Gradient Well-Formedness in Optimality Theory[*]

Bruce P. Hayes

UCLA

March 1998

Abstract

A minor modification in the framework of Optimality Theory (Prince and Smolensky 1993) is suggested which enables it to model phenomena where consultant intuitions are gradient, falling somewhere between complete well-formedness and complete ill-formedness. The proposal consists of assigning to certain constraints *bands of values* along a reified continuum of constraint strictness. When a particular form can be generated only by assigning a constraint a strictness value within a designated "fringe" of the strictness band, the grammar generates the form marked with an intermediate degree of well-formedness.

The proposal is tested against data involving light and dark /l/ in American English, using a set of gradient intuitions obtained from ten native speaker consultants. A

[*] Thanks to reviewer Carlos Gussenhoven for a number of thoughtful comments which have been incorporated into this article.

rationale from language learning is then posited for why well-formedness intuitions are so frequently gradient.

Gradient Well-Formedness in Optimality Theory

Bruce P. Hayes

UCLA

March 1998

1. The Gradience Problem

Virtually every generative linguist has had the following experience: a given linguistic entity (sentence, novel word, pronunciation) is presented to a native speaker and judged to be neither fully well-formed nor fully unacceptable. In such instances, consultants often say things like "I *guess* I could say that," "It's all right but not perfect," "It's pretty bad but not completely out," and the like.

Such judgments often get reified: the data are sorted into grammatical and ungrammatical categories, and an analysis is developed whose rules or principles generate all and only the grammatical outcomes. This procedure is controversial, as Schütze's (1996) comprehensive review points out. Critics of generative grammar might take the existence of gradient well-formedness judgments as an indication that the entire enterprise is misconceived: that discrete "categories," "rules," and "constraints" are just illusions suffered by the linguist. In this eliminativist view, gradient well-formedness judgments constitute evidence that generative linguistics must be replaced by something very different, something much "fuzzier."

Other scholars, also reviewed by Schütze, have maintained that well-formedness really is categorical (surface appearances to the contrary), and that gradience is merely the result of performance factors that obscure the judgment process. However, while such factors certainly do exist, much of the patterning of gradient judgments is based on authentic structural aspects of the linguistic material being judged. This makes it unlikely that it could be described as mere performance (see Schütze 1996, 63-64; and 3.1.4 below).

A third possibility, the one advocated here, is to claim that gradient well-formedness judgments are on the whole authentic: abstracting away appropriately from performance factors, carefully elicited gradient judgments really do reflect the internalized knowledge of the native speaker. What has been lacking, in this view, is the right theoretical tools to model grammars that can generate outputs with varying degrees of well-formedness.

My specific suggestion is that within Optimality Theory (Prince and Smolensky 1993), it is possible to devise such grammars. The modification in the theory that is needed is strikingly minor, and is quite independent of the choice of formal representations and constraints used in the grammar.

The actual proposal has been briefly presented in Hayes and MacEachern (forthcoming). The material that article deals with is perhaps somewhat peripheral, being drawn from the theory of metrics, particularly the part of the theory concerned with rhythmic form. The goal of the present article is to show that the approach we adopted can be extended to a strictly linguistic example. I will also attempt to increase the plausibility of the proposal by offering a rationale, based on language learning, for why gradient well-formedness exists and why it occurs in particular areas.

2. The Proposal

I will go through the proposal in a series of steps, each one a closer approximation to the final version.

2.1 Free Ranking

Consider first the issue of free variation, where a single input yields multiple outputs. A commonly adopted approach to free variation in Optimality Theory is to suppose that certain pairs of constraints may be *ranked freely*. Each variant outcome is obtained by fixing the free rankings of a grammar in a particular way.[1] In this view, a grammar is not a monolithic ranking, generating a single set of outputs, but a set of rankings, of which some are obligatory and some are free. If one examines the complete set of possible total rankingsof the constraints of such a grammar, it will be found that the set of "subgrammars" thus defined will collectively generate all and only the forms permitted in the language, including all the free variants.[2]

2.2 Strictness Bands

It will be useful in what follows to consider rankings not as simple arrangements of constraint pairs, but rather as the result of the constraints each taking on a range of values on an abstract continuum. This is a concept explored independently by Boersma (this volume). We can speak of each constraint possessing a *strictness band*, and depict the bands graphically as follows:

(1) more strict ⟷ less strict

CONSTRAINT A:

CONSTRAINT B:

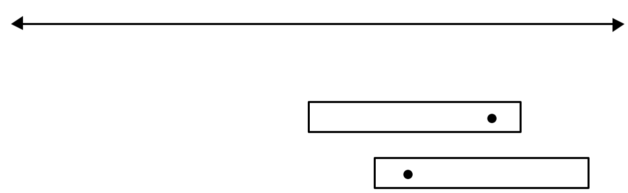

Within each band, I have given a *selection point*, which is defined as the particular value of strictness taken on by a constraint on a given speaking occasion. It can be seen that as long as the strictness

[1] A nice example of the empirical effectiveness of the free-ranking approach may be found in Davidson and Noyer 1997).

[2] One should add that, in any real-life grammar, it will turn out that there are quite a few pairwise rankings that simply don't matter. They may without harm be considered as free (vacuous free ranking); or they may be assigned a ranking arbitrarily; the same empirical outcomes will be observed.

bands of two constraints overlap, then both rankings of the two constraints will be available for the generation of outputs. In (1), the selection points are such that on the particular speaking occasion involved, outputs will be generated that respect a ranking of CONSTRAINT A over CONSTRAINT B. Likewise, on other occasions the selection points could require a ranking of B over A:

(2)

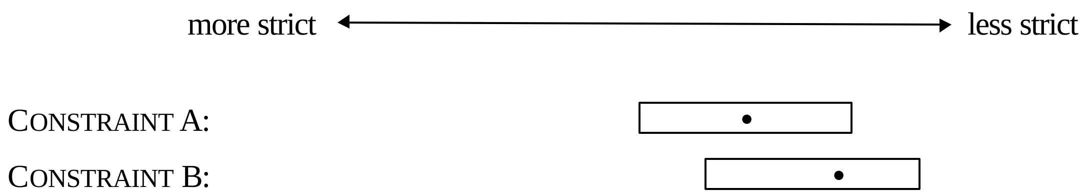

We may also suppose that in speech perception, the listener may explore the strictness bands until a set of selection points is found that appropriately matches the phonetic form to a suitable underlying phonological representation (for discussion see Boersma 1997).

Where bands fail to overlap, the width of the bands is of course vacuous, since the ranking that results will be the same no matter where the selection points fall.

It is easy to imagine extensions of the strictness band notion beyond just accounting for optionality. Suppose, for instance, that we adopt the assumption that on any given speaking occasion, the selection of a strictness value within a band is made at random. In such cases, the *relative frequency* of output forms will vary accordingly. The issue of modeling text frequencies is discussed further in Boersma (this volume) and below in section 5.2.

2.3 Fringes

The actual problem at hand, gradient well-formedness, can be treated by further amplifying the strictness-band idea. Let us suppose that, at least in the crucial cases, the range of a constraint is not firmly delimited. Formally, we can model this idea by positing *fringes*: at the edge of a constraint's strictness band, we add special blocks labeled with traditional well-formedness diacritics such as "?" and "??". Selection points may occur within a fringe, but only at the cost of the degree of ill-formedness indicated. For instance, the sample diagram in (3):

(3)

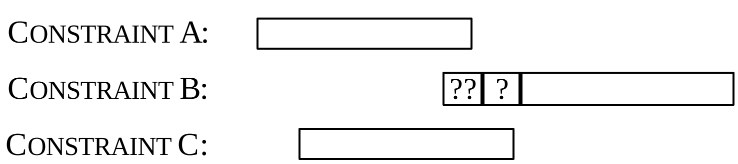

can be interpreted as follows: (a) Ordinarily, Constraint B is *outranked* by constraints A and C. (b) However, it is *somewhat possible* for B to outrank C. This will occur if the selection point for B occurs quite close to the left edge of its "?" fringe, and the selection point for C occurs quite close to the right edge of its strictness band as a whole. Forms that can be generated only with this ranking are intuited to be mildly ill-formed (“?”). (c) It is only *marginally possible* for B to outrank A. This will occur if the selection point for B occurs close to the left edge of its “??” fringe, and that for A close to the right edge of its strictness band. Forms that can be generated only with this ranking are intuited to be considerably ill-formed, though they are not completely excluded (“??”).[3]

As for where the fringes come from: I conjecture that they arise as part of the acquisition process, in cases where the input data do not suffice to establish firmly what the upper or lower bounds of a constraint’s strictness band are. This is discussed further in section 4.

How do we evaluate the well-formedness of a given linguistic form under the proposed model? I assume that in judging a given form, a consultant will normally assign it *the highest rating possible* under the grammar. Thus, suppose a given form emerges from the grammar in two ways: (a) with “?” attached, using a choice of strictness values that employs a “?” fringe of one of the strictness bands; (b) perfectly, using a choice of strictness values that is drawn solely from the central (non-fringe) portions of the bands. This is illustrated below:

(4) a. CONSTRAINT A:
CONSTRAINT B:

b. CONSTRAINT A:
CONSTRAINT B:

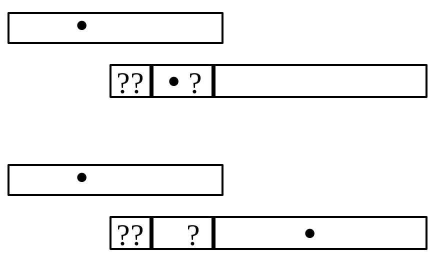

I assume that consultants normally avoid “perverse” parsings of the input like (4a), which assign it some degree of ill-formedness, when there exist alternative selection points ((4b)) that generate the form as perfect.

The proposal here is strongly reminiscent of Schütze’s (1996, 172, 189) view of the grammaticality judgment process. Schütze suggests that when consultants judge marginal forms, “constraints [can] be selectively relaxed when an initial parse fail[s]. Once a parse [is] eventually found this way, ... the nature and degree of constraint relaxation [is] reflected in their ungrammaticality ratings.” The difference here,

---

[3] Carlos Gussenhoven sensibly notes that it is unlikely that strictness ranges include discrete boundaries between “?”, “??”, and “*”. I would agree with this view completely, and posit that fringes are indeed themselves gradient in the ill-formedness they impose. Ideally, what one wants is not a set of discrete blocs, but rather a continuous, hump-shaped function with the constraint strictness continuum as its x axis and “goodness” values as its y axis.

Despite this, throughout the discussion below I will continue to assume an artificial categorization into discrete “?” and “??” fringes. It will be seen that this idealization greatly facilitates analysis. Moreover, I believe the idealization is harmless. In scientific modeling, step functions are frequently used to approximate continuous functions, provided the “grain” imposed by the steps is innocuously fine. The “?” and “??” blocs used here can be regarded as a suitable approximation to a continuous “goodness” function, because the data that would be needed to distinguish the predictions of the continuous function from the stepwise approximation would be far more refined than what is available here.

following the general approach of OT, is that constraints *per se* are not relaxed, but rather their mutual rankings.

This is the proposal. In the remainder of this article, I will try to provide some empirical support for it with a case study, worked out under the assumptions just made. I will then address some general issues that the proposal raises.

3. Case Study: English Light and Dark /l/

From time to time I have pondered the distribution of the light and dark allophones of /l/ in my idiolect of American English. I have usually found the judgments difficult to make and rather gradient in character. Thus the light versus dark /l/ problem seemed a reasonable area to take on as a case study.

3.1 The /l/ Data

Below, I give what I take to be crucial background facts, followed by the results of a survey of well-formedness judgments.

3.1.1 The Phonetics of /l/

The phonetic quality of /l/ in English dialects varies to a startling degree, as I have discovered by eliciting forms from speakers from different countries and regions. However, most dialects some kind of distinction between "light" and "dark" /l/. In my own speech, which I find to be similar to that of many other Americans in the relevant respects, both light and dark /l/ are rather velarized (back tongue body position). This has been noticed previously for American dialects by Wells (1982, 490). The light allophone [l] is less backed than the dark, and has obligatory tongue blade contact in the denti-alveolar region. Dark [ɫ] is backer; and in the more casual speech registers, if it is not prevocalic, it can lose its tongue blade contact entirely, becoming a kind of high back vocoid with lateral tongue body compression.[4]

A phonetic matter that will be important here is the highly noticeable allophony found on vowels that precede dark [ɫ]. In my own speech, such vowels are backed, receive a schwa off-glide if front or high, and are otherwise monophthongized. /ɑ/ is slightly rounded; and the "true diphthongs" /a͡ɪ, a͡ʊ, ɔ͡ɪ/ get schwa off-glides. The chart in (5) gives a survey; related observations may be found in Wells (1982, 487).

[4] In various dialects such as Cockney (Wells 1982) or Adelaide Australian (Borowsky and Horvath, forthcoming), the dark /l/ can lose its laterality entirely, becoming a simple back vowel. This phenomenon does not occur, I believe, in the American varieties studied here.

(5)

Default Quality	**Example**		**Pre-[ɫ] Allophone**	**Example**	
[iː]	*tree*	[tɹiː]	[i͡ə]	*feel*	[fi͡əɫ]
[ɪ]	*fit*	[fɪt]	[ɪ͡ə]	*fill*	[fɪ͡əɫ][5]
[e͡ɪ]	*pay*	[pe͡ɪ]	[e͡ə]	*pail*	[pe͡əɫ]
[ɛ]	*set*	[sɛt]	[ɛ͡ə]	*sell*	[sɛ͡əɫ][5]
[æ]	*pat*	[pæt]	[a͡ə]	*pal*	[pa͡əɫ]
[ʉ͡u]	*do*	[dʉ͡u]	[u͡ə]	*fool*	[fu͡əɫ]
[o͡ʊ]	*foe*	[fo͡ʊ]	[o̱ː]	*foal*	[fo̱ːɫ]
[ʊ]	*put*	[pʊt]	[ʊ͡ə]	*pull*	[pʊ͡əɫ][5]
[ʌ]	*but*	[bʌt]	[ʌ̱]	*dull*	[dʌ̱ɫ]
[ɔ͡ə]	*saw*	[sɔ͡ə]	[ɔ̱]	*Saul*	[sɔ̱ɫ]
[ɑ]	*Pa*	[pɑ]	[ɒ̱]	*all*	[ɒ̱ɫ]
[a͡ɪ]	*tie*	[ta͡ɪ]	[aɪ͡ə]	*tile*	[taɪ͡əɫ]
[a͡ʊ]	*cow*	[ka͡ʊ]	[aʊ͡ə]	*cowl*	[kaʊ͡əɫ]
[ɔ͡ɪ]	*boy*	[bɔ͡ɪ]	[ɔɪ͡ə]	*boil*	[bɔɪ͡əɫ]

The most careful phonetic study yet done of backness in English /l/ was carried out by Sproat and Fujimura (1993), who gathered X-ray microbeam data. An important claim that Sproat and Fujimura make is that there is no categorial distinction between light and dark /l/ in English, but only a phonetic continuum. This claim strikes me as controversial. In particular, the data Sproat and Fujimura gathered did not include the most crucial cases for demonstrating categories, namely those cited below under (14). Inspection of the data in their article suggests, to me at least, that there are two analytical possibilities that remain tenable: (a) an outright phonetic continuum, as Sproat and Fujimura propose, or (b) two phonetic categories that are partially obscured by free variation and near-neutralizing lenition. It is the latter interpretation that will be adopted here, as it appears to be necessary to adopt this assumption in order to give a coherent account of the Paradigm Uniformity effects noted below.

3.1.2 A Survey of Well-Formedness Judgments

After exploring the pattern of /l/ allophony using just my own intuitions, I attempted to obtain higher quality data by conducting an organized survey of ten native speakers of American English.[6] All of my consultants were volunteers and were known to me; half were linguists and half were not. As

[5] A caution regarding [ɪ͡ə], [ɛ͡ə], and [ʊ͡ə]: like their corresponding default allophones, these nuclei are short, and the diphthongization is not nearly as perceptible for them as it is for other vowels.

[6] I discovered, somewhat late in the process, that there is a recent literature that provides valuable guidance to linguists conducting such surveys: Schütze (1996), Bard, Robertson, and Sorace (1996), and Cowart (1997). Much of what I have done here could have been done better by following the prescriptions of these works.

All of the survey data may be downloaded from the author's Web site, listed at the end of this article.

subsequent statistical testing showed no essential differences between the two groups,[7] the results are pooled below. Each consultant was presented with 17 words, sometimes placed in a particular sentence context to make their meaning clear. Each word contained an /l/, and was pronounced by the author in two ways, with light [l] and dark [ɫ]. The consultants were asked to rate *both* pronunciations on an integer scale ranging from 1 to 7, with 1 designating "sounds just right, perfectly normal in my dialect of English" and 7 designating "sounds awful, I would never say it that way." In addition, the consultants were invited, if it seemed appropriate, to check boxes labeled as follows: "*casual*: to the extent it's acceptable, it's acceptable in casual speech only"; and "*formal*: to the extent that it's acceptable, it's acceptable in formal, careful speech only."

On the whole, the consultants found the task somewhat difficult, and indeed the variance in their responses was rather high. The view taken here is that *in any individual instance*, the judgments were indeed subject to apparently random influences. However, when averaged over all the consultants, the results formed a quite coherent pattern.[8] I will report the results of the survey interspersed below amid a general, structurally-oriented description of the facts.

3.1.3 Data Pattern

In *pretonic* position (that is, immediately before a stressed vowel), it seems utterly obligatory to produce a light [l]. Representative examples, as I would pronounce them, are the following:

(6) **Pretonic: Obligatory Light [l]**

light [ˈ**l**a͡ɪt], *Lee* [ˈ**l**iː], *Lou* [ˈ**l**uː], *aloud* [əˈ**l**a͡ʊd], *balloon* [bəˈ**l**uːn], *apply* [əˈp**l**a͡ɪ]

This is likewise true for word-initial /l/ when the following vowel is stressless:

(7) **Word-Initial Pre-Atonic: Obligatory Light [l]**

Lamarck [**l**əˈmɑɹk], *Louanne* [**l**uˈæn]

If any of the forms of (6) or (7) is pronounced with dark [ɫ] instead, the result sounds laughable, like a phonetic exercise.

This is my own judgment, but it is clearly supported by the results of the survey, which included the words *light* and *Louanne*, ranked on the 1-7 scale. The crucial numbers given are the mean of the

[7] Here are the result of a factorial ANOVA test. For well-formedness judgments, there was no significant main effect of the linguist/non-linguist difference (Fisher's PLSD post-hoc test: p = .501), nor a significant interaction between type of word tested and linguist status (p = .891). For judgments of casualness, there was a small main effect (linguists tended to judge the forms overall as more casual, p = .019), but no significant interaction between word type and linguist status (p = .528).

[8] Cowart (1997, Chap. 2) provides an interesting review and defense of the use of such data patterns in judgment experiments.

judgment of all ten consultants. The column labeled "σ" gives the standard deviation, the accepted measure for describing how much the consultants differed from one another.

(8)

Word	with [l]:		with [ɫ]:	
	mean	σ	mean	σ
light	1.30	0.48	6.10	1.10
Louanne	1.10	0.32	5.55	1.74
average for both words	1.20	0.41	5.83	1.44

There is also an environment where dark [ɫ] is just as obligatory: preconsonantal and prepausal position.

(9) **Preconsonantal: Obligatory Dark [ɫ]**

fault [ˈfɒ̠ɫt], *help* [ˈhɛ͡əɫp], *shelter* [ˈʃɛ͡əɫtɚ]

(10) **Prepausal: Obligatory Dark [ɫ]**

feel [ˈfi͡əɫ], *whole* [ˈho̠ːɫ]

Here, too, the substitution of the wrong allophone (light [l]) produces, I think, a comic effect; in this case, a rather cruel one of mocking a foreign speaker whose native language has only light [l]. The data from the 10 consultants is summarized below:

(11)

Word	with [l]:		with [ɫ]:	
	mean	σ	mean	σ
bell	6.60	0.97	1.20	0.42
help	6.60	0.97	1.05	0.16
average	6.60	0.94	1.12	0.32

The next environment to consider is intervocalic pre-atonic position. This environment, sometimes called "ambisyllabic," yields special allophones for quite a few English consonants. For /l/ in the target dialect, it evokes free variation between light and dark:

(12) **Intervocalic Pre-Atonic: Free Variation**

Greeley	[ˈgriːli, ˈgri͡əɫi]
Bailey	[ˈbe͡ɪli, ˈbe͡əɫi]
mellow	[ˈmɛloʊ, ˈmɛ͡əɫoʊ]
Hayley	[ˈhe͡ɪli, ˈhe͡əɫi]
Mailer	[ˈme͡ɪlɚ, ˈme͡əɫɚ]

The consultant survey examined four such words, and in all, found fair acceptability for both variants, the light one being slightly preferred:

(13)

Word	**with [l]:**		**with [ɫ]:**	
	mean	σ	mean	σ
(Norman) Mailer	2.00	1.33	2.00	1.33
Hayley (Mills)	1.55	0.96	3.05	1.83
(Horace) Greeley	1.80	1.32	2.70	1.77
(Mayor) Daley	2.25	1.48	2.80	1.62
average	1.90	1.26	2.64	1.63

The realization of intervocalic pre-atonic /l/ is also influenced by morphology. Suppose first that the /l/ is the first segment of a suffix. Here, the preference is rather strongly shifted to a light [l], and I will mark the dark [ɫ] forms with a question mark:

(14) **Suffix-initial: light acceptable, dark "?"**

free-ly [ˈfriː**l**i], ?[ˈfri͡ə**ɫ**i]
dai-ly [ˈde͡ɪ**l**i], ?[ˈde͡ə**ɫ**i]
gray-ling [ˈgɹe͡ɪ**l**ɪŋ], ?[ˈgɹe͡ə**ɫ**ɪŋ]
eye-let [ˈa͡ɪ**l**ət], ?[ˈa͡ɪə**ɫ**ət]

The forms examined in the survey were as follows:

(15)

Word	**with [l]:**		**with [ɫ]:**	
	mean	σ	mean	σ
gray-ling [9]	1.39	0.49	3.17	2.32
gai-ly	1.45	0.76	3.65	2.14
free-ly	1.85	1.25	3.20	1.81
average	1.57	0.87	3.34	2.03

Contrariwise, supposing that a vowel-initial suffix is added to a stem ending in /l/, the form with dark [ɫ] is preferred, and the light [l] form deserves the "?":

[9] To assistance intuitions of morphological relatedness, this was elicited in the frame: "A grayling is a kind of trout with a gray color." The words *mail-er* and *hail-y* were elicited with similar frames.

(16) **Stem-final before vowel-initial suffix: dark acceptable, light "?"**

(touchy-)feel-y	[ˈfi͡əɫi], ?[ˈfiː**l**i]
heal-ing	[ˈhi͡əɫɪŋ], ?[ˈhiː**l**ɪŋ]
mail-er *'one who mails'*	[ˈme͡əɫɚ], ?[ˈme͡ɪ**l**ɚ]

Here are the forms tested with the consultants:

(17)

Word	**with [l]:**		**with [ɫ]:**	
	mean	σ	mean	σ
mail-er	2.80	2.20	2.00	1.41
hail-y	4.00	1.80	1.56	1.01
gale-y	3.39	2.42	2.28	1.86
(*touchy-*)*feel-y*	2.00	1.49	2.20	1.87
average	3.01	2.06	2.01	1.54

Such influences of morphology have been observed before in other dialects: see Wells (1982, 312-313 for vernacular London English, Simpson (1980) for Australian English, and Gimson (1970, 202) for Standard British (RP) in a limited (postatonic) environment.

Stem-final /l/ can also become prevocalic when the stem precedes a vowel-initial word, as in *mail it*. The result here seems to be an exaggerated version of the preceding case: the /l/ "wants" quite strongly to be dark. This preference appeared in the single such form checked in the consultant survey:

(18)

Word	**with [l]:**		**with [ɫ]:**	
	mean	σ	mean	σ
mail it	4.40	1.71	1.10	0.32

One remaining word was elicited in the survey: *antler*, as an example of postconsonantal, pre-atonic /l/. The consultants strongly preferred light [l], but I was surprised by the relatively high level of acceptability of the dark [ɫ] variant:

(19)

Word	**with [l]:**		**with [ɫ]:**	
	mean	σ	mean	σ
antler	1.40	0.70	3.55	1.54

The relevant contrast here is with *Louanne* above, where the /l/ is likewise pre-atonic and non-post-vocalic, but the judgment of the dark variant is far harsher (3.55 versus 5.55; the difference is highly significant statistically). For now, I will leave *antler* aside, taking it up again in section 3.2.6.

We can now turn to the most crucial aspect of the survey: the demonstration that the differences in judgment between the various categories are not just the result of random fluctuations (which admittedly are present in abundance), but are an authentic effect of the underlying structural differences. To show this, I reduced the judgments for each datum (i.e., each individual word elicited from each individual consultant) to a single number, namely the light [l] judgment value subtracted from the dark [ɫ] value. The higher this number is, the greater the preference for light [l], with the overall range going from -6 to +6. Gathering all such numbers for the different categories of cases, I then carried out an ANOVA test, using the Fisher's PLSD post-hoc test for significance. This establishes the probability *p* that the difference in values could have been obtained by accident, given the amount of random variation present. Here are the results:

(20)

Word set	**average difference score**	**significance**
a. *light, Louanne*	4.62	
		p < .0001
b. *gray-ling, gai-ly, free-ly*	1.78	
		p = .0527
c. *Mailer, Hayley, Greeley, Daley*	0.74	
		p = .0006
d. *mail-er, hail-y, gale-y, feel-y*	-0.97	
		p = .0031
e. *mail it*	-3.30	
		p = .0021
f. *mail, help*	-5.47	

It can be seen that all the results but one are highly significant statistically, and that the remaining one is near-significant. Keeping the near-significant outcome as a case where further checking should be done, I will assume for present purposes that all differences given here should be accounted for in an adequate analysis.[10]

To keep the size of the problem under control, I will further reduce the numerical data of the survey to the traditional categories "✓", "?", "??", and "*". The categories assigned, with the survey numbers used to justify them, are given below:

[10] Some further statistics: for non-adjacent categories on the scale of (20), all comparisons came out significant; p < .0001. Further, of the word-to-word comparisons *within* groups, only one came out significant, namely *feel-y* versus *hail-y* (p = .0137). This suggests that most of the relevant structural differences have probably been located.

The individual consultants gave so few judgments that few subject-internal results reached significance, but the profiles of individual subjects tend to resemble (20).

(21)

	as light		as dark	
light, Louanne	✓	(1.20)	*	(5.83)
gray-ling, gai-ly, free-ly	✓	(1.57)	?	(3.34)
Mailer, Hayley, Greeley, Daley	✓	(1.90)	✓	(2.64)
mail-er, hail-y, gale-y, feel-y	?	(3.01)	✓	(2.01)
mail it	??	(4.40)	✓	(1.10)
bell, help	*	(6.60)	✓	(1.12)

In principle, one might analyze more finely, but given the uncertainties and high standard deviations, it seemed advisable to work with a fairly coarse well-formedness grid.

3.1.4 The "Performance" Issue

Before going on the to the analysis of these data, I wish first to address a crucial potential objection mentioned above. Suppose that linguistic well-formedness really is an all-or-nothing matter, but that the judgments we get are filtered through various performance mechanisms. It is the performance mechanisms, not the grammar itself, which results in the gradient intuitions. If this is so, then it is not legitimate to attempt to analyze the gradient data of (21); rather, we should be reducing them (somehow) to two categories and developing a model of the judgment process itself to account for the numbers observed.

A problem with this view, as Schütze (1996, 64) has pointed out, is that patterns of gradient well-formedness often seem to be driven by the very same principles that govern absolute well-formedness. This holds true, for instance, for the phenomena under discussion here. Thus, for RP British English, Gimson (1970) and Wells (1982) report a data pattern that is reminiscent of what is described above, with gravitation of light and dark /l/ to pre-and post-vocalic positions and certain morphological effects similar to what we have seen. But there is also an important difference: RP does not allow dark [ɫ] in words like *free-ly* or *gray-ling* at all.[11] This categorial prohibition is rather likely to be based on the same principles that govern the subtler intermediate judgment of the American speakers.

I conclude that the proposed attribution of gradient well-formedness judgments to performance mechanisms would be uninsightful. Whatever "performance" mechanisms we adopted would look startlingly like the grammatical mechanisms that account for non-gradient judgments. For this reason, I will assume that the competence model itself should generate gradient judgments, and will now turn to the task of accounting for the data in (21).

3.2 Analysis

[11] Thanks to Peter Ladefoged, a native speaker, for confirming this judgment.

3.2.1 Constraints

Following traditional views in phonetics, we can plausibly attribute the variation in /l/ to conflicting principles based on articulation and perception. The loss or diminution of alveolar closure in dark [ɫ] seems fairly plainly a case of *lenition*, a process widely thought to be grounded in the conservation of articulatory effort.[12] I state the relevant constraint as follows:

(22) /l/ IS DARK

This leaves open exactly what "darkness" is in English: roughly, it should be characterized as involving a lenited, delayed, or absent tongue blade closure. Often, in compensation, there is an especially backed tongue body position.

A defect of the formulation in (22) is that it is categorial, not gradient. The expression of a constraint that would require darkness in /l/ in gradient fashion would require further theoretical development that goes beyond the scope of this article. I believe the central point at hand can be made, however, with the artificially categorized constraint given here.

The other constraints on /l/ are perceptual in origin, and are more subtle in character. An important finding of Sproat and Fujimura's (1993) X-ray microbeam study is that English dark [ɫ] is temporally *asymmetrical*: it begins with a tongue body backing gesture and then in most cases continues with the blade-raising gesture. In light [l], the blade raising gesture is invariant and robust, and often comes somewhat earlier than the tongue body backing gesture.

Sproat and Fujimura's findings tie into a recent proposal made by Steriade (1997), who studies the phonotactics of *temporally-asymmetrical* segments, such as aspirated and preaspirated stops, ejectives, preglottalized sonorants, and (surprisingly) retroflexes. Her general finding is that if the acoustic cues for a particular consonant lie on one side of the consonant, then there is a very strong tendency for phonologies to require that side of the consonant to be vowel-adjacent. Thus aspirated stops are often limited to prevocalic position, preaspirated stops to postvocalic position.

Applying Steriade's principle to the present case, one would expect that languages would be likely to limit dark [ɫ] to postvocalic position. It is plain that [ɫ] imposes a massive degree of coarticulation on the preceding vowel, as is attested in the allophone chart of (5). Plausibly, this coarticulation plays a major role in rendering dark [ɫ] identifiable. If dark [ɫ] were allowed to occur non-postvocalically, it would be harder to detect, and would in particular risk being confused with /w/, to which it is acoustically similar. Thus in Steriade's general view, there is good support for a constraint of the following character:

(23) DARK [ɫ] IS POSTVOCALIC

[12] See Kirchner (to appear) for extended discussion of constraints covering lenition.

Let us now consider the opposite side of the /l/, namely its articulatory release. Observations of prevocalic [l] on spectrograms show an important effect of the more-robust tongue blade gesture: at release, it produces a rapid amplitude rise and vivid formant shifts. These plausibly would form strong cues to the presence of prevocalic light [l]. The expenditure of articulatory effort in forming a light [l] is thus most effective in this context. The grammatical reflex of these considerations is postulated to be the following:

(24) PREVOCALIC /l/ IS LIGHT

Somewhat more tentatively, I further conjecture that a *stressed* vowel forms the best acoustic backdrop for the formant transitions created in a light /l/. It is certainly true that stressed vowels license another important temporally-asymmetrical class of sounds in English, namely the aspirated stops, so this seems a plausible assumption. This leads to the similar, but generally stricter constraint below:

(25) PRETONIC /l/ IS LIGHT

The informal predicates "light", "pretonic", and "prevocalic" used here should be assumed to have suitable expressions in a formalized theory. The exact formulation chosen is noncrucial to present purposes.

In some of the data (cases (14) and (16)), the presence of a morpheme boundary before or after the /l/ is crucial. For these, I will assume constraints that limit alternation among surface allomorphs, of the type propounded by McCarthy and Prince (1995), Benua (1995), Kenstowicz (1996, 1997), Steriade (1996), Burzio (1997), and others. Referring to these as Paradigm Uniformity constraints, I state the relevant such constraint here as follows, in schematic form:

(26) PARADIGM UNIFORMITY

Morphologically derived forms may not deviate from their bases in Property X.

Let us leave aside for a moment the issue of what Property X is in the present case. What is crucial is that in (say) *touchy-feel-y*, surface [ˈfi͡əɫi] possesses very much the same phonetic form as that which is observed in the base form *feel* [ˈfi͡əɫ]. This is not the case in the imperfect rival form ?[ˈfiːli]. Similarly, in *freely* [ˈfɹiːli], the segments show very much the same phonetic quality that one observes in related forms like *free* [ˈfɹiː] and *barely* [ˈbɛɹli]; this is not so in the imperfect rival ?[ˈfɹi͡əɫi].

Imitation of base forms by derived forms is commonplace in phonology, and has in pre-OT approaches been analyzed with notions like cyclicity and word-internal boundaries (Chomsky and Halle 1968). It is not surprising that we should find them at work in /l/ allophony.

Let us now try to specify the identity of Property X, the property that is being conserved across the paradigm. At first glance, it looks like X might be /l/ darkness itself: thus, the dark /l/ of *feel* [ˈfi͡əɫ] must be carried over into the derived form *feel-y*. However, for forms like *freely* with /l/-initial suffixes, this view is more dubious since the suffix *-ly* has no isolation form that could serve as the light [l] base.

A more perspicuous analytic path would be to impose Paradigm Uniformity on *the quality of the stem vowel*, which as noted in (5), is always quite different before dark [ɫ] than elsewhere. Under this view, the diphthongized vowel of *feel* ([ˈf**i͡ə**ɫ]) is carried over into *touchy-feel-y* ([ˈf**i͡ə**ɫi]), and the non-diphthongized vowel of *free* [ˈfɹ**iː**] is carried over into *freely* ([ˈfɹ**iː**li]).

For this to work, the system must include undominated constraints specifying the appropriate matchup of light and dark /l/ with the vowel allophones that go with them, as noted above in (5). This will exclude any candidates like *[ˈf**i͡əl**i] or *[ˈf**iːɫ**i]. As the exact formulation of these constraints is not crucial, I will not attempt it here.[13]

Summing up, I will assume that in the present case, Paradigm Uniformity requires morphologically derived forms to possess the vowel quality of their bases, and that this indirectly regulates the distribution of light and dark /l/.

There is one further elaboration needed for the Paradigm Uniformity phenomena seen here. Consider that both *mail-er* ([ɫ, ?l]) and *mail it* [ɫ, ??l] involve Paradigm Uniformity effects, with carry-over of the obligatory dark [ɫ] of *mail* into larger morphosyntactic constructions. But the effect is *stronger* in the phrasal construction than in the suffixed form: specifically, light [l] gets a "??" (consultant average 4.40) in *mail it* but only a "?" (= 3.01) in *mailer*. Why should this be so?

It is commonly observed that phonological alternation tends to be inhibited in relatively larger phonological domains. This typological observation has been translated into various theoretical approaches in various ways; thus, for example, in one version of Lexical Phonology (Kiparsky 1985, 87), rules are held to be gradually "turned off" as one reaches later levels of the grammar. Similarly, hierarchies of boundaries (McCawley 1968) or "P-structure" (Selkirk 1981) are set up to permit less rule application (and thus less alternation) at higher levels.

In the present approach, employing Optimality Theory and Paradigm Uniformity, an appropriate implementation of this idea would be to suppose that the Paradigm Uniformity constraints are *a priori* stricter for higher levels; for example, stricter in phrases than in words. For the case at hand, we can suppose that there are separate constraints of Paradigm Uniformity for phrasal versus morphological contexts, with the former ranked within UG as necessarily stricter than the latter. Thus:

---

[13] Carlos Gussenhoven raises the important question of whether a phonological constraint can refer to surface vowel allophones before [ɫ], which may be coarticulatory in origin. There are two possibilities here: (a) as Gussenhoven suggests, the pre-[ɫ] allophones might be truly phonological, derived perhaps by feature spreading; (b) paradigm uniformity has access at least in some degree to phonetic representation.

(27) PARADIGM UNIFORMITY(VOWEL QUALITY, PHRASAL) »
(28) PARADIGM UNIFORMITY (VOWEL QUALITY, MORPHOLOGICAL)

With this distinction in place, we now have the constraints that will be needed to derive the correct outcomes. The constraints are listed below with the abbreviations that will be used.

(29) /l/ IS DARK	/l/ IS DARK
DARK [ɫ] IS POSTVOCALIC	[ɫ] IS /V___
PREVOCALIC /l/ IS LIGHT	PREVOCALIC:[l]
PRETONIC /l/ IS LIGHT	PRETONIC:[l]
PARADIGM UNIFORMITY(VOWEL QUALITY, PHRASAL)	PU(PHRASAL)
PAR. UNIFORMITY (V QUALITY, MORPHOLOGICAL)	PU(MORPHOL)

What I have tried to argue in this section is that each constraint is *principled*, being grounded either in general concepts of phonetically driven phonology or in widely attested typological patterns of paradigmatic alternation.

3.2.2 Generating One Set of Outcomes

For purposes of presentation only, let us temporarily commit the methodological sin of overidealization, and generate an invariant set of outcomes, with no free variation and no gradient well-formedness. These reified outcomes may be schematized as follows: [l] in *light* (pretonic position), *Louanne*, (initial position), *Greeley* (medial pre-atonic) and *freely* (suffix-initial), and [ɫ] in *bell* (prepausal), *help* (preconsonantal), *feel-y* (stem-final before suffix), and *mail it* (word-final before vowel-initial word). To generate these outcomes, it suffices to rank the two PARADIGM UNIFORMITY constraints over PREVOCALIC /l/ IS LIGHT, with the latter dominating /l/ IS DARK. The outcomes emerge as follows:

(30)		PRETON-IC:[l]	[ɫ] IS /V_	PU(PHRA-SAL)	PU(MOR-PHOL)	PREVO-CALIC:[l]	/l/ IS DARK
a. light	☞ [**l**a͡ɪt]						*
	*[**ɫ**a͡ɪt]	*!	*			*	
b. Louanne	☞ [**l**uˈæn]						*
	*[**ɫ**uˈæn]		*!			*	
c. gray-ling	☞ [ˈɡɹe͡ɪ-**l**ɪŋ]						*
	*[ˈɡɹe͡ə-**ɫ**ɪŋ]				*!	*	
d. Greeley	☞ [ˈɡɹiː**l**i]						*
	*[ɡɹi͡ə**ɫ**i]					*!	
e. mail-er	☞ [me͡ə**ɫ**-ɚ]					*	
	*[me͡ɪ**l**-ɚ]				*!		*
f. mail it	☞ [ˈme͡ə**ɫ** ɨt]					*	
	*[ˈme͡ɪ**l** ɨt]			*!			*
g. help	☞ [ˈhɛ͡ə**ɫ**p]						
	*[ˈhɛ**l**p]						*!

Discussion of this tableau is slightly tricky, since some of the rankings turn out to be crucial only later on, when we cover free variation and gradient well formedness. Tentatively, we can say the following.

- Light [l] in (30d) *Greeley* is forced by the dominance of PREVOCALIC /l/ IS LIGHT over the lenition constraint /l/ IS DARK.
- Light [l] in (30a) *light* shows the same effect, though as it happens the light outcome here is forced by higher-ranking constraints as well: PRETONIC /l/ IS LIGHT and DARK [ɫ] IS POSTVOCALIC.
- Light [l] in *Louanne* (30b) also follows from PREVOCALIC /l/ IS LIGHT » /l/ IS DARK, but again the high-ranking DARK [ɫ] IS POSTVOCALIC also suffices to rule out the dark [ɫ] candidate.
- The same thing is true for (30c) *gray-ling*, where the higher-ranked constraint is PARADIGM UNIFORMITY (VOWEL QUALITY, MORPHOLOGICAL).
- In (30g) *help* (similarly *bell*), the constraint violations are in a subset relation, so no matter what ranking is adopted, dark [l] will win.
- The cases with the greatest interest here are (30e,f) *mail-er* and *mail it*. In these, the PARADIGM UNIFORMITY constraints (morphological and phrasal, respectively) force the appearance of dark [ɫ], despite the appearance of the /l/ in prevocalic position. The morphological base that drives both PU constraints is *mail*, which itself receives obligatory dark [ɫ] for the same reasons as *help*.

3.2.3 Accounting for Free Variation

Moving onward from this preliminary sketch toward a more accurate model, we can next account for the fact that monomorphemic forms with intervocalic pre-atonic /l/ (e.g. *Greeley*) show free variation.

To do this, let us assume that PREVOCALIC /l/ IS LIGHT and /l/ IS DARK are ranked freely. This means that all outputs derived under the tableau of (30), where PREVOCALIC /l/ IS LIGHT outranks /l/ IS DARK, will still be obtainable, but in addition we will obtain whatever outcomes derive from the opposite ranking, as shown below:

(31)		PRETON-IC:[l]	[ɫ] IS /V_	PU(PHRA-SAL)	PU(MOR-PHOL)	/l/ IS DARK	PREVO-CALIC:[l]
a. light	☞ [la͡ɪt]					*	
	*[ɫa͡ɪt]	*!	*				*
b. Louanne	☞ [luˈæn]					*	
	*[ɫuˈæn]		*!				*
c. gray-ling	☞ [ˈgɹe͡ɪ-lɪŋ]					*	
	*[ˈgɹe͡ə-ɫɪŋ]				*!		*
d. Greeley	☞ [gɹi͡əɫi]						*
	*[ˈgɹiːli]					*!	
e. mail-er	☞ [me͡əɫ-ɚ]						*
	*[me͡ɪl-ɚ]				*!	*	
f. mail it	☞ [ˈme͡əɫ ɨt]						*
	*[ˈme͡ɪl ɨt]			*!		*	
g. help	☞ [ˈhɛ͡əɫp]						
	*[ˈhɛlp]					*!	

As it happens, the only outcome that is altered under the new ranking is that for monomorphemic *Greeley* (31d), which now comes out with a dark [ɫ], since the lenitional constraint /l/ IS DARK outranks PREVOCALIC /l/ IS LIGHT. In all other instances, switching these two constraints makes no difference to the outcome: either a higher-ranking constraint has already decided the issue (*light*, *Louanne*, *gray-ling*, *feel-y*), or the loser has a superset of the violations of the winner (*help*), so that no reranking would ever make a difference.

Here are two further ranking arguments that emerge from tableau (31b):

- It can be seen that the continued appearance of light [l] in (31b) *Louanne* justifies the invariant ranking DARK [ɫ] IS POSTVOCALIC » /l/ IS DARK, a ranking which could not be justified in the previous tableau.
- Likewise, the continued appearance of light [l] in (31b) *gray-ling* follows from the invariant ranking PU(MORPHOL) » /l/ IS DARK.

3.2.4 The Analysis in Strictness Bands

Recall from above that in the present proposal, free constraint rankings result from overlapping strictness bands. We can translate the analysis so far into a strictness band approach. It suffices here to assign the two constraints PREVOCALIC /l/ IS LIGHT and /l/ IS DARK to overlapping bands at the bottom of the scale, with all other constraints assigned to a group at the top of the scale. There must be no overlap of the two groups.

(32)

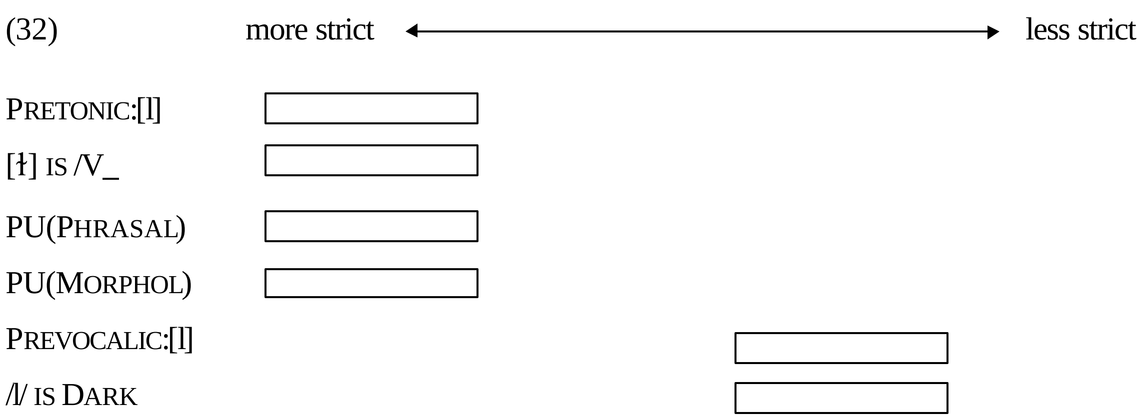

As it happens, the overlaps in the strictness bands are non-crucial in all cases except for the two laxest constraints; for the four strict constraints, the data considered so far would be compatible with many other arrangements as well. These could include, for instance, point-like strictness values, with zero widths.

3.2.5 Accounting for Gradient Well-Formedness

Let us now take the final step in developing the full analysis, namely accounting for the gradient well-formedness of certain forms. To start, recall the cases that yielded intermediate well-formedness judgments in the consultant survey.

- Forms like *gray-ling* tend to preserve the pure, uncoarticulated quality of the stem vowel ([e͡ɪ] of *gray* [ˈgɹe͡ɪ]), rather than adopting the altered, backed quality ([e͡ɪə]) that is elsewhee found before dark [ɫ]. Since the constraints (unstated here) that correlate vowel quality and /l/ darkness are undominated, this means that *gray-ling* prefers light [l]: [ˈgɹe͡ɪlɪŋ], ?[ˈgɹe͡ɪəɫɪŋ].

- Forms like *mail-er* tend to preserve the dark [ɫ] and coarticulated vowel quality of their bases; thus [ˈmeɪ͡əɫɚ] is preferred over ?[ˈme͡ɪlɚ], because of the base form *mail* [ˈmeɪ͡əɫ].

- Forms like *mail it* work exactly the same way, only the judgment is strong; thus [ˈmeɪ͡əɫ ɨt], ??[ˈme͡ɪl ɨt].

The basic approach to be followed here starts out as follows: the two constraints which drive the appearance of light and dark /l/ in the crucial context (intervocalic, pre-atonic) are PREVOCALIC /l/ IS LIGHT and /l/ IS DARK. These two opposed constraints are ranked freely, as is demonstrated by the variation in monomorphemic forms like *Greeley*. But in suffixed forms, neither constraint is quite "strong enough" to override the Paradigm Uniformity constraints that enforce light and dark /l/ based on what appears in the isolation stem.

But thus far we have considered only fringeless constraints. Once fringes are added in, greater descriptive accuracy become possible. Let us suppose that the *fringes* of PREVOCALIC /l/ IS LIGHT and /l/ IS DARK, unlike their central ranges, extend upward into the areas occupied by the Paradigm Uniformity constraints. If this is so, it becomes possible to analyze the gradient well-formedness judgments. The specific arrangements of bands and fringes required are as follows.

- Only the upper "?" fringe of PREVOCALIC /l/ IS LIGHT, and not its central region, extends high enough to permit rankings in which PREVOCALIC /l/ IS LIGHT dominates PARADIGM UNIFORMITY (VOWEL QUALITY, MORPHOLOGICAL). This means that forms crucially derived with this ranking (specifically, *mai*[l]-*er* and similar cases) will receive a "?" under the analysis.

- Likewise, only the upper "?" fringe of /l/ IS DARK, and not its central region, extends high enough to permit rankings in which /l/ IS DARK dominates PARADIGM UNIFORMITY (VOWEL QUALITY, MORPHOLOGICAL). This means that forms crucially derived with this ranking (specifically, *gray*-[ɫ]*ing* and similar cases) will also receive a "?".

- Only the upper "??" fringe of PREVOCALIC /l/ IS LIGHT (not its central region, nor even its "?" fringe) extends high enough to permit rankings in which PREVOCALIC /l/ IS LIGHT dominates PARADIGM UNIFORMITY (VOWEL QUALITY, PHRASAL). It is for this reason that ??*mai*[l] *it* receives its highly marginal status.

A further comment: the fact that one need use only the "?" fringe of PREVOCALIC /l/ IS LIGHT to rank it above PARADIGM UNIFORMITY (VOWEL QUALITY, MORPHOLOGICAL), but the "??" fringe to rank it above PARADIGM UNIFORMITY (VOWEL QUALITY, PHRASAL), is unlikely to be an accident. It will be recalled that, by a general and well attested (but so far unexplained) principle of phonology, phrasal paradigm uniformity is always ranked more highly than morphological paradigm uniformity;[14]

[14] I assume that this holds true for any particular choice of selection points, not necessarily for the bands themselves.

thus it is alway “harder” for a given constraint to outrank a phrasal paradigm constraint than to outrank a morphological one. It is in light of this general background that we should not be surprised to see ??*mai*[l] *it* receive worse judgements than ?*mai*[l]*er*.

Summing up this preview: the intermediately well formed cases are posited to receive this status because they force the ranking of the basic constraints favoring light and dark /l/ into regions where they can dominate the relevant Paradigm Uniformity constraints. Such rankings require the use of the “?” and “??” fringes. To make the analysis fully explicit, here is the full set of strictness bands and fringes posited here:

(33) more strict ⟵⟶ less strict

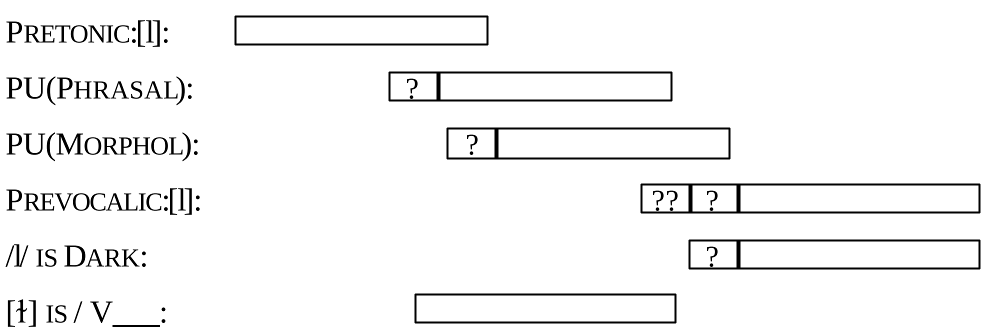

To obtain predictions about well-formedness from these bands and fringes in a rigorous way takes a bit of work. First, one must determine the *pairwise rankings* that result from the strictness bands, including those that are free rankings, or which arise when selection points occupy fringes. Since six constraints give rise to 15 pairwise rankings, the following is an exhaustive list:

(34) a. Ranking is **irrelevant** because there are no empirical consequences:

PRETONIC:[l] w.r.t. [ɫ] IS / V___	(Both exclude [ɫ].)
PRETONIC:[l] w.r.t. PREVOCALIC:[l]	(Both exclude [ɫ].)
[ɫ] IS / V___ w.r.t. PREVOCALIC:[l]	(Both exclude [ɫ].)
[ɫ] IS / V___ w.r.t. PU(MORPHOL)	(No bases are required to start with [ɫ].)
[ɫ] IS / V___ w.r.t. PU(PHRASAL)	(No bases are required to start with [ɫ].)
PU(MORPHOL) w.r.t. PU(PHRASAL)	(Constraints apply in disjoint domains.)

b. Ranking is **free**, because central ranges of strictness bands overlap:

PREVOCALIC:[l] w.r.t. /l/ IS DARK: *Gree*[l,ɫ]*ey*, etc.

c. Ranking shown is possible only by placing a selection point within the **"?" fringes** of a constraint (opposite ranking is freely available):

PREVOCALIC:[l] » PU(MORPHOL): *mai*[?l, ɫ]*er*, etc.
/l/ IS DARK » PU(MORPHOL): *gray*-[l, ?ɫ]*ing*, etc.

d. Ranking shown is possible only by placing a selection point within the **"??" fringes** of a constraint (opposite ranking is freely available):

PREVOCALIC:[l] » PU(PHRASAL): *mai*[??l, ɫ] *it*, etc.

e. Ranking is **obligatory**, because strictness bands fail to overlap even in the fringes:

[ɫ] IS / V___ » /l/ IS DARK: [l,*ɫ]*ouanne*, etc. (but see below, section 3.2.6)

f. Ranking **cannot be determined** with data given (see section 3.2.6 for discussion):

PRETONIC:[l] w.r.t. /l/ IS DARK
PRETONIC:[l] w.r.t. PU(MORPHOL)
PRETONIC:[l] w.r.t. PU(PHRASAL)
/l/ IS DARK w.r.t PU(PHRASAL)

Once these specific rankings have been established, then one can assign "*", "??", and "?" by means of a careful search. For each pairing of input form and output candidate, one seeks to find: (a) if *any* choice of selection points within the bands yields a grammar in which the output candidate defeats all rivals and is thus generated; (b) if so, which choice minimizes the use of "??" fringes and "?" fringes, in that order of priority. I assume that the rating given by consultants will generally correspond to the best available choice of strictness points.

Carrying out this check is rather tedious, and I will not recapitulate here the work that was done. , other than to list in a footnote an algorithm that eases the task.[15] The relevant tableaux may be

[15] (a) Define **R** as the set of constraint rankings assumed at any particular stage of computation to be obligatory. To begin, examine the strictness bands of the constraints and set **R** equal to the cases where a ranking *always* holds, even when the selection points are permitted to occur within "??" and "?" strictness bands (these are the rankings that result from total non-overlap of strictness bands).

(b) For each possible output **O**, determine whether **O** can be derived from its input (that is, whether **O** wins the standard OT competition among rivals for output status) under at least one constraint ranking that is consistent with the rankings of **R**. If **O** cannot be so derived, assign it "*".

(c) Add to **R** all pairwise constraint rankings that are possible only when a selection point occupies a "??" fringe. Repeat step (b), but this time, if an output O cannot be derived, assign it "??". All outputs assigned "*" earlier on keep their "*" status.

(d) Add to **R** all pairwise constraint rankings that are possible only when a selection point occupies a "?" fringe. Repeat step (c), this time assigning "?" to all nongenerated outputs. (As before, all outputs assigned "*" or "??" earlier on keep their status.)

downloaded from the Web site listed at the end of this article. From the tableaux, it can be shown that the result that emerges is indeed was what I had originally aimed at in the analysis, namely:

(35)

	as light	**as dark**
light, Louanne	✓	*
gray-ling, gai-ly, free-ly	✓	?
Mailer, Hayley, Greeley, Daley	✓	✓
mail-er, hail-y, gale-y, feel-y	?	✓
mail it	??	✓
bell, help	*	✓

To summarize the various cases:

- It is indeed possible to generate ?*gray*-[ɫ]*ing*, ?*mai*[l]-*er*, and ??*mai*[l] *it*, but only by placing a selection point within the "?" fringe of /l/ IS DARK, the "?" fringe of PREVOCALIC /l/ IS LIGHT, and the "??" fringe of PREVOCALIC /l/ IS LIGHT, respectively.

- The essentially free variation between *Gree*[l]*ey* and *Gree*[ɫ]*ey* works just as it did in (31) above.

- In *light*, the strict domination of /l/ IS DARK by PRETONIC /l/ IS LIGHT forces a "*" for the dark [ɫ] candidate.

- Likewise, complete ill-formedness for dark [ɫ] in *Louanne* is guaranteed by strict domination of /l/ IS DARK by DARK /l/ IS POSTVOCALIC.[16]

- *Bell* and *help* receive obligatory dark [ɫ]'s, because there is no constraint in the system that would force light ones.

In the last of these cases, it is striking that the constraint forcing dark outcomes is one of the two weakest in the system, namely /l/ IS DARK. Yet the intuition that light-[l] *bell* or *help* is ill-formed is very strong, a point verified by the consultants. This demonstrates, if demonstration were needed, that the "lowness" in the tableaux at which a form is excluded has nothing to do with its degree of well-formedness.

(e) Assign fully well-formed status to all outputs that can be generated at stage (d).

This procedure will duly sort the outputs into those which cannot be derived at all, those which can be derived only by using the "??" strictness fringes, those which can be derived only by using the "?" strictness fringes, and those which can be derived using just the central strictness bands.

What makes the procedure workable is factorial typology software, which finds all outcomes of the constraint set at each stage, as the rankings are gradually tightened. Factorial typology software exists in forms prepared by Raymond and Hogan (1993) and by the author (available from the Web site listed at the end of this article.)

[16] See section 3.2.6 for further discussion of this form, however.

3.2.6 Residual Cases

There are a few cases that require additional discussion. None of these bears on the critical cases (of the *Greeley*, *gray-ling*, and *mail-er* classes) on which the treatment of gradient well-formedness rests.

(a) The data obtained from the consultant survey did not establish the ranking of PRETONIC /l/ IS LIGHT with PU(MORPHOL). This ranking hinges on the well-formedness of words like *stylistic* (from *style* [staɪ͡əɫ]) or *alcoholism* (from *alcohol* [ˈa͡əɫkəˌhɒɫ]). If such words tolerate light and dark /l/ equally, then these two constraints should overlap in their central ranges; but if light [l] is preferred then the lower end of PRETONIC /l/ IS LIGHT should overlap only with the upper "?" fringe of PU(MORPHOL). In laying out the strictness bands of (33), I made the latter assumption, based on my own judgment.

(b) Again, guessing from my own intuitions, I conjecture that forms such as *mail it* ??[me͡ɪl ɨt], with phonological alternation between contextual [me͡ɪl] and the isolation form [me͡əɫ], occur only before clitic pronouns. Forms where the alternation crosses a large phonological juncture, like *mail Italian* (*books*) *[ˈme͡ɪləˈta͡əɫjən...], seem inconceivable to me; Thus, it seems likely that there is a still stricter PU constraint that likewise forbids alternation, but is limited to higher levels of phonological juncture, such as the Phonological Phrase (Selkirk 1981, Nespor and Vogel 1986, Hayes 1989). This constraint is plausibly undominated.

(c) Forms like *allow*, with medial pretonic /l/, were not included in the survey, but based on my own judgment I suspect that they would emerge with obligatory light [l]. This outcome is derived in the present grammar by having PRETONIC /l/ IS LIGHT fully outrank /l/ IS DARK.

(d) A difficulty for the analysis is the surprisingly high acceptability of dark [ɫ] in *antler*, as determined in the survey (value: 3.55). This is considerably better than the 5.55 awarded to dark-[ɫ] *Louanne*, and the difference is indeed statistically significant. But the two words have virtually identical /l/ environments: non-postvocalic, pre-atonic with no Paradigm Uniformity effects.

I see two possible ways around the problem. First, it may be that the survey data for *Louanne* were unreliable: Sproat and Fujimura (1993) collected a couple of tokens of "*Likkovsky*", analogous to *Louanne*, with surprisingly dark /l/ (see their data for speaker CS, p. 302). I think it likely that my own renditions of *Louanne* and *antler*, in performing the survey, were skewed towards a really grotesquely dark [ɫ] in *Louanne* and only a moderately dark one in *antler*.[17] This naturally would have skewed the results. If this is right, then we should adjust the ranking of DARK /l/ IS POSTVOCALIC

[17] This is less of a problem for the really crucial cases of the *Greeley*, *mail-er*, and *gray-ling* classes, because for these a technique was available to achieve more natural pronunciations. Specifically, I pronounced light and dark /l/ by aiming at a *different word*, where necessary, from what the consultant was listening for. Thus, for example, to say *Greeley* with guaranteed light /l/ I thought to myself "*gree-ly*; in a manner which is *gree*"; and to say it with a guaranteed dark [ɫ] I thought to myself "*greel-y*; covered with *greel*". In contrast, saying *Louanne* with a dark [ɫ] was much harder to do naturally, and felt much more like a phonetic exercise.

downward slightly, so that it overlaps the upward "?" or "??" fringes of /l/ IS DARK. This would let in both *ant*[ɫ]*er* and [ɫ]*ouanne* on a "?" or "??" basis.

A second possibility is that more careful study would prove that dark-[ɫ] *Louanne* really is worse than dark-[ɫ] *antler*. If this is so, a possibility to consider is that stem-initial position calls for allophones that are especially salient acoustically, as Kohler (1990, 88) suggests. That something like this is right is also indicated by phonological work (Steriade 1993; Beckman 1995, 1997; Casali 1996) showing that stem-initial position favors particularly high-ranked faithfulness constraints; the same might well be true for salient allophones.

None of these questions is resolvable with the data presently gathered. However, for the crucial cases demonstrating gradient well-formedness (see the preceding section), the data and analysis seem rather more secure.

3.2.7 Variation in the Consultants' Intuitions

So far, I have simply pooled the opinions of my ten consultants, as if there were no difference in their opinions. For one case, however, it appears that there may have been a genuine split in their behavior.

Under the analysis proposed here, the judgment that one gives for forms like *mail-er* is logically independent of the judgment one gives for forms like *gray-ling*: light [l] will be tolerated in *mail-er* to the extent that one is able to rank PREVOCALIC /l/ IS LIGHT over PU(MORPHOL); whereas dark [ɫ] will be tolerated in *gray-ling* to the extent that one is able to rank /l/ IS DARK over PU(MORPHOL).

My own judgment, or at least the one I brought to the project originally, is that *gray*-[ɫ]*ing* with dark [ɫ] is considerably worse than *mai*[l]-*er* with light [l]; in fact I would rate *gray*-[ɫ]*ing* as "??" and *mai*[l]-*er* as "?". That this is not an isolated intuition is borne out by the consultants' individual judgments on the *gray-ling* and *mail-er* classes of forms. Here, we find two groups; one which agrees with me in finding *mai*[l]-*er* forms better than *gray*-[ɫ]*ing* forms, and one which goes in the opposite direction:

(36) a. *mai*[l]-*er* etc. better than *gray*-[ɫ]*ing* etc.

	Average for *mai*[l]-*er* words	Average for *gray*-[ɫ]*ing* words
DP	2.17	2.67
FC	2.67	4.00
KH	1.50	4.33
MS	3.00	4.33
PH	2.50	3.67

b. *gray*-[ɫ]*ing* etc. better than *mai*[l]-*er* etc.

	Average for *mai*[l]-*er* words	Average for *gray*-[ɫ]*ing* words
MP	4.75	3.33
MG	1.75	1.00
VA	4.75	4.00

The remaining two consultants ranked the two classes about equally.

It is straightforward to model these cases. As an application of the method described in fn. 15 has verified, a grammar that gives a "??" instead of a "?" to *gray*-[ɫ]*ing* etc. (= (36a)) can be obtained by slightly lowering the range of /l/ IS DARK, so that only its "??" fringe overlaps the range of PU(MORPHOL); and a grammar that gives a "??" instead of a "?" to *mai*[l]-*er* etc. can be obtained with an analogous weakening of PREVOCALIC /l/ IS LIGHT. These detail adjustments for particular dialects can be seen below:

(37)

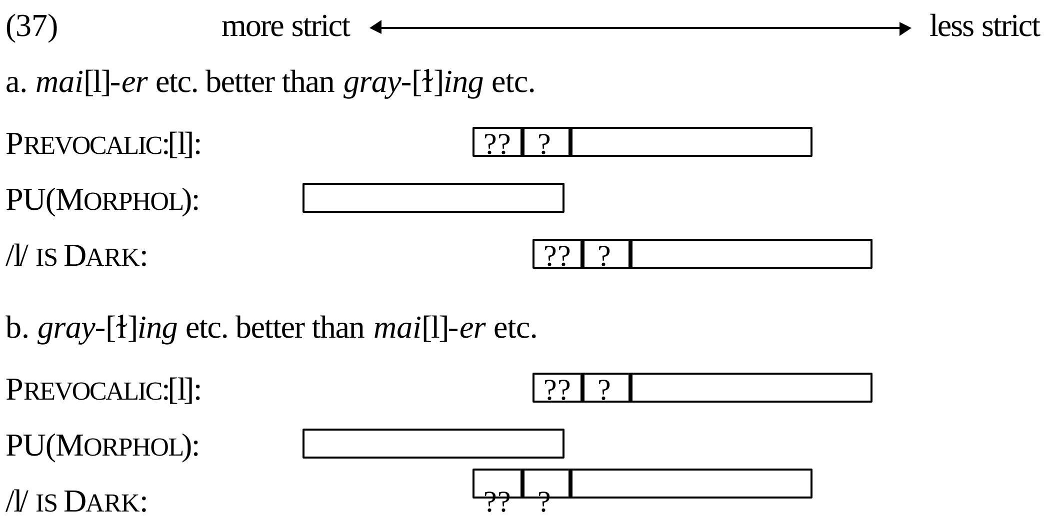

One would expect such grammars also in principle to differ on how they treat other forms as well; and indeed, the subjects that preferred (say) *gray*-[l]*ing* to *gray*-[ɫ]*ing* also tended to prefer *Gree*[l]*ey* to *Gree*[ɫ]*ey*, and so on.[18]

[18] For instance, across consultants there is a positive correlation between how the *gray-ling* forms and the *Greeley* forms were judged; for the statistic (dark judgments - light judgments) used above, r^2 = .638. Other relevant correlations were weaker but mostly positive.

3.2.8 Judgments of Speech Style

Recall that the questionnaire given to the consultants permitted them to judge the casualness" or "formality" of individual examples. They often had some difficulty in doing this; and not all tried to make such judgments. Nevertheless, a pattern did emerge in the judgments: in general, it is the output with *light* [l] that is felt to be formal relative to the one with dark [ɫ]. This result was statistically significant: testing the full set of judgments of light [l] against the full set of judgments of dark [ɫ], and treating a consultant's check mark in the "casual" box as -1 and a check mark in the "formal" box as 1 (with no check being counted as 0), we find means of .07 for light [l] versus -.11 for dark [ɫ]. This difference is shown to be statistically significant (p = .0002) by a Fisher's PLSD test. The difference emerged primarily in the *Greeley* and *gray-ling* word classes.[19]

I believe that the association of light [l] with formality, and dark [ɫ] with casualness, is not random, but reflects the nature of the constraints themselves. Recall that, where variation is possible, light [l] reflects relatively high ranking of the constraint PREVOCALIC /l/ IS LIGHT, while dark [ɫ] reflects relatively high ranking of /l/ IS DARK. Recall further (from section 3.2.1) that both of these constraints are grounded in phonetic principles: /l/ IS DARK is a lenitional constraint, whose teleology is speaker-centered, involving the maximization of articulatory ease. PREVOCALIC /l/ IS LIGHT is grounded in speech perception, calling forth an articulation that will render /l/ more identifiable for the listener (in a context where the more costly articulation will be more effective).

As phoneticians have suggested (Kohler 1990, Lindblom 1990), formal speech is characteristically listener-centered speech, aimed at constructing a maximally-decodable acoustic signal; whereas casual speech is speaker-centered, aimed at specifying an articulatory program that is easy to execute. Within the approach established here, it is easy to imagine the grammatical reflex of these strategies: in formal speech, the selection point for PREVOCALIC /l/ IS LIGHT will be chosen from the upper part of its range, and the selection point for /l/ IS DARK will be chosen from the lower part. In casual speech, exactly the opposite will tend to occur.

Assuming that these speaker strategies are also decodable to some extent by listeners, we obtain the outlines of how the consultants might have judged casualness versus formality: hearing [l] or [ɫ], they tacitly deduced the rankings that were needed to obtain it, figured out the relative positions of the selection points that yielded these rankings, then obtained the casualness/formality ratings based on the selection points.

[19] It seems likely that government of variation by speech style is only possible where both variants are reasonably well formed; thus consultants did not have strong opinions about casualness in words like *light* or *help* where there was really only one option available in the first place.

The upshot is that the ranges of constraint strictness appear to possess some internal structure, with different ends annotated (following general principles of speaker-oriented and listener-oriented speech) for formal versus casual register. [20]

4. The Origin of Gradient Well-Formedness[21]

An issue only briefly mentioned thus far is where gradient well-formedness intuitions come from: specifically, the problem of language acquisition for gradient well-formedness.

There are two primary sources of linguistic knowledge on which the language-acquiring child can draw: the *a priori* principles of Universal Grammar, and the ambient data. If there are cases where UG principles firmly dictate an outcome, the judgment the learner acquires by adulthood will surely be quite firm. Likewise, where the relevant phenomenon is not at all innately driven, but the ambient data are very robust, there is also good reason to think that the judgments will be firm. For instance, the morphological principle specifying *-ing* suffixation in the formation of English present participles is utterly exceptionless in the input data for English. Accordingly, it seems inconceivable that the present participle of a hypothetical verb like *blick* could be anything other than *blicking*.

The place to look for the origin of gradient judgments, I believe, is in the remaining set of cases, where (as far as we can tell) no principle of UG forces the outcome, and the data for acquisition are sparse. Here are some examples:

- The present example: light and dark /l/ clearly participate in massive free variation, with the data further obscured by the variable phonetic realization of both categories (Sproat and Fujimura 1993). It seems very likely that the input data for most English learners do not suffice to establish firm rankings of /l/ IS DARK and PREVOCALIC /l/ IS LIGHT with respect to the Paradigm Uniformity constraints.

- In a pilot study, I have found it possible at least tentatively to model gradient well-formedness judgments for a corpus of novel deverbal Latinate *-able* adjectives (like ?*obfuscable*, from *obfuscate*) currently under study by Donca Steriade. The judgments here are blatantly gradient, a fact that plausibly reflects the extreme sparseness of input data for what is a quite learned and unusual process of word formation.

- More generally, examination of the literature on theoretical syntax reveals a research strategy that is constantly leading to cases with gradient judgments: syntacticians consistently try to push the study of a given syntactic phenomenon into all contexts where it occurs or could logically be imagined to occur. Thus, for example, Wh- extraction for a given language is tested for all syntactic environments in which a noun phrase is ever permitted. This is good scientific strategy, since it gives us data of the

[20] There is one further data puzzle here: why didn't the *mail-er* class show significant formality effects? I suspect that there may have been a conflict present with regard to what should be counted as "formal": a light [l] renders the /l/ phoneme more identifiable, but a dark [ɫ] renders the morphological structure clearer; each factor is in principle helpful to the listener.

[21] This section presents several views that have been arrived at independently by Paul Boersma (1997).

fullest possible scope. However, the very act of pushing a phenomenon to its limits is often what leads to gradient well-formedness judgments. As before, I conjecture that the cases involved here are precisely those in which the combined effects of UG and input data do *not* suffice to establish a clear pattern of well-formedness. That is, the sentences that syntacticians ask are often so unusual that they correspond to areas in which the data given to the child did not suffice to create categorial well-formedness judgments.

If the view taken here is right, the way to gain insight into gradient well-formedness is to examine what language learning would look like in the difficult areas; namely, those with sparse data and no inviolable UG principles to help. Along these lines, let us consider one promising approach to learning within OT.

4.1 The Tesar/Smolensky Ranking Algorithm

Tesar and Smolensky (1993, 1996) have devised an algorithm, called Constraint Demotion, which given (a) a set of well-formed outputs of a grammar; (b) a set of ill-formed rival candidates for each input; and (c) a set of constraints, will locate a ranking for the constraints that derives all and only the well-formed outputs. The algorithm is proved always to succeed, provided that such a ranking exists. Further, it is very efficient, working in quadratic time.

I believe that this algorithm is a good first step toward an ultimate formal theory of language learning. However, one should also understand its limitations. Most saliently, the algorithm cannot learn any grammar that generates free variation. The reason is that it is founded on the crucial assumption that if [A] is the output deriving from the underlying form /A/, then any other form [A′] distinct from [A] must be ill-formed. The algorithm uses the putative ill-formedness of [A′] as negative evidence, forcing the demotion of constraints that would select [A′] as the output instead of [A].

My own judgment is that it is unlikely that we will ever find a language whose grammar does not generate free variation. This view is based on my experience in working with consultants, reading grammars, and studying the work of sociolinguists. If I am right, the inability of the Constraint Demotion algorithm to learn free variation must be considered a quite general failure, not an isolated difficulty.

What do we want a learner to do when it meets with free variation? Plausibly, it should do what OT analysts currently do, namely to posit a small stratum of crucially freely-ranked constraints. The size of the freely-ranked stratum should be as small as possible, so that it confines the variation to the cases that are justified by the data and not overgenerate.

4.2 Robust Algorithms and Their Consequences for Gradience

There is one other characteristic that an improved constraint-ranking algorithm ought to have: robustness in the face of noisy input data. Consider, for example, a speech error collected by Joseph Stemberger (1983, 32):

(38) *in a first floor [dl]orm—dorm room

Stemberger claims, based on his own error corpus, that cases that violate legal English segment sequencing occur with "a reasonable frequency." Thus, it is easy to imagine a form like (38) being uttered in the presence of a small child who is still in the process of learning the English segment-sequencing system. Taken at face value, the datum given is a clear invitation to the child to rerank the relevant Faithfulness constraints above whatever markedness constraint it is that forbids initial [dl]. But to my knowledge, all adult native speakers of English immediately recognize *dlorm* as impossible. Thus, the occasional hearing of a phonotactically-illegal sequence from speech errors during childhood does not prevent language learners from developing an appropriate sense of phonotactic well-formedness by the time they have become adults.

From this we can conclude the following: the child learning language must consider *all* the data that she hears on a provisional basis. Only a reasonably healthy dosage of [dl] forms should ever induce her to consider a constraint ranking that would permit [dl] as a legal word onset.

One way to implement such a scheme would be to have input data adjust the rankings of constraints *incrementally*: hearing a datum that would imply a novel ranking leads the ranker to make a modest adjustment, only slightly advancing the fringes and central domains of the relevant constraints.[22] Plausibly, under such a regime, the hearing just once or twice in childhood of [dl] would not suffice to adjust the relevant constraint rankings to the point where [dl] emerged as well formed in the adult grammar.

From this perspective, I think it is clear why gradient well-formedness intuitions might arise where input data are sparse or otherwise uninformative: an incremental learner would never get enough data to eliminate the fringes that it set up as tentative hypotheses.

A practical conclusion of this section is that computational linguists now have a major opportunity to contribute to progress in learning theory. The development of algorithms that do what the Tesar/Smolensky algorithm does, but with variable and sparse input data, would be a very significant advance, and justifies, I believe, a serious research effort.[23]

5. Conclusions: Advantages of Analyzing Gradiently

To conclude, I will discuss some possible virtues of the proposal made here.

5.1 Conservatism

[22] It seems likely, as Janet Pierrehumbert has pointed out to me, that fringes and central domains are adjusted *together*. In this view, fringes are a predictable adjunct to every constraint, rather than being learned separately.

[23] Since originally writing the above paragraph I have encountered the algorithm of Boersma (1997; this volume), which aspires to exactly the goals I have stated. We are currently submitting to the algorithm a number of empirical tests intended to assess its capacities as thoroughly as possible.

First, from the viewpoint of mainstream Optimality Theory, the proposal is rather conservative, because all previous forms of constraint ranking are compatible with it. The kinds of rankings posited in earlier Optimality-theoretic work may be seen as a subset (indeed, a very important subset) of the rankings countenanced here. This was shown above under (32) and (33): to the existing inventory of invariant rankings (nonoverlapping bands along the strictness continuum) and free ranking (overlapping strictness bands), the present proposal adds only the fringes. Constraints that overlap only in their fringes are assumed to have a strongly preferred, but not quite obligatory ranking. It is easy to imagine a large grammar in which quite a few of the constraints have nonoverlapping ranges, corresponding to empirical domains in which free variation is lacking and intuitions are robust.

The upshot is that nothing in the current OT literature is incompatible with what has been suggested here. For linguists who think that OT has led to considerably deepened understanding of linguistic phenomena (and I am among them), this is an important point.

The conservatism of the present proposal may be contrasted with the view that gradient judgments imply "gradient principles" or "gradient representations," a view taken for instance by Ross (1972).[24] My own tentative view is that such an approach is unlikely to be fruitful: a gradient representation must be gradient in all contexts in which it appears, which is likely to make it hard to analyze those particular cases in which well-formedness judgments come out completely clear. In contrast, the more flexible mechanism of gradient constraint ranking can easily generate gradient outcomes in some cases and clear outcomes in others, as was shown here.

5.2 Text Frequencies

As noted above in section 3.2.4, fine-tuning of the constraint ranking bands allows us to make predictions about the relative frequencies with which rival output types are used. Hayes and MacEachern (to appear) have been able to do this in some detail for a problem in metrics (quatrain types), in which the data matched for textual frequency was a corpus of 1000 Appalachian folk song quatrains; see also Boersma (1997; this volume) for further frequency modeling.[25]

5.3 Generality

A further advantage that could be asserted for the present approach is that it deals with gradient well-formedness in a completely general way, one intimately tied up with the structure of the theory itself; rather than involving *post hoc* additions to the theory in particular areas. The pervasiveness of

[24] "A number of phenomena ... suggest that the traditional distinction between verbs, adjectives, and nouns—a distinction which is commonly thought of as discrete—should be modified. I will postulate, instead of a fixed, discrete inventory of syntactic categories, a quasi-continuum..." (Ross 1972, 316)

[25] We are here preceded by Anttila (forthcoming), who models the frequency of certain unusual Finnish inflected forms in an amazingly simple way, simply by counting the number of grammars (under free ranking) that generate each output. At least for the case of English quatrain structure, it appears that it may be difficult to get Anttila's idea to generalize to other data. See also Boersma (this volume) for an algorithm that matches strictness ranges to corpus frequency.

gradient well-formedness judgments in language suggests that a fully general approach is likely to be the correct one.

6. Coda: Why This is Worth Trying

To assert a commonplace: there is little point in analyzing overidealized data (here: sorted into "✓" and "*") if you possess a theory that permits you to analyze accurate data. Supposing that some particular principles can be deployed with categorial ranking to analyze a particular data set, reified into "✓" and "*" categories (e.g. as in (30) above); and that the same principles can be more subtly deployed under the scheme laid out here to capture the full range of judgments. It would seem that in the latter case the principles at stake have passed a stricter empirical test.

This is a nontrivial result, because there is good evidence that at present linguistics is *not difficult enough*. Specifically, there are usually multiple theoretical approaches to a given problem that work reasonably well. The task of sorting them out is unlikely to succeed unless we submit all approaches to the most stringent empirical testing available.

To conclude, I will cite a rather disarming quotation from Pesetsky (1997, 151), whose view concerning one particular aspect of linguistic research seems quite pertinent here. I have harmlessly altered a couple of words to maximize relevance:

> "A useful tactic when considering novel ideas is to look—not for problems that have already received satisfying solutions—but for problems that are largely unsolved. This is not as easy as it sounds. One learns to live with one's unsolved problems, and with time, one becomes so used to the unsolved that it almost comes to look solved ... Attempts to deal with [the particular question Pesetsky is discussing] in the traditional program of [linguistic] theory have been fairly unsuccessful, and [linguists] have become used to the unsuccessful proposals that are popular."

Gradient well-formedness may be one of the most pervasive overlooked-but-unsolved problems in linguistics. The proposal made here is presented in hopes that it may help bring this problem back into the light and, perhaps, serve as the key to solving it.

References

Anttila, Arto (forthcoming) "Deriving variation from grammar," to appear in F. Hinskens, R. van Hout, and Leo Wetzels, eds., *Variation, Change and Phonological Theory*, John Benjamins, Amsterdam. [Also: Rutgers Optimality Archive 63, http://ruccs.rutgers.edu/roa.html]

Bard, Ellen Gurman, Dan Robertson, and Antonella Sorace (1996) "Magnitude estimation of linguistic acceptability," *Language* 72, 32-68.

Beckman, Jill N. (1995) "Shona height harmony: markedness and positional identity," in Jill N. Beckman, Laura Walsh Dickey, and Suzanne Urbanczyk, eds., *Papers in Optimality Theory*, University of Massachusetts Occasional Papers 18, 53-75.

Beckman, Jill N. (1997) *Positional Faithfulness*, Ph.D. dissertation, University of Massachusetts, Amherst. Rutgers Optimality Archive 234, http://ruccs.rutgers.edu/roa.html].

Benua, Laura (1995) "Identity effects in morphological truncation," in Jill N. Beckman, Laura Walsh Dickey, and Suzanne Urbanczyk, eds., *Papers in Optimality Theory*, University of Massachusetts Occasional Papers 18, 77-136. [Also: Rutgers Optimality Archive 74, http://ruccs.rutgers.edu/roa.html].

Boersma, Paul (1997) "How we learn variation, optionality, and probability," Rutgers Optimality Archive 221, http://ruccs.rutgers.edu/roa.html.

Boersma, Paul (this volume) "Learning a grammar in Functional Phonology"

Borowsky, Toni and Barbara Horvath (forthcoming) "L-vocalization in Australian English," to appear in F. Hinskens, R. van Hout, and Leo Wetzels, eds., *Variation, Change and Phonological Theory*, John Benjamins, Amsterdam.

Burzio, Luigi (1997) "Multiple correspondence," paper read at the BCN Workshop on Conflicting Constraints, Groningen.

Casali, Roderic F. (1996) *Resolving Hiatus*, Ph.D. dissertation, UCLA, Los Angeles, CA.

Chomsky, Noam and Morris Halle (1968) *The Sound Pattern of English*, Harper and Row, New York.

Cowart, Wayne (1997) *Experimental Syntax: Applying Objective Methods to Sentence Judgments*, SAGE Publications, Thousand Oaks, CA.

Davidson, Lisa and Rolf Noyer (1997) "Loan phonology in Huave: nativization and the ranking of Faithfulness constraints," *Proceedings of the Fifteenth West Coast Conference on Formal Linguistics*, ed. Brian Agbayani and Sze-Wing Tang, Center for the Study of Language and Information, Stanford, CA, pp. 65-79.

Gimson, A. C. (1970) *An Introduction to the Pronunciation of English* (2nd ed.), St. Martin's Press, New York.

Hayes, Bruce (1989) "The Prosodic Hierarchy in meter," in Paul Kiparsky and Gilbert Youmans, eds., *Rhythm and Meter*, Academic Press, Orlando, FL, pp. 201-260.

Hayes, Bruce, and Margaret MacEachern (to appear) "Folk verse form in English," To appear in *Language*; currently available at Web site listed below.

Kenstowicz, Michael (1996) "Base Identity and Uniform Exponence: alternatives to cyclicity," in Jacques Durand and Bernard Laks, eds., *Current Trends in Phonology: Models and Methods*, Vol. 1, European Studies Research Institute, University of Salford, pp. 363-393. [Also: Rutgers Optimality Archive 103, http://ruccs.rutgers.edu/roa.html].

Kenstowicz, Michael (1997) "Uniform Exponence: exemplification and extension," paper read at the 1997 Johns Hopkins University workshop on Optimality Theory.

Kiparsky, Paul (1985) "Some consequences of Lexical Phonology," *Phonology Yearbook* 2, 85-138.

Kirchner, Robert (to appear) *Lenition in Phonetically-Driven Optimality-Theoretic Phonology,* Ph.D. dissertation, UCLA, Los Angeles, CA.

Kohler, Klaus J. (1990) "Segmental reduction in connected speech in German: phonological facts and phonetic explanations," in William J. Hardcastle and Alain Marchal, eds., *Speech Production and Speech Modelling*, Kluwer, Dordrecht, pp. 69-92.

Lindblom, Björn (1990) "Explaining phonetic variation: a sketch of the H&H theory," in William J. Hardcastle and Alain Marchal, eds., *Speech Production and Speech Modelling*, Kluwer, Dordrecht, pp. 403-439.

McCarthy, John and Alan S. Prince (1995) "Faithfulness and reduplicative identity," in Jill N. Beckman, Laura Walsh Dickey, and Suzanne Urbanczyk, eds., *Papers in Optimality Theory, University of Massachusetts Occasional Papers* 18, 249-384. [Also: Rutgers Optimality Archive 60, http://ruccs.rutgers.edu/roa.html].

McCawley, James D. (1968) *The Phonological Component of a Grammar of Japanese,* Mouton, The Hague.

Nespor, Marina and Irene Vogel (1986) *Prosodic Phonology*, Foris, Dordrecht.

Pesetsky, David (1997) "OT and syntax: movement and pronunciation," in Diana Archangeli and D. Terence Langendoen, eds., *Optimality Theory: An Overview*, Blackwell, Oxford, pp. 134-170.

Prince, Alan S., and Paul Smolensky (1993) *Optimality Theory: Constraint Interaction in Generative Grammar*. Ms., Rutgers University and University of Colorado. To appear: Cambridge, MA: MIT Press.

Raymond and Hogan (1993) "The Optimality interpreter," Rutgers Optimality Archive 130, http://ruccs.rutgers.edu/roa.html.

Ross, John R. (1972) "The category squish: Endstation Hauptwort," *Chicago Linguistic Society* 8, 316-328.

Schütze, Carson T. (1996) *The Empirical Base of Linguistics: Grammaticality Judgments and Linguistic Methodology*, University of Chicago Press, Chicago.

Selkirk, Elisabeth O. (1980) "Prosodic domains in phonology: Sanskrit revisited," in Mark Aronoff and Mary-Louise Kean, eds., *Juncture*, Anma Libri, Saratoga, CA.

Simpson, Jane (1980) "Cyclic syllabification and a first cycle rule of vowel-rounding in some dialects of Australian English," ms., Dept. of Linguistics and Philosophy, MIT.

Sproat, Richard and Osamu Fujimura (1993) "Allophonic variation in English /l/ and its implications for phonetic implementation," *Journal of Phonetics* 21, 291-311.

Stemberger, Joseph P. (1983) "Speech errors and theoretical phonology: a review," paper distributed by the Indiana University Linguistics Club, Bloomington, IN.

Steriade, Donca (1993) "Positional neutralization," talk presented at NELS 24, University of Massachusetts, Amherst.

Steriade, Donca (1996) “Paradigm Uniformity and the phonetics-phonology boundary,” paper presented at the Fifth Conference in Laboratory Phonology, Northwestern University, Evanston, IL.
Steriade, Donca (1997) “Phonetics in phonology: the case of laryngeal neutralization,” ms., Dept. of Linguistics, UCLA, Los Angeles, CA.
Tesar, Bruce B. and Paul Smolensky (1993) “The learnability of Optimality Theory: an algorithm and some basic complexity results,” Rutgers Optimality Archive 2, http://ruccs.rutgers.edu/roa.html.
Tesar, Bruce B. and Paul Smolensky (1996) “Learnability in Optimality Theory,” Rutgers Optimality Archive 110, http://ruccs.rutgers.edu/roa.html.
Wells, John C. (1982) *Accents of English* (in three volumes), Cambridge University Press, Cambridge.

Department of Linguistics
UCLA
Los Angeles, CA 90095-1543

bhayes@humnet.ucla.edu

http://www.humnet.ucla.edu/humnet/linguistics/people/hayes/hayes.htm