our model (Albright and Hayes 2002), which did not aspire to learn non-local environments, and thus could use simple edge-in alignment.

the segments that are **most similar** to each another. For instance, (7) gives an intuitively reasonable alignment for [č^hò:jìn] and [č'íř]:

$$(7) \quad \begin{array}{cccccc} \text{č}^h & \text{ò:} & \text{j} & \text{i} & \text{n} & \\ | & & & | & | & \\ \text{č}' & & & \text{i} & \text{ř} & \end{array}$$

Good alignments have two properties: they match phonetically similar segments like [č^h] and [č'], and they avoid leaving too many segments unmatched. To evaluate the similarity of segments, we employ the similarity metric from Frisch, Pierrehumbert and Broe (2004). To guarantee an optimal pairing, we use a cost-minimizing string alignment algorithm (described in Kruskal 1999) that efficiently searches all possible alignments for best total similarity.

Seen in detail, the process of collapsing rules is based on three principles, illustrated in (8) with the collapsing of the rules $\emptyset \rightarrow \text{ši} / [_ _ \text{k}^h \text{é} \text{š} \text{k} \grave{\text{a}}:]$ and $\emptyset \rightarrow \text{ši} / [_ _ \text{t}^h \text{ă} \text{š}]$.

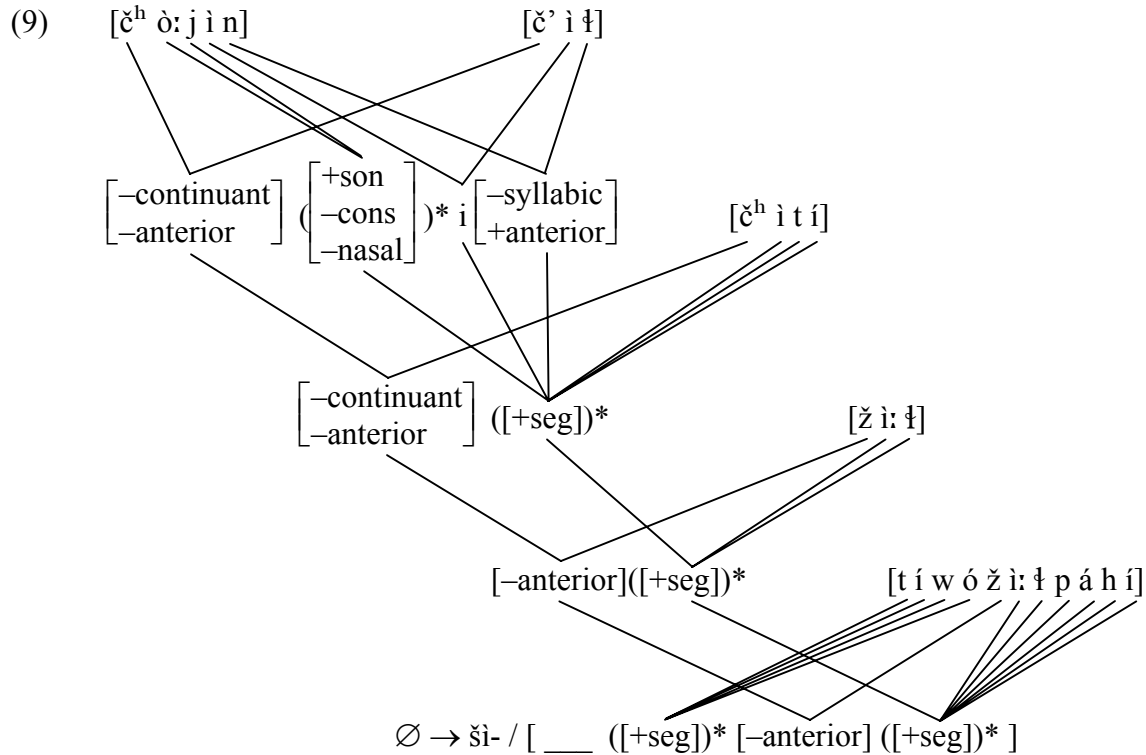


Paired feature matrices are collapsed by constructing a new matrix, containing all of their shared features (see step 1). Next, any material in one rule that is unmatched to the other is designated as optional, represented by parentheses (step 2). Finally, sequences of consecutive optional elements are collapsed together into a single expression of the form $(\mathcal{F})^*$, where \mathcal{F} is the smallest natural class containing all of the collapsed optional elements (step 3).

The process is iterated, generalizing the new rule with the other words in the learning data; the resulting rules are further generalized, and so on. Due to memory limitations, it is necessary

periodically to trim back the hypothesis set, keeping only those rules that perform best.⁵ Generalization terminates when no new “keeper” rules are found.

We find that this procedure, applied to a representative set of words, discovers the environment of non-local sibilant harmony after only a few steps. One path to the correct environment is shown in (9):



The result can be read: “Prefix [ši-] to a stem consisting of any number of segments followed by a nonanterior segment, followed by any number of segments”. (Note that $[-\text{anterior}]$ segments in Navajo are necessarily sibilant.) In more standard notation, one could replace $([\text{+seg}])^*$ with a free variable X, and follow the standard assumption that non-adjacency to the distal word edge need not be specified, as in (10):

$$(10) \quad \emptyset \rightarrow \text{ši-} / \text{___} X [-\text{anterior}]$$

We emphasize that at this stage, the system is only generating hypotheses. The task of using these hypotheses to construct the final grammar is taken up in section 5.

⁵ Specifically: (a) for each word in the training set, we keep the most reliable rule (in the sense of Albright and Hayes 2002) that derives it; (b) for each change, we keep the rule that derives more forms than any other.

4. Testing the Approach: A Simulation

We now show that, given representative learning data, the system just described can discover the rule environments needed for Navajo sibilant harmony. As noted above, our learning simulation involved artificial Navajo *s*-perfectives, created by attaching appropriate prefix allomorphs to whole Navajo words (as opposed to stems). Selecting 200 words at random,⁶ we attached prefixes to the bases as follows, following Sapir and Hoijer's characterization: (a) if the base began with a nonanterior sibilant, we prefixed [šì-] (there were 19 of these in the learning set); (b) if the base contained but did not begin with a nonanterior sibilant, we made two copies, one prefixed with [šì-], the other with [sì-] (37 of each); (c) we prefixed [sì-] to the remaining 144 bases.

Running the algorithm just described, we found that among the 92 environments it learned, three were of particular interest: the environment for obligatory local harmony ((11a)); the environment that licenses distal harmony ((11b); note that this includes local harmony as a special case); and the vacuous "environment" specifying the default allomorph [sì-] ((11c)).

(11) a. Obligatory local harmony

$\emptyset \rightarrow [\text{šì-}] / \text{ ___ } [-\text{anterior}]$

b. Optional distal harmony (= (10))

$\emptyset \rightarrow [\text{šì-}] / \text{ ___ } X [-\text{anterior}]$

c. Default [sì-]

$\emptyset \rightarrow [\text{sì-}] / \text{ ___ } X$

The remaining 89 environments are discussed below.

5. Forming a Grammar

These environments can be incorporated into an effective grammar by treating them not as rules, as just given, but rather as Optimality-theoretic constraints of morphology (Boersma 1998, Russell 1999, Burzio 2002, MacBride 2004). In this approach, rule (11a) is reconstrued as a constraint: "Use [šì-] / ___ [-anterior] to form the *s*-perfective". This constraint is violated by forms that begin with a nonanterior segment, but use something other than [šì-] to form the *s*-perfective. The basic idea is illustrated below, using hypothetical monosyllabic roots:

⁶ From the entire database of 3023 words, we selected 2000 words at random, dividing this set into ten batches of 200 words each. To confirm the generality of our result, we repeated our simulation on each 200-word training sample. Due to space considerations, we report here the results of only one of the ten trials; the remaining nine were essentially the same in that they all discovered the environments in (11). The primary difference between trials was the precise formulation of the other, unwanted constraints (section 6), but in every case, such constraints were correctly ranked below the crucial constraints, as in (13).

(12)	Morphological Base	Candidates that obey USE [ši-] / ___ [-anterior]	Candidates that violate USE [ši-] / ___ [-anterior]
	[šáp]	[ši-šáp]	*[si-šáp], *[mù-šáp], etc.
	[táp]	all	none

It is straightforward to rank these constraints in a way that yields the target pattern, as (13) and (14) show:

(13) USE [ši-] / ___ [-ant] >> { USE [si-] / ___ X, USE [ši-] / ___ X [-ant] } >> all others
ranked in free variation

(14)a.

/sì-čʰit/	USE [ši-] / ___ [-ant]	USE [ši-] / ___ X [-ant]	USE [si-] / ___ X
☞ ši-čʰit			*
* sì-čʰit	*!	*	

b.

/sì-té:ž/	USE [ši-] / ___ [-ant]	USE [ši-] / ___ X [-ant]	USE [si-] / ___ X
☞ ši-té:ž			*
☞ sì-té:ž		*!	

For (14b), the free ranking of USE [ši-] / ___ X [-ant] and USE [si-] / ___ X produces multiple winners generated in free variation (Anttila 1997).

6. Unwanted Constraints

The 89 constraints not discussed so far consist largely of complicated generalizations that happen to hold true of the learning data. One example is shown in (15):

(15) USE sì- / ___ ([-round])* $\left[\begin{array}{l} +\text{anterior} \\ +\text{continuant} \end{array} \right]$ ([-consonantal])*

As it happens, this constraint works for all 37 forms that meet its description in the learning data. However, it makes profoundly incorrect predictions for forms outside the learning data, such as the legal but non-existing stem /čálá/:

(16) USE sì- / ___ ([-round])* $\left[\begin{array}{l} +\text{anterior} \\ +\text{continuant} \end{array} \right]$ ([-consonantal])*

sì-

 \wedge
 č á

 |
 l

 |
 á

If ranked high enough, this constraint will have the detrimental effect of preventing [ši-čálá] from being generated consistently. We will call such inappropriate generalizations as “junk” constraints.

One possible response is to say that the learning method is simply too liberal, allowing too many generalizations to be projected from the learning data. We acknowledge this as a possibility, and we have experimented with various ways to restrict the algorithm to more sensible generalizations. Yet we are attracted to the idea that constraint learning could be simplified—and rely on fewer a priori assumptions—by letting constraints be generated rather freely and excluding the bad ones with an effective evaluation metric. Below, we lay out such a metric, which employs the Gradual Learning Algorithm.⁷

7. The Gradual Learning Algorithm

The Gradual Learning Algorithm (GLA; Boersma 1997, Boersma and Hayes 2001) can rank constraints in a way that derives free variation and matches the frequencies of the learning data. The GLA assumes a stochastic version of Optimality Theory, whereby each pair of constraints {A, B} is assigned not a strict ranking, but rather a probability: “A dominates B with probability P”. Thus, the free ranking given in (13) above would be captured by assigning the constraints USE [si-] / ___ X and USE [ši-] / ___ X [-ant] a 50-50 ranking probability.

Any such theory needs a method to ensure that the pairwise probabilities assigned to the constraints are mutually consistent. In the GLA, this is done by arranging the constraints along a numerical scale, assigning each constraint a **ranking value**. On any particular occasion that the grammar is used, a **selection point** is adopted for each constraint, taken from a Gaussian probability distribution with a standard deviation fixed for all constraints. The constraints are sorted by their selection points, and the winning candidate is determined on the basis of this ranking. Since pairwise ranking probabilities are determined by the ranking values,⁸ they are guaranteed to be mutually consistent.

8. The Need for Generality

Let us now consider the application of the GLA to Navajo. Naively, one might hope that when the constraints are submitted to the GLA, the junk will settle to the bottom. However, what one actually finds is that the junk constraints get ranked high. Although USE [ši-] / ___ [-ant] does indeed get ranked on top, the crucial constraints USE [ši-] / ___ X [-ant] and USE [si-] / ___ X are swamped by higher-ranking junk constraints, and rendered largely ineffective. The result is a grammar that performs quite well on the training data (producing something close to the right

⁷ A reviewer points out that another approach to weeding out unwanted generalizations is to train the model on various subsets of the data, keeping only those generalizations that are found in all training sets (cross-validation). Although this technique could potentially eliminate unwanted generalizations (since each subset contains a potentially different set of such generalizations), it could not absolutely guarantee that they would not be discovered independently in each subset. Given that such constraints make fatal empirical predictions, we seek a technique that reliably demotes them, should they arise.

⁸ A spreadsheet giving the function that maps ranking value differences to pairwise probabilities is posted at <http://www.linguistics.ucla.edu/people/hayes/GLA/>.

output frequencies for every stem), but fails grossly in generating novel forms. The frequencies generated for novel forms are determined by the number of high ranking junk constraints that happen to fit them, and do not respect the distribution in (11).

The problem is a classic one in inductive learning theory. If a learning algorithm excessively tailors its behavior to the training set, it may learn a patchwork of small generalizations that collectively cover the learning data. This does not suffice to cover new forms—which, after all, is the main purpose of having a grammar in the first place!

Why does the GLA fail? The reason is that it demotes constraints only when they prefer losing candidates. But within the learning data, the junk constraints generally prefer only winners—that is precisely why they emerged from the inductive generalization phase of learning. Accidentally true generalizations thus defeat the GLA as it currently stands. What is needed is a way for the GLA to distinguish accidentally true generalizations from linguistically significant generalizations.

9. Initial Rankings Based on Generality

Boersma (1998) suggested that for morphology, initial rankings should be based on generality—the more general the constraint, the higher it is ranked before learning takes place. It turns out that this insight is the key to solving the Navajo problem. What is needed, though, is a way to characterize generality in numerical terms. There are various possible approaches; Chomsky and Halle (1968), for example, propose counting symbols (fewer symbols = greater generality). Here, we adopt an *empirical* criterion: a morphological constraint is maximally general if it encompasses all of the forms that exhibit its structural change. We use the fraction in (17):

$$(17) \quad \frac{\text{number of forms that a constraint applies to}}{\text{total number of forms exhibiting the change that the constraint requires}}$$

In the 200-word Navajo simulation discussed above, some representative generality values are:

(18) Constraint	Relevant forms	Forms with this change	Generality
USE [ši-] / ___ [-anterior]	19	56 [ši-] forms	.339
USE [ši-] / ___ X [-anterior]	56		1
USE [si-] / ___ X	181	181 [si-] forms	1
Constraint (15) (“junk” constraint)	37		.204

The idea, then, is to assign the constraints initial ranking values that reflect their generality, with more general constraints on top. If the scheme works, all the data will be explained by the

most general applicable constraints, and the others will remain so low that they never play a role in selecting output forms.


In order to ensure that differences in initial rankings are large enough to make a difference, the generality values from (17) were rescaled to cover a huge probability range, using the formula in (19):

$$(19) \text{ For each constraint } c, \text{ initial ranking value}_c = 500 \times \frac{\text{Generality}_c - \text{Generality}_{\min}}{\text{Generality}_{\max} - \text{Generality}_{\min}}$$

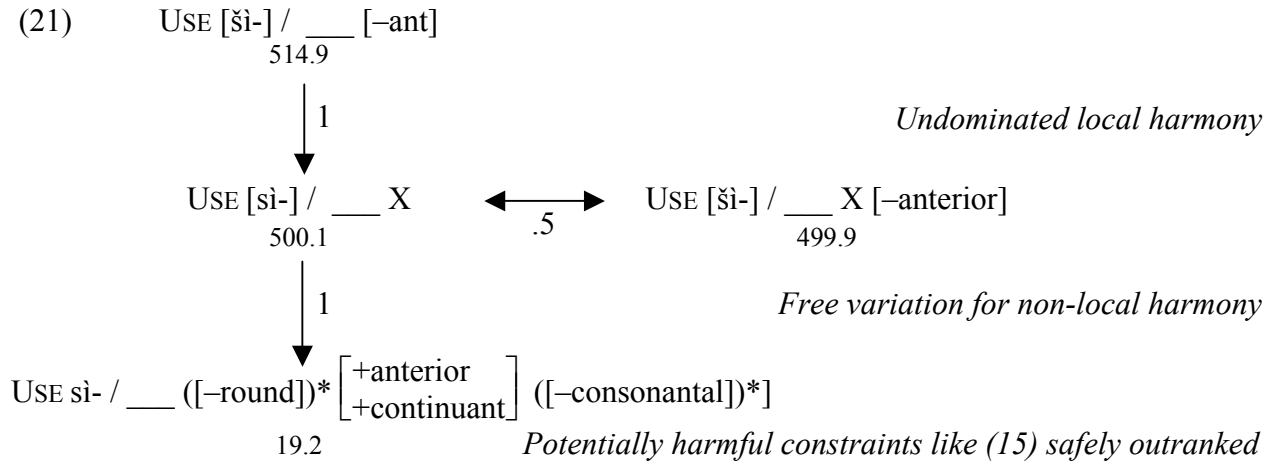
where Generality_{\min} is the generality of the least general constraint, and Generality_{\max} is the generality of the most general constraint.

10. Employing Generality in a Learning Simulation

We implemented this scheme and ran it multiple times on the Navajo pseudodata described above. For one representative run, it caused the relevant constraints (including here just one representative “junk” constraint (15)), to be ranked as follows:



The final grammar is depicted schematically in (21), where the arrows show the probabilities that one constraint will outrank the other. When the difference in ranking value exceeds about 10, the probability that the ranking will hold is essentially 1 (strict ranking).



This approach yields the desired grammar: all of the junk constraints (not just (15)) are ranked safely below the top three.

The procedure works because the GLA is error-driven. Thus, if junk constraints start low, they stay there, since the general constraint that does the same work has a head start and averts any errors that would promote the junk constraints. Good constraints with specific contexts, on the other hand, like “USE [šì-] / ____ [-ant]”, are also nongeneral—but appropriately so. Even if they start low, they are crucial in averting errors like *[sì-šáp], and thus they are soon promoted by the GLA to the top of the grammar.

We find, then, that a preference for more general statements in grammar induction is not merely an aesthetic bias; it is, in fact, a necessary criterion in distinguishing plausible hypotheses from those which are implausible, but coincidentally hold true in the learning sample.

11. Analytic Discussion

While the Navajo simulation offers a degree of realism in the complexity of the constraints learned, hand analysis of simpler cases helps in understanding why the simulation worked, and ensures that the result is a general one.

To this end, we reduce Navajo to three constraints, renamed as follows: (1) USE [sì-], which we will call **DEFAULT**, (2) the special-context USE [šì-] / ____ X [-ant], which we will call **CONTEXTUAL [šì-]**, and (3) the accidentally-exceptionless (15), which we will call **ACCIDENTAL [sì-]**. **ACCIDENTAL [sì-]** is exceptionless because the relevant forms in the training data happen not to contain non-anterior sibilants.

Suppose first that all harmony is optional (50/50 variation). Using the normal GLA, all constraints start out with equal ranking values, set conventionally at 100. The constraints **CONTEXTUAL [šì-]** and **DEFAULT** should be ranked in a tie to match the 50/50 variation. During learning (see Boersma and Hayes 2001, 51-54), these two constraints vacillate slightly as the GLA seeks a frequency match, but end up very close to their original value of 100. **ACCIDENTAL [sì-]** will remain at exactly 100, since the GLA is error driven and none of the three constraints favors an incorrect output for the training data that match **ACCIDENTAL [sì-]** (**DEFAULT** and

ACCIDENTAL [sì-] both prefer [si-], which is correct; and CONTEXTUAL [ši-] never matches these forms). Thus, all three constraints are ranked at or near 100. This grammar is incorrect; when faced with novel forms like (16) that match all three constraints, CONTEXTUAL [ši-] competes against two equally ranked antagonists, deriving [ši-] only a third of the time instead of half.

Initial rankings based on generality (section 9) correct this problem. Given that DEFAULT and CONTEXTUAL [ši-] cover all [si-] and [ši-] forms respectively, they will be assigned initial ranking values of 500. Define the **critical distance** C as the minimum difference between two constraints that is needed to model strict ranking. (Informal trials suggest that a value of about 10.5, which creates a ranking probability of .9999, is sufficient). It is virtually certain that the initial ranking value for ACCIDENTAL [sì-] will be far below $500 - C$, because accidentally true constraints cannot have extremely high generality, other than through an unlikely accident of the learning data. Ranking proceeds as before, with DEFAULT and CONTEXTUAL [ši-] staying around 500 and ACCIDENTAL [sì-] staying where it began. The resulting grammar correctly derives 50/50 variation, because ACCIDENTAL [sì-] is too low to be active.

Now consider what happens when the data involve no free variation; i.e. [ši-] is the outcome wherever CONTEXTUAL [ši-] is applicable. When initial rankings are all equal, [ši-] forms will cause CONTEXTUAL [ši-] to rise and DEFAULT to fall, with their difference ultimately reaching C (CONTEXTUAL [ši-]: $500 + C/2$; DEFAULT: $500 - C/2$). Just as before, ACCIDENTAL [sì-] will remain ranked where it started, at 500. The difference of $C/2$ between CONTEXTUAL [ši-] and ACCIDENTAL [sì-], assuming $C = 10.5$, will be 5.25, which means that when the grammar is applied to novel forms matching both constraints, [si-] outputs will be derived about 3% of the time. This seems unacceptable, given that the target language has no free variation. Again, the incorrect result is avoided under the initial-ranking scheme of section 9, provided that ACCIDENTAL [sì-] is initially ranked at or lower than $500 - C/2$, which is almost certain to be the case.

In summary, schematized analysis suggests that the Navajo result is not peculiar to this case. The effect of accidentally true generalizations is strongest when optionality is involved, but they pose a threat even in its absence. Initial rankings based on generality avoid the problem by keeping such constraints a critical distance lower than the default, so they can never affect the outcome.

12. The Realism of the Simulation

In this section we address two possible objections to our model.

12.1 Phonological Rules vs. Allomorph Distribution

Navajo sibilant harmony is typically analyzed as a phonological process, spreading [–anterior] from right to left within a certain domain. The grammar learned by our model, on the other hand, treats harmony as allomorphy ([sì-] vs. [ši-]), and cannot capture root-internal harmony effects. Thus, it may be objected that the model has missed the essential nature of harmony.

In this connection, we note first that harmony is often observed primarily through affix allomorphy—either because there is no root-internal restriction, or because the effect is weaker

within roots, admitting greater exceptionality. For these cases, allomorphy may be the only appropriate analysis. For arguments that root-internal and affixal harmony often require separate analyses, see Kiparsky (1968).

More generally, however, there still remains the question of how to unify knowledge about allomorphy and root-internal phonotactics. Even when affixes and roots show the same harmony patterns, we believe that understanding the distribution of affix allomorphs could constitute an important first step in learning the more general process, provided there is some way of bootstrapping from constraints on particular morphemes to more general constraints on the distribution of speech sounds. We leave this as a problem for future work.

12.2 Should arbitrary constraints be generated at all?

Another possible objection is that a less powerful generalization procedure would never have posited constraints like (15) in the first place. Indeed, if all constraints come from Universal Grammar (that is, are innate), the need to trim back absurd ones would never arise. Against this objection can be cited work suggesting that environments sometimes really are complex and synchronically arbitrary (Bach and Harms 1972; Hale and Reiss 1998; Hayes 1999; Blevins 2004). For instance, in examining patterns of English past tenses, we found that all verbs ending in voiceless fricatives are regular, and that speakers are tacitly aware of this generalization (Albright and Hayes 2003). Not only are such patterns arbitrary, but they can also be rather complex (see also Bybee and Moder 1982). Regardless of whether such generalizations are learned or innate, it seems likely that any model powerful enough to handle the full range of attested patterns will need a mechanism to sift through large numbers of possibly irrelevant hypotheses.

13. Modeling Gradient Productivity: The Fate of Reliability Metrics

As noted above, one of our long-term goals is to understand how gradient productivity arises when the learner confronts conflicting data. The results above challenge our earlier views, and in this section we lay out ways in which our previous approach might be revised.

Earlier versions of our model evaluated contexts according to their accuracy, or **reliability**, defined as the ratio of the number of forms a rule derives correctly, divided by the total number of forms to which the rule is applicable. We have found in many cases that we could model native speaker intuitions of novel forms by using the reliability of the best rule that derives them (adjusted slightly, in a way to be mentioned below). However, the results of our Navajo simulations show that accuracy alone is not an adequate criterion for evaluation, since assiduous rule discovery can sometimes find accidentally-true (and thus perfectly accurate) generalizations which nonetheless lead to disaster if trusted. The Navajo example illustrates why it is not enough to evaluate the accuracy of each generalization independently; we must also consider whether generalizations cover forms that are better handled by a different generalization.⁹

⁹ A related problem, in which overly broad generalizations appear exaggeratedly accurate because they contain a consistent subset, is discussed in Albright and Hayes (2002).

Another possible failing of the reliability approach is that it is ill-suited to capture special case/“elsewhere” relations (Kiparsky 1982). The environment for [ši-] in Navajo is difficult to express by itself, but easy as the complement set of the [si-] environments. In Optimality theory, “elsewhere” is simply the result of constraint ranking: a context-sensitive constraint outranks the default. Unfortunately for the reliability-based approach, default environments like (11c) often have fairly high reliability (181/237 in this case)—but that doesn’t mean that they should be applied in the special-allomorph context (e.g. of (11a)).

In light of this, it is worth considering why we adopted reliability scores in the first place. Ironically, the reason likewise involved accidentally-true generalizations, but of a different kind.

One of the phenomena that compelled us to use reliability scores was the existence of **small-scale patterns** for irregulars, seen, for example, in English past tenses. As Pinker and Prince (1988) point out, when a system includes irregular forms, they characteristically are not arbitrary exceptions, but fall into patterns, e.g. *cling-clung*, *fling-flung*, *swing-swung*. These patterns have some degree of productivity, as shown by historical change (Pinker 1999) and “wug” (nonce-word) testing (Bybee and Moder 1983, Prasada and Pinker 1993, Albright and Hayes 2003).¹⁰

The problem is that our algorithm can often find environments for these minor changes that are exceptionless. For example, the exceptionless minor change in (22) covers the four verbs *dig*, *cling*, *fling*, and *sling*.¹¹

$$(22) \text{I} \rightarrow \Lambda / \text{X} \left[\begin{array}{l} +\text{cor} \\ +\text{ant} \\ +\text{voice} \end{array} \right] \text{---} \left[\begin{array}{l} +\text{dorsal} \\ +\text{voice} \end{array} \right]]_{[+\text{past}]}$$

The GLA, when comparing an exceptionless constraint against a more general constraint that suffers from exceptions, always ranks the exceptionless constraint categorically above the general one. For cases like Navajo, where the special constraint was (11a) and the general constraint was (11c), the default constraint for [si-], this ranking is entirely correct, capturing the special/default relationship. But when exceptionless (22) is ranked categorically above the constraints specifying the regular ending for English (such as USE [-d]), the prediction is that novel verbs matching the context of (22) should be exclusively irregular (i.e., *blig* → *blug*, not **bligged*). There is evidence that this prediction is wrong, from wug tests on forms that match (22). For instance, the wug test reported in Albright and Hayes (2003) yielded the following judgments (scale: 1 worst, 7 best):

¹⁰ We restrict our discussion to phonological patterns; for discussion of patterns based possibly on semantic, rather than phonological similarities, see Ramscar (2002). In principle, the approach described here could be easily extended to include constraints that refer to other kinds of information; it is an empirical question what properties allomorphy rules may refer to.

¹¹ This is the largest set of $\text{I} \rightarrow \Lambda$ verbs that yields an exceptionless generalization. There are other subsets, such as *cling*, *fling*, and *sling*, that also lead to exceptionless generalizations, and these are also generated by our model. The problem that we discuss below would arise no matter which set is selected, and would not be solved by trying to, for example, exclude *dig* from consideration.

(23)	Present stem	Choice for Past	Rating
a.	<i>blig</i> [blɪg]	<i>blug</i> [blʌg] <i>bligged</i> [blɪgd]	4.17 5.67
b.	<i>splɪŋ</i> [splɪŋ]	<i>splʌŋ</i> [splʌŋ] <i>splɪŋd</i> [splɪŋd]	5.45 4.36

The regular forms are almost as good or better than the forms derived by the exceptionless rule.

We infer that numbers matter: a poorly attested perfect generalization like (22) is not necessarily taken more seriously than a broadly attested imperfect generalization like USE [-d]. For Navajo, strict ranking is appropriate, since the special-environment constraint (11a) that must outrank the default (11c) is robustly attested in the language. In the English case, the special-environment constraint is also exceptionless, but is attested in only four verbs, yet the GLA—in either version—ranks it on top of the grammar, just as in Navajo.

It can now be seen why in our earlier work we avoided constraint interaction and employed reliability scores instead. With reliability scores, it is simple to impose a penalty on forms derived by rules supported by few data—following Mikheev (1997), we used a statistical lower confidence limit on reliability. Thus, for a wug form like *blig*, two rules of comparable value compete: the regular rule (has exceptions, but vast in scope) vs. (22) (no exceptions, but tiny in scope). Ambivalence between the two is a natural consequence.

If reliability statistics are not the right answer to this problem, what is? It seems that the basic idea that rules based on fewer forms should be downgraded is sound. But the downgrade needn't be carried out based on reliability scores—it might also be made part of the constraint ranking process. In particular, we propose that the basic principles of the GLA be supplemented with biases that exert a downward force on morphological constraints that are supported by few data, using statistical smoothing or discounting.

As of this writing we do not have a complete solution, but we have experimented with a form of **absolute discounting** (Ney, Essen and Kneser 1994), implemented as follows: for each constraint C, we add to the learning data an artificial datum that violates C and obeys every other constraint with which C is in conflict. Under this scheme, if C (say, (22) above) is supported by just four forms, then an artificially-added candidate would have a major effect in downgrading its ranking. But if C is supported by thousands of forms (for example, the constraint for a regular mapping), then the artificially added candidate would be negligible in its effect.

We found that when we implemented this approach, it yielded reasonable results for the English scenario just outlined: in a limited simulation consisting of the regulars in Albright and Hayes (2003) plus just the four irregulars covered by (22), regular *splɪŋd* was a viable competitor with *splʌŋ*, and the relationships among the competing regular allomorphs remained essentially unchanged.

There are many ways that small-scale generalizations could be downgraded during learning. We emphasize that the development of a well-motivated algorithm for this problem involves not

just issues of computation, but an empirical question about productivity: when real language learners confront the data, what are the relative weights that they place on accuracy vs. size of generalization? Both experimental and modeling work will be needed to answer these questions.¹²

14. Conclusion

The comparison of English and Navajo illustrates an important problem in the study of gradient well-formedness in phonology. On the one hand, there are cases like English past tenses, in which the learner is confronted with many competing patterns and must trust some generalizations despite some exceptions. In such cases, gradient well-formedness is rampant, and the model must retain generalizations with varying degrees of reliability. On the other hand, there are cases like Navajo sibilant harmony, in which competition is confined to particular contexts, and the learner has many exceptionless generalizations to choose from. In these cases, the challenge is for the model to choose the “correct” exceptionless patterns, and refrain from selecting an analysis that predicts greater variation than is found in the target language.

We seek to develop a model that can handle all configurations of gradient and categoricalness, and we believe the key lies in the trade-off between reliability and generality. We have shown here how our previous approach to the problem was insufficient, and proposed a new approach using the GLA, modified to favor more general constraints. The precise details of how generality is calculated, and how severe the bias must be, are left as a matter for future research.

¹² An unresolved question that we cannot address here is whether a bias for generality can be applied to all types of phonological constraints, or just those that govern allomorph distribution. It is worth noting that for certain other types of constraints, such as faithfulness constraints, it has been argued that specific constraints must have higher initial rankings than more general ones (Smith 2000). At present, we restrict our claim to morphological constraints of the form “USE X”.

References

- Albright, Adam (2002) 'Islands of reliability for regular morphology: Evidence from Italian', *Language* 78, 684-709.
- Albright, Adam, Argelia Andrade and Bruce Hayes (2001) 'Segmental environments of Spanish diphthongization', *UCLA Working Papers in Linguistics* 7, 117-151.
- Albright, Adam and Bruce Hayes (2002) 'Modeling English past tense intuitions with minimal generalization', in Mike Maxwell, ed., *Proceedings of the 2002 Workshop on Morphological Learning*. Philadelphia: Association for Computational Linguistics.
- ___ (2003) 'Rules vs. analogy in English past tenses: a computational/experimental study', *Cognition* 90: 119-161.
- Anttila, Arto (1997) 'Deriving variation from grammar', in Frans Hinskens, Roeland van Hout and Leo Wetzels (eds.), *Variation, change and phonological theory*. Amsterdam: John Benjamins, pp. 35-68.
- Bach, Emmon and Robert Harms (1972) 'How do languages get crazy rules?', in Stockwell, Robert and Ronald Macauley (eds.), *Linguistic Change and Generative Theory*. Bloomington, Indiana: Indiana University Press, pp. 1-21.
- Blevins, Juliette (2004) *Evolutionary Phonology*. Cambridge: Cambridge University Press.
- Boersma, Paul (1997) 'How we learn variation, optionality, and probability', *Proceedings of the Institute of Phonetic Sciences of the University of Amsterdam* 21: 43-58.
- Boersma, Paul (1998) 'Typology and acquisition in functional and arbitrary phonology', Ms., University of Amsterdam. http://www.fon.hum.uva.nl/paul/papers/typ_acq.pdf.
- Boersma, Paul and Bruce Hayes (2001) 'Empirical tests of the Gradual Learning Algorithm', *Linguistic Inquiry* 32: 45-86.
- Burzio, Luigi (2002) 'Surface-to-surface morphology: when your representations turn into constraints', in Paul Boucher (ed.), *Many Morphologies*. Somerville, MA: Cascadilla Press, pp. 142-177
- Bybee, Joan and Carol L. Moder (1983) 'Morphological classes as natural categories', *Language* 59: 251-270.
- Chomsky, Noam and Morris Halle (1968) *The Sound Pattern of English*. New York: Harper and Row.
- Frisch, Stefan, Janet Pierrehumbert and Michael Broe (2004) 'Similarity avoidance and the OCP', *Natural Language and Linguistic Theory* 22: 179-228.
- Hale, Mark and Charles Reiss (1998) 'Formal and empirical arguments concerning phonological acquisition', *Linguistic Inquiry* 29: 656-683.
- Hayes, Bruce (1999) 'Phonological restructuring in Yidj and its theoretical consequences', in Ben Hermans and Marc van Oostendorp (eds.), *The Derivational Residue in Phonological Optimality Theory*. Amsterdam: John Benjamins, pp. 175-205.
- Kiparsky, Paul (1968) 'How abstract is phonology?' Bloomington: Indiana University Linguistics Club. Reprinted 1982 in *Explanation in Phonology*. Dordrecht: Foris, pp. 119-163.
- Kiparsky, Paul (1982) 'Lexical morphology and phonology', in The Linguistics Society of Korea (ed.), *Linguistics in the Morning Calm: Selected Papers from SICOL-1981*. Seoul: Hanshin Publishing Co., pp. 3-91.

- Kruskal, Joseph (1999) 'An overview of sequence comparison', in David Sankoff and Joseph Kruskal (eds.), *Time Warps, String Edits, and Macromolecules: The Theory and Practice of Sequence Comparison*, 2nd ed. Reading, MA: Addison-Wesley, pp. 1-44.
- MacBride, Alex (2004) *A Constraint-Based Approach to Morphology*. Ph.D. dissertation, UCLA, <http://www.linguistics.ucla.edu/faciliti/diss.htm>.
- Mikheev, Andrei (1997) 'Automatic rule induction for unknown-word guessing', *Computational Linguistics* 23:405–423.
- Ney, Hermann, Ute Essen and Reinhard Kneser (1994) 'On structuring probabilistic dependencies in stochastic language modeling', *Computer Speech and Language* 8: 1–28.
- Pinker, Steven (1999) *Words and Rules: The Ingredients of Language*. New York: Basic Books.
- Pinker, Steven and Alan Prince (1988) 'On language and connectionism: Analysis of a parallel distributed processing model of language acquisition', *Cognition* 28: 73–193.
- Prasada, Sandeep and Pinker, Steven (1993) 'Generalization of regular and irregular morphological patterns', *Language and Cognitive Processes* 8: 1–56.
- Ramscar, Michael (2002) 'The role of meaning in inflection: Why the past tense does not require a rule', *Cognitive Psychology* 45, 45-94.
- Russell, Kevin (1999) *MOT: Sketch of an Optimality Theoretic Approach to Morphology*. Ms., <http://www.umanitoba.ca/linguistics/russell/>.
- Sapir, Edward and Harry Hoijer (1967) *The Phonology and Morphology of the Navajo Language*. Berkeley: University of California Press.
- Smith, Jennifer (2000) 'Positional faithfulness and learnability in Optimality Theory', in Rebecca Daly and Anastasia Riehl (eds.), *Proceedings of ESCOL 99*. Ithaca: CLC Publications, 203–214.
- Young, Robert W., William Morgan Sr. and Sally Midgette (1992) *Analytical Lexicon of Navajo*. Albuquerque: University of New Mexico Press.

Author contact information:

Albright: albright@mit.edu
Hayes: bhayes@humnet.ucla.edu