Predicting the unpredictable: Capturing the apparent semi-regularity of rendaku voicing in Japanese through harmonic grammar*

Eric Rosen
University of California, Santa Cruz & University of British Columbia

1 Introduction

Semi-regular phonological processes occur often in natural language. For example, rendaku voicing in Japanese fails to occur in a seemingly random fashion among roughly 16% of certain classes of compounds. This presents an analytical challenge for generative theories with exceptionless rules or categorical constraints: irregularity of any kind must arise within lexical representations, not the grammar. For example, the compound in (1) predictably voices by the rule of rendaku voicing McCawley (1968, inter alia) but the compound in (2) doesn't. There is no known phonological distinction between *kuma* and *yama* that enables a rule or constraint to explain the difference.

(1)	kuma	+	te	→	kuma-de
	'bear'		'hand'		'rake'

(2)	yama	+	te	→	yama-te
	'mountain'		'hand'		'mountainside'

Kawahara (2015) confronts the question of whether semi-regular processes such as rendaku voicing in Japanese should be considered phonological or purely lexical, in view of its many exceptions. He addresses an apparently lexicalist view of rendaku taken by Vance (2014), and presents a number of arguments in favor of its status as a bona-fide phonological process. In describing what is at stake with respect to the question of whether rendaku is lexical or phonological, he points out that 'rendaku and its properties have been extensively used for phonological argumentation, and that rendaku has been made famous among the field of phonological theory in general' (Kawahara, 2015, p. 4). As he discusses, the irregularity of rendaku undermines its phonological status inasmuch as its production in a generative framework will be short-circuited by lexical listings that apply to a whole compound in cases where the process fails to apply. Given the fact that the same word will voice in some compounds but not in others, it is not enough to list something about the voicing of a word in its independent lexical entry – at least it so appears. In familiar types of generative frameworks, it would seem necessary for the lexicon to separately list any compound, such as *yama-te* above, that happens to block the process of voicing.

*I wish to thank Paul Smolensky, Matt Goldrick, Armin Mester, Nick Kalivoda, Doug Pulleyblank, Junko Itô, Grant McGuire, Pranav Anand, Deniz Rudin, other members of the UCSC Phlunch group, the audience at the Berkeley Linguistic Society 2016 conference and Yohei Rosen for helpful comments and suggestions. I also thank Mark Schmidt for valuable mentorship in machine learning. All errors are my own. Please send any inquiries to errosen@mail.ubc.ca.

Not only does rendaku exhibit irregularity McCawley (1968, inter alia), but, as is well known, the preference or dispreference for a given word to undergo voicing in a compound shows gradience across the lexicon, as will be shown on page 3 in (3) and (4) below, where we see a continuum that goes from complete resistance to rendaku in some words all the way across the scale to a state of exceptionless voicing by other words (Haruo Kubozono, p.c., 2000; Irwin, 2015, inter alia). There is no way in standard rule-based (e.g. Chomsky and Halle, 1968) or constraint-based frameworks (e.g. Prince and Smolensky, 1993) in which features are binary or privative rather than scalar, to give a word a feature that will determine its precise degree of preference for voicing.[1]

This work proposes a new analysis of rendaku that solves this problem, allowing the correct output forms to be generated with little or no specification of voicing for particular, 'exceptional' compounds. I adopt the framework of Gradient Symbolic Computation (Smolensky and Goldrick, 2015, henceforth GSC), a type of Harmonic Grammar (Legendre et al., 1990a) that allows weighted constraints and features with continuous activation levels. In this analysis, rendaku voicing occurs by the coalescence of two stem-specific, partially activated [+*voi*] features that occur as attributive affixes on compound-forming stems: a variation of the junctural morpheme for rendaku proposed by Itô and Mester (1998). Only when the additive combination of these features exceeds some threshold t does voicing occur. In the above examples, $[+voi]_{kuti} + [+voi]_{kuse} > t > [+voi]_{sake} + [+voi]_{kuse}$. The contribution of *both* conjuncts to voicing captures not only the well-known gradient continuum of voicing preference/dispreference among second conjuncts but also a lesser-known gradient effect of first conjuncts on voicing.

Adopting the principle of Minimum Description Length (Goldsmith, 2009) I will show that GSC can provide a better model of this semi-regular phenomenon than other frameworks by reducing the degree of lexicalization with minimum cost to the grammar. Moreover, computer-simulated algorithms show that this proposed grammar is learnable. This analysis holds promise that the GSC model can shed new light on the lexicalization versus grammaticalization question with respect to other semi-regular processes.

This analysis also accounts for the gradient nature of rendaku, where the preference or dispreference for voicing of a given word can be shown to be due to the activation strength of its underlying voicing feature. An intrinsic part of that explanation will be the proposal that not only do words that occur as the second conjunct of a compound exhibit gradient preferences for voicing, but the first conjunct in a two-member compound also arguably exhibits the same kind of gradient preference for triggering voicing in the word that follows it. This phenomenon is not as easy to see unless it is viewed as part of the overall interaction between the voicing activation values of the two conjuncts. Evidence that supports the hypothesis of gradience for triggering as well as for voicing will be demonstrated by creating a noncontradictory hierarchy of voicing feature activation values – a hierarchy that will be shown to be statistically very unlikely if the first conjuncts of the set of compounds were to just occur randomly.

2 Examples of the gradient nature of rendaku voicing

A database of 921 noun-noun compounds was used in this analysis. The first conjuncts (henceforth 'N1's') occurred among a set of 233 Yamato (native Japanese) nouns and the second conjuncts (henceforth 'N2's') among a set of 306 Yamato nouns. The database was limited to compounds of native Yamato origin in which the total moraic count does not exceed four (i.e. compounds of the

[1] In Rosen (2001) I proposed a partial solution to some of these problems by having the lexical listing of a compound occur representationally as pointers to the listings of the individual constituents, but it addressed the issue of gradience only in a very coarse-grained way.

form $1\mu + 1\mu$, $1\mu + 2\mu$, $1\mu + 3\mu$, $2\mu + 1\mu$, $2\mu + 2\mu$, or $3\mu + 1\mu$).[2] Compounds in which rendaku voicing is blocked because of Lyman's Law or because of belonging to the class of coordinative or dvandva compounds, which also block rendaku (Martin, 1987, p. 9), were also excluded. Common placenames and family names, which show evidence of being lexicalized as single words, were also omitted. If we examine the frequency at which a noun experiences rendaku voicing within this set, we find the following examples of a range of voicing frequencies among nouns that occur in at least six compounds:

(3)

N2	saki	kusa	te	hara	ki	kawa
Gloss	tip	grass	hand	field	tree	skin
Freq. of rendaku	0	0.12	0.21	0.25	0.33	0.5
Num. of examples	16	17	19	8	18	14

(4)

N2	tori	hune	hito	hue
Gloss	bird	boat	person	flute
Freq. of rendaku	0.84	0.93	1.0	1.0
Num. of examples	13	21	16	12

(5) shows the triggering frequency of some N1s as the first conjuncts of two-member compounds. Fewer nouns occur as abundantly in this position, since the number of nouns that can occur as the nonhead of the compound will be less semantically and pragmatically limited than for the head noun.

(5)

N1	niwa	me	mizu	ura	yama	yoko	hana2
Gloss	garden	eye	water	back	mountain	side	flower
Freq. of rendaku	0.16	0.5	0.7	0.75	0.81	0.9	1.0
Num. of examples	6	12	17	12	22	10	10

The following graph shows the voicing behavior of the compounds in the dataset, with N1, the first conjunct, represented as a distance along the y-axis and N2's on the x-axis, arranged in order of voicing tendency on both axes. Red dots represent a compound that voices, blue dots one that doesn't. The graph is subdivided into three rows and three columns in each dimension according to the overall voicing behavior of the stems. The graph shows the patterning of gradient voicing preferences among stems. Overall, 16% of compounds in this set fail to voice. In the lowest row and leftmost column are N1's and N2's that always block voicing; in the highest row and rightmost column, N1's and N2's that always participate in voicing. But in the middle rectangle we find both blue dots and red dots. The fact that clustering of colors occurs on this graph is a clue towards the analysis of rendaku that shall be presented.

The next section introduces the Gradient Symbolic Computation framework under which this analysis will be developed.

[2]The database was limited to compounds that fall within this length limit because of evidence that compounds of greater length are much less resistant to rendaku voicing – what Rosen (2001, 2003) calls the 'prosodic size factor' (see also Kawahara and Sano, 2014). There is a similar reason for limiting the database to Yamato morphemes only. If we adopt the stratification of the Japanese lexicon proposed by Itô and Mester (1995), compounds show increasing resistance to rendaku as we move outside of the Yamato stratum of the lexicon to other strata.

(6)

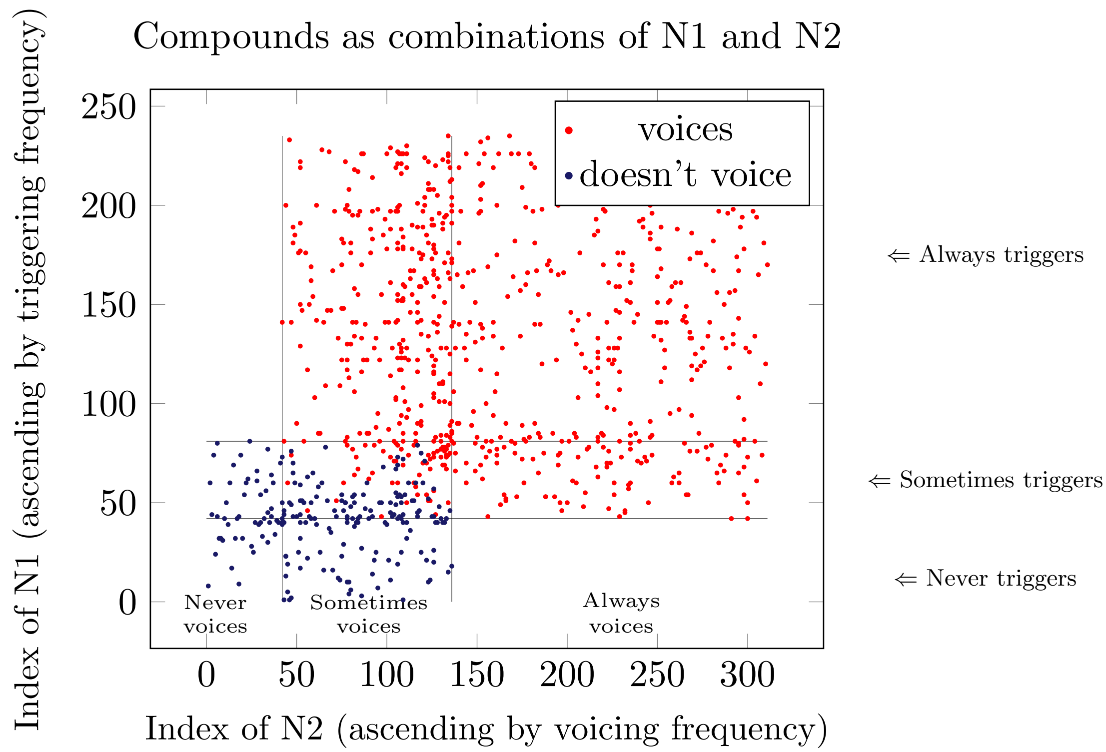

3 The Gradient Symbolic Computation framework (Smolensky and Goldrick, 2015)

This grammar architecture consists of two levels: (a) a symbolic level of discrete symbol structures in which symbols such as representations of phonemes are assigned to a set of roles such as positions in a string; (b) a subsymbolic or subconceptual level which is a kind of neural connectionist network that consists of distributed representations in which a given binding between a filler such as a phoneme and a role such as a position is distributed over the whole network. Specifically, there is an activity pattern for a given filler-role binding which is calculated by the tensor product of activity vectors that encode the filler and the role at this subsymbolic level. The way the symbolic level is derived algorithmically from the subsymbolic level involves two important factors: (i) a gradually decreasing 'temperature' factor T of added Gaussian noise that creates a simulated annealing process for optimization and (ii) a gradually increasing quantization factor q that forces the output to be discrete (at least to some degree, although blended representations are still possible).

At the symbolic level, the grammar belongs to the class of Harmonic Grammars (e.g. Smolensky, 1986; Legendre et al., 1990a,b; Goldsmith, 1993; Pater, 2009), where, as Smolensky and Goldrick describe, 'the grammatical wellformedness of a symbol structure S is measured by a grammar-Harmony function $H_{\mathcal{G}}$. $H_{\mathcal{G}}(S)$ is the weighted sum of S's violations of constraints on co-occurrence of filler/role bindings'. In this symbol structure, structural positions are occupied by blends of symbols that can have partial levels of activation.

Of particular interest to us here is the way, in this kind of harmonic grammar, the interaction of partially-weighted constraints with partially activated input features will derive particular outputs. The analysis of rendaku to be presented here parallels, in certain ways, an analysis of French liaison in the GSC framework by Smolensky (2015), to which the reader is referred.

4 A GSC account of rendaku voicing

The apparent semi-regularity of Japanese rendaku becomes grammatical and explainable if we adopt the hypothesis that there are partially activated [+voice] features at the edges of morphemes whose activation values reflect, for the first conjunct, that morpheme's inclination to trigger rendaku in the following stem, and, for the second conjunct, that morpheme's inclination to undergo rendaku. The combined effects of the two coalescing, partially activated features will determine whether a [+voice] feature surfaces.

This requires a slight modification of the proposal of Itô and Mester (1998) of rendaku as a junctural prefix:

(7) [yoko] + $[\text{voi}]_\rho$ + [tsuna] → yokodzuna 'horizontal rope' (Itô and Mester, 1998, p. 29)

In the present account the rendaku morpheme is both a prefix and a suffix, whose features coalesce to one feature in the output as shown in (8). In (9), coalescence occurs but fails to produce voicing, which shall be explained forthwith.

(8)

kuma-$[\textit{voi}]_{\rho_1}$	+	$[\textit{voi}]_{\rho_2}$-te	→	kuma- $\overset{\rho_{1,2}}{\overset{\vert}{t}}$e = kuma-**d**e
'bear'		'hand'		'rake'

(9)

yama-$[\textit{voi}]_{\rho_1}$	+	$[\textit{voi}]_{\rho_2}$-te	→	yama- $\overset{\rho_{1,2}}{\overset{\vert}{t}}$e = yama-**t**e
'mountain'		'hand'		'mountainside'

If a particular N1 triggers voicing in a significant majority of compounds, it will have a strongly activated voicing feature that reflects its tendency (modulo the effects of the activation on N2) to trigger voicing in a compound. The same is true for N2's, where a strongly activated N2 reflects its type-frequency of voicing in various compounds.

At the symbolic level, the GSC framework has weighted constraints that act on features that can be partially activated. We shall see that the effects of weighted constraints such as Max and Dep, whose general properties are familiar from Optimality Theory (Prince and Smolensky, 1993), are such that there is a threshold for the sum of the activations of the [+*voi*] features on N1 and N2, above which voicing will occur and below which it will not.

The following harmonic tableaux show how voicing is determined by the effects of weighted Max and Dep constraints on partially-activated [+*voi*] features, to which, for the time being, we assign hypothetical values. We give Max a weight of 1 and Dep a weight of −1. A Max constraint creates positive harmony: its weight times the activation of the feature in question that surfaces. A Dep constraint creates negative harmony: its (negative) weight times the difference in activation values between a feature's underlying form and surface form. The winning candidate is the one with the highest harmony value.

(10)

	1	1	-1	
$kuma([+voi]\cdot(0.4))\ ((0.225)\cdot[+voi])te$	Max[+voi]$_1$	Max[+voi]$_2$	Dep[+voi]	H
☞ **kuma-de**	0.4	0.225	-0.375	**0.25**
kuma-te				0

- Dep violation is $1-(0.4+0.225)=-0.375$.

(11)

	1	1	-1	
$yama([+voi] \cdot (0.225))$ $((0.225) \cdot [+voi])te$	Max[+voi]$_1$	Max[+voi]$_2$	Dep[+voi]	H
yama-de	0.225	0.225	-0.55	-0.1
☞ **yama-te**				**0**

- Dep violation is $1 - (0.225 + 0.225) = -0.55$.
- The rendaku suffix (0.225) on *yama* 'mountain' is posited to have a lower activation value than the one on *kuma* 'bear' (0.4).
- So the combined weights are not enough to surpass the threshold and cause voicing.

In general, we can show that the threshold for voicing will be given by $\frac{D}{M+D}$, where M is the weight of Max and $-D$ is the weight of Dep.

Having rendaku voicing determined by whether the sum of voicing feature activation levels on N1 and N2 surpasses some threshold will depend completely on whether there can be a strict, noncontradictory domination hierarchy of voicing activation values on morphemes that reflects their triggering and voicing tendencies. Consider, for example, the following interleaved set of morphologically minimal pairs of compounds, shown in the diagram below, where the first two are repeated from (8) and (9) above. Under our hypothesis that voicing is determined by the combined strength of [+*voi*] rendaku affixes on N1 and N2, these examples establish a hierarchy of activation values for these affixes shown in the boxes below.

(12)

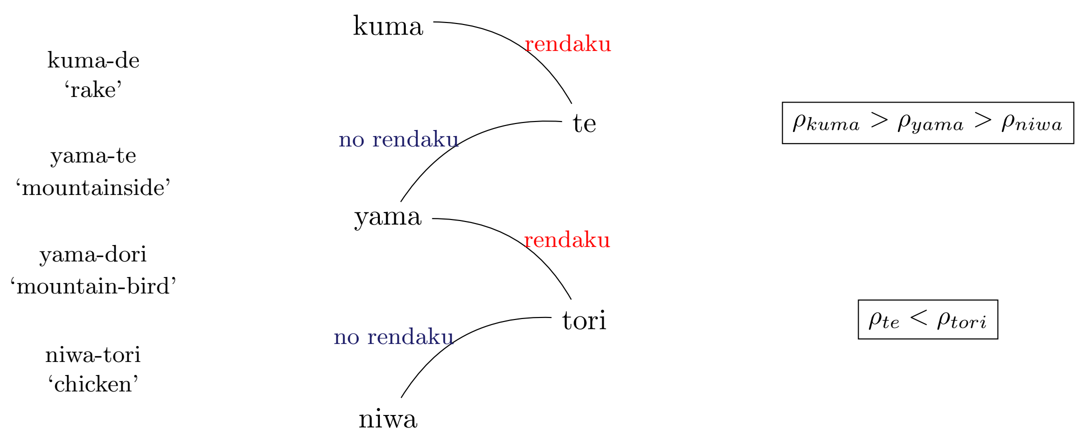

From this data we establish, for example, that $\rho_{kuma} > \rho_{yama} > \rho_{niwa}$ for stem rendaku suffixes.

So if *kuma* 'bear' and *niwa* 'garden' both combined with the same stem (e.g. *kinu* 'silk') to form compounds, the following morphologically minimal pair should be impossible:

(13) $\left.\begin{array}{l}\text{*kuma-kinu}\\ \text{*niwa-ginu}\end{array}\right\}$ Contradicts the hierarchy $\boxed{\rho_{kuma} > \rho_{niwa}}$

These are not real compounds, but if they existed, we predict that they could not have this kind of voicing contrast. If such contradictions are found to exist in the data, then our hypothesis – that voicing is determined by the sum of innate voicing activation levels of affixes on stems – will not be viable. But if no such contradictions can be found, and if such contradictions *are* likely to occur in a dataset of randomly voiced compounds, then the current hypothesis is supported. The following algorithm was used to search for contradictions in the data among both N1 rendaku suffixes and N2 rendaku prefixes. (Subscripts on N's here are to distinguish between different stems, not to distinguish between the first and second conjunct of a compound.)

1. Find all immediate domination instances from morphologically minimal pairs:
   $\rho_{N_1} > \rho_{N_2}$ iff N_1N_3 voices and N_2N_3 doesn't voice.
   e.g.: kuma-de vs. yama-te (previous page)
2. By transitivity of domination, ($\rho_{N_1} > \rho_{N_2}$ and $\rho_{N_2} > \rho_{N_3} \rightarrow \rho_{N_1} > \rho_{N_3}$), create a dominance tree.
3. Search the tree, depth first, for contradictions.

The following is a fragment of what such a domination tree would look like:

(14)

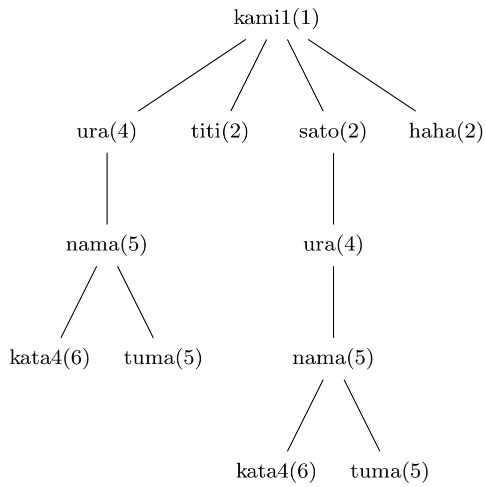

It is significant that no contradictions turned up in the data set. We might ask how likely it is that a lack of contradictions would occur randomly. To try to answer this question, the following computer simulation was carried out:

1. Go through the list of compounds, randomly altering the voicing on N2, so that voicing occurs 84% of the time.
2. Do this 10 times.
3. Check each list for contradictions in domination for ρ_{N_1} and ρ_{N_2}.

Each run of the simulation produced 5 to 11 contradictions, suggesting that the lack of contradiction in the real data reflects some real property.

Having established that there exists a set of non-contradictory activation levels that is consistent with the data, we next need to consider two related questions: what is an actual set of levels that can correctly derive the voicing of compounds in the data, and how can this set of activations be learnable?

A number of learning algorithms were tested by computer simulation, each with certain advantages and disadvantages. It would exceed the limitations on length of this paper to describe all of them, so I shall focus on one that has certain advantages with respect to Minimum Description Length, which shall be discussed in the following section, after which, a learning algorithm for activation levels will be presented.

5 Minimum Description Length (Goldsmith, 2011)

Minimum Description Length is an evaluation metric by which we shall argue that deriving the correct voicing outcome in compounds through the grammatical combination of partial voicing features is preferable to a model that regards the surface forms as resulting from a choice of a lexically listed allomorph. Our objective is to avoid unnecessary lexicalization of the phonology of compound words in cases where the grammar can determine correct outputs from underlying features. We want to find an optimal middle ground between two extremes: (a) an unwieldy lexicon that unnecessarily gives a full phonological description of every morphologically complex output form and (b) and unwieldy grammar that overfits the data. For the purpose of making a scientific judgment of how well a particular linguistic model fits the data that is being studied, Goldsmith (2002, 2011) and Goldsmith and Riggle (2012) develop the principle of Minimum Description Length or MDL, based on work by Rissanen (1989), which gives us a way of choosing the optimal trade-off between the lexical specification of output forms and a grammar that generates output forms from the sparsest possible set of input forms. Formally, MDL is calculated in bits as the sum of the two quantities we want to minimize: the negative log probability of the grammar plus the negative log probability of the data given the grammar: $argmin(-log_2 p(g) - log_2 p(D|g))$, where $p(D, g) = p(g)p(D|g)$ and thus $-log_2 p(D, g) = -log_2 p(g) - log_2 p(D|g)$. With respect to the present compound data we can calculate the description length of the data using the information-theoretic principle that data D can be encoded with n bits when the n is the positive log probability that the grammar assigns to D (Goldsmith and Riggle, 2012).

In calculating the MDL of a grammar that derives morphologically complex output forms from a composition of simple lexemes, we need to consider the question of which combinations of morphemes can combine: for example, which stems can combine with which affixes in the case of affixation or which stems can combine with which other stems in the case of compounding. To represent these kinds of combinatorial possibilities, Goldsmith proposes *signatures*: 'structures indicating which stems may appear with which affixes'.

In the case of the compounds we are considering, suppose that whether or not rendaku voicing occurred is determined through the pairing of allomorphs of stems in a signature structure, where allomorphs can be voiced or nonvoiced. A nonexhaustive signature for stems that can follow *mugi* 'barley' in a compound might look something like the following, where a voiced allomorph of the second conjunct occurs in the first and third entries.

(15)
$$\{\text{mugi `barley'}\} \left\{ \begin{array}{c} \text{bue `flute'} \\ \text{ko `flour'} \\ \text{batake `field'} \\ \text{mesi `food'} \end{array} \right\}$$

In the case of the proposed analysis of rendaku voicing in the GSC framework, a signature can also list, for a given stem, the level of activation of a rendaku affix it combines with. If every stem has an affix with a slightly different activation level than any of the others, then no affix could share its signature with more than one stem. On the other hand, if the voicing of compounds could be derived with a small number of discrete levels of activation on affixes, where each stem has a signature pointer a certain level, the burden on the lexicon will be less since the listing of signatures can be made more compact. Suppose, for example that a number of stems that can occur as N1's all associate with a rendaku affix with the same activation level, which we can for the time being call 'strong' then a signature for that set of stems might look like the following:

(16)
$$\left\{ \begin{array}{c} \text{kuma `bear'} \\ \text{kusa `grass'} \\ \text{umi `sea'} \\ \vdots \end{array} \right\} \{\rho_{\text{strong}}\}$$

With respect to MDL, then, it becomes important to find a learning algorithm that can find a small number of discrete activation levels for rendaku affixes rather than a different level for the affix that occurs with each stem. The next section introduces an algorithm that can achieve this.

6 A learning algorithm for activation levels of affixes

The following algorithm was tested through computer simulation.

1. Initialize the activation level of each N1 and N2 affix to zero, based on the hypothesis that stems will be encountered first in their simplex form, where there is no evidence for any [+*voi*] activation.
2. Initialize the weights of Max and Dep constraints at 1 and −1 respectively.
3. Initialize η, the stepsize for Max and Dep increments and for random noise perturbations at 0.01.
4. Initialize t, a temperature factor for simulated annealing on Max and Dep weights at 1.
5. For each iteration:
    (a) For each compound in the database:
        i. If voicing occurs and $a_{N1} + a_{N2} +$ Gaussian noise $< \frac{D}{M+D}$ (threshold):
            A. Increment Dep by stepsize $\times$ $a_{N1} + a_{N2} - 1 +$ Gaussian noise $\times$ temperature
            B. Increment Max by stepsize $\times$ $a_{N1} + a_{N2} +$ Gaussian noise $\times$ temperature
            C. Randomly choose N_i (one of N1 or N2) to increment first by some set amount (e.g. 0.05)
            D. If $a_{N_i} <$ threshold (i.e. won't voice in simplex word), increment a_{N_i}
            E. If still $a_{N1} + a_{N2} +$ Gaussian noise $< \frac{D}{M+D}$ (threshold), increment other affix's activation
        ii. If voicing does not occur and $a_{N1} + a_{N2} +$ Gaussian noise $> \frac{D}{M+D}$ (threshold):
            A. Follow the above 4 steps but in opposite direction.
        iii. If either N1 or N2 activation on its own with added Gaussian noise is above the threshold, decrement its activation by 0.05.
    (b) Drop the temperature factor to one-fourth of its value
    (c) Drop the stepsize for Max and Dep to 99.9% of its value
    (d) Stop if all compound activations gave a correct voicing value.

The results of a computer simulation of this algorithm were convergence after 482 iterations with the following final results.[3]

(17)

Max	1.07
Dep	0.85
threshold	0.44
Correct predictions	885
Incorrect predictions	0
N1 activation levels	8
N2 activation levels	7

[3]Other algorithms are possible. For example, an informal algorithm that factors in pitch accent can be shown to reduce the number of activation levels to three for each of N1 affixes and N2 affixes; however, this algorithm requires the set of compounds to be looked at somewhat globally rather than learning strictly from encounters of one compound at a time. Although this algorithm may be associated with model that has a slightly better MDL than a seven-level model, length limitations on this paper preclude a full discussion of such a model.

We shall now compare the model with activation levels derived from the algorithm described above with models that lexically list exceptional compounds.

7 Feature activation values and minimum description length

We shall consider here five possible models that account for the rendaku data at hand and compare them by each MDL:

Full lexical specification of compounds with each compound being listed separately in the lexicon.

Two separate allomorphs (with voicing and without) for all the N2's that show alternations in voicing and a single allomorph for the N2's that do not alternate. Compounds are derived through *signatures* (see above) that structure the combinations of co-occurring allomorphs

GSC account: 7 activation level classes The grammar generates output forms for compounds in the GSC framework as described above in Section 4 with partially activated features for voicing that can have seven different possible values.

Signatures with single allomorphs plus lexical listing of exceptions For all the compounds to which rendaku voicing applies, have signatures with single, voiced allomorphs for each of the N2's involved and lexically list the remaining nonundergoing compounds with no signatures.

For each of these models we need to calculate the following quantities whose sum will be the MDL for that model with respect to rendaku voicing of the dataset. Notice that we are abstracting away from a model of the whole language, but comparing how various models deal with the semi-regular process of rendaku voicing in a certain set of noun-noun compounds.

Cost of URs Unigram bit cost, phoneme by phoneme, of listing the UR or each simplex N1 and N2. Since these simplex forms can all exist on their own as monomorphemic words, in any model the allomorph that does not undergo rendaku needs to be listed anyway, so given that we are comparing models, we shall only be concerned here with the cost of listing voiced allomorphs.

Cost of fully-listed compounds Unigram bit cost, phoneme by phoneme, of listing each full compound on its own.

Cost of pointers to signatures If signatures are part of the model, the bit cost of all the pointers from signatures to the allomorphs involved.

Cost of partially activated feature representation If partially activated voicing features on N1's and N2's are part of the model, the bit cost of representing each feature. This cost will depend on the number of activation level classes.

Cost of grammar description The description cost of the relevant part of the grammar that derives output forms.

(18)

Summary of bit cost for various models						
Model	pointers to signatures	voicing activation values	grammar (network model)	extra allomorphs	listed compounds	total
Full specification	0	0	0	0	22,506	22,506.0
Full allomorphs	5,123.0	0	0	2,903.3	0	8,026.3
7 levels	2,963.9	1,248.0	363.1	0	0	4,575.0
Voicing through signatures & listed exceptions	2,352.3	2,903.3	0	0	6.781.9	11,737.7

In the foregoing calculation of MDL for various models to account for patterns of rendaku voicing, the GSC model with seven levels of activation came out with a substantially better MDL score than the other models, all of which account for the semi-regularity of rendaku through some kind of lexicalization for explaining exceptions.[4]

The following tableaux give derivations of two compounds, that were given above with hypothetical activations in (10) and (11), now based on activation levels that were derived from the algorithm given above in Section 6.

(19)

kuma([+*voi*] · (0.35)) ((0.15) · [+*voi*])*te*	1.07 Max[+voi]$_1$	1.07 Max[+voi]$_2$	−0.85 Dep[+voi]	H
☞ **kuma-de**	0.35	0.15	-0.425	**0.075**
kuma-te				0

- Dep violation is $(1 - (0.35 + 0.15)) \times 0.85 = -0.425$.

(20)

yama([+*voi*] · (0.25)) ((0.15) · [+*voi*])*te*	1.07 Max[+voi]$_1$	1.07 Max[+voi]$_2$	−0.85 Dep[+voi]	H
yama-de	0.25	0.15	-0.51	-0.11
☞ **yama-te**				**0**

- Dep violation is $(1 - (0.25 + 0.15)) \times 0.85 = -0.51$.

8 Partially-activated [+voice] features as attributive affixes

The underlying, partially-activated [+voice] feature that we are proposing here can be seen as an extension of the idea of morphological paradigms. Some researchers have argued (e.g. Itô and Mester, 1998) that rendaku occurs because of a junctural morpheme of the form [+voice] which effects some kind of compositional linking between the two morphemes, in the same way that the particle *no* does when it occurs after nouns in Japanese, allowing them to compose syntactically with certain following elements. The junctural morpheme hypothesis is supported by historical

[4] As pointed out by Paul Smolensky (p.c.) the idea of partitioning affixes on N1 and N2 into subclasses of different strengths could also be achieved in categorical OT through varying the rankings of Max-voi and *voi constraints to different subgrammars, each of which corresponds to a different activation level in a GSC account. The OT model appears to require an increased description length for a grammar that has a more complex constraint ranking, with separate faithfulness and markedness constraints with respect to voicing for each lexical subclass. A very rudimentary bitwise comparison of the MDL for an OT account versus a GSC account gives a bit cost of 264.7 for an OT model and 136.5 for a GSC model. Quite independent of the choice between the two models, what enables us to avoid over-lexicalization in accounting for rendaku's semi-regularity is the previously unexplored idea of combining the effects of both N1 and N2 on voicing, which can be done with either model.

evidence that this [+voice] feature is in fact the reflex of the particle *no* that occurred between two nouns in the same configuration at an earlier stage of the language (Martin, 1987). A slight modification of Itô and Mester's proposal is that the junctural morpheme is realized as both a suffix and a prefix on nouns rather than an independent morpheme or clitic in the same way that independent morphemes may evolve into suffixes or prefixes on words with which they often associate, forming morphological paradigms. Given the productivity of the way that Japanese lexical morphemes combine to form compounds, it would not be unreasonable to suppose that the surface realization of a noun with an added [+voice] is a kind of attributive affix. The alternation of the affixless form of a noun when it occurs in simplex form with the affixed form in compounds can be compared to the same kind of contrast between surface realizations of adjectives in German. Adjectives in German are inflected with either strong or weak declension suffixes when they occur attributively but are unsuffixed when they occur predicatively:

(21) ein klein-er Hund
a.M.NOM.SG small-M.NOM.SG dog
'a small dog'
(inflection on attributive adjective)

(22) Der Hund ist klein
the.M.NOM.SG dog be.3.SG small
'The dog is small'
(no inflection on predicative adjective)

So in Japanese we can think of an underlying form /asi-v_1 + v_2-kata/ 'foot+shape: footprint' as having partially-activated voicing features v_1, v_2 that realize a morphological paradigm whose allomorphs determine the attributive and nonattributive forms of the noun.

9 Previous accounts of rendaku's gradient, semi-regular nature

Vance (2014) places the burden of exceptions or gradience on the lexicon, not on the grammar. He cites an account in Ohno (2000) based on analogy, in which stems exhibit the same rendaku behavior if they are semantically or phonologically similar, for example, *siro* 'white' is proposed to behave similarly with respect to rendaku as *kuro* 'black', and *mimi* 'ear' to *hana* 'nose'. In Ohno's account, *kami* 'hair' a robust voicer (e.g. *kabe-gami* 'wall-paper') happens not to be voiced in *kuro-kami* 'black hair'. So he concludes that semantically similar *siro* 'white' should also block voicing in 2nd conjunct, which is what he finds in a psycholinguistic experiment in which novel compound *siro-kami* 'white hair' is chosen over a voiced version by subjects.

He also finds that *ti* 'blood' is a robust **non**-voicer except in *hana-zi* 'nosebleed'. So semantically similar *mimi* 'ear' should also block voicing in a 2nd conjunct. His evidence is novel compound *mimi-zi* 'ear-blood' which was chosen over a non-voiced version because of analogy with *hana-zi* 'nosebleed'.

Some problems with an analogical account are the following. 'Analogy' is a descriptive term to which it is difficult to assign the status of a rigorous principle of a formal grammar. And classifying stems by semantic category in a way that can predict voicing across the board seems impossible, given the following examples:

(23) • *tori* 'bird' doesn't voice after *niwa* 'garden' or *mizu* 'water'

- It does voice after *tutu* ‘pipe’, *yama* ‘mountain’, *hina* ‘chick’, *umi* ‘sea’ (note semantic relation to ‘water’ above), *huyu* ‘winter’, *oya* ‘parent’, *muku* ‘yew’, *natu* ‘summer’, *koma* ‘shogi piece’.

This leaves us with the question of how to semantically distinguish the two lists. GSC, on the other hand, captures Ohno’s examples in a more formal way through voicing activations that resulted from the learning simulation discussed in Section 6 above. For example, *siro*-[+*voi*] $\cdot$ 0.15 and *kuro*-[+*voi*] $\cdot$ 0.05 both had low activations in addition to the fact that they happen to have some semantic similarities.

10 Summary and conclusion

The dataset used in this study was purposely restricted in order to control the number of variables involved in what we are testing. By limiting the data to noun-noun compounds only that fall within a certain prosodic size limit and the same lexical stratum,[5] we can control for the possible effects on rendaku voicing of (a) syntactic category of the members of the compound (b) the effects of prosodic size and (c) lexical stratum, all of which have been shown to have an effect on whether or not rendaku voicing occurs (Kubozono, 2005, inter alia). Further study could look at ways in which these factors can be incorporated into an analysis within the same framework.

Notice also that even though the feature activation values we are positing are gradient, the predictions they make are categorical inasmuch as the output of N1+N2 for a given compound is not estimated to vary from speaker to speaker or utterance to utterance – at least within what we can take to be standard dialect.[6,7]

The GSC framework, by virtue of the way that gradience and blending of symbolic structures are an intrinsic part of its system, allows us to approach and understand gradient irregular patterns such as rendaku voicing in a way that is difficult or impossible with some other systems. The foregoing analysis of the irregularity of rendaku voicing arguably solves the problem of having its irregular patterns at cross-purposes with the productivity of generative rules or constraints. For the seven-level model, no lexical exception needs to be specified in a lexical listing for the whole compound. Another result of this approach is that it creates a more equal division of labor between the lexicon and the rest of the grammar in accounting for observed patterns in the language. Some of the gradient patterns we observe, such as the variable behavior among morphemes to undergo this process can only be accounted for through the way the lexical listing of individual morphemes acts in tandem with the generative process. This is because there are no robust correlations between phonological properties of the morphemes and their willingness to undergo the rendaku process, that could be accounted for by some phonological rule or constraint.

In this analysis, then, some of the patterning that we see in the language stems from grammatical processes and some also results from patterning in the lexicon. The lexicon in this view is therefore

[5]Paul Smolensky (p.c.) suggests that these proposed activation level classes can be seen as further subdividing the lexicon into substrata.

[6]Judgments that were made here with respect to whether or not voicing occurs for a given compound were made to conform with listed pronunciations in the NHK (Japanese Broadcasting Corporation) Pronunciation and Accent Dictionary (Nippon Hoosoo Kyookai (Japanese Broadcasting Corporation), 1998).

[7]Some may imagine that this kind of Harmonic Grammar is overly prone to making the kinds of ‘vague, soft, or non-categorical predictions’ that Goldsmith et al. (to appear, p. 4) observes that some linguists may incorrectly imagine will be the result of probabilistic models – models that share much in common with Harmonic grammars. The GSC model is not expressly a probabilistic model, but the Harmony function operates in many ways like the probability function of probabilistic and information theoretic grammars such as those proposed by Goldsmith (1993, 2002) and Goldsmith and Riggle (2012).

not completely random in the way that a framework such as OT (Prince and Smolensky, 1993) would seem to insist on.[8] At the same time, there is no evidence that the approach to the gradience of rendaku voicing described in this paper creates a 'duplication problem', since the lexicon and the rest of the grammar are arguably performing separate, nonoverlapping functions in determining voicing.

References

Chomsky, Noam, and Morris Halle. 1968. *The sound pattern of English*. Harper and Row.

Goldsmith, John. 1993. *The Last Phonological rule: Reflections on Constraints and Derivations*, chapter Harmonic Phonology, 21–60. Studies in Contemporary Linguistics.

Goldsmith, John. 2002. Probabilistic models of grammar: Phonology as information minimization. *Phonological Studies* 21–46.

Goldsmith, John. 2009. *Analogy in Grammar: Form and Acquisition*, chapter Morphological analogy: Only a beginning, 137–163. Oxford University Press.

Goldsmith, John. 2011. The evaluation metric in generative grammar. Paper presented at the 50th anniversary celebration for the MIT Department of Linguistics, December 2011.

Goldsmith, John, Alex Clark, Nick Chater, and Amy Perfors. to appear. *Empiricist Perspectives on Language Learnability*, volume 36 of *Recherches linguistiques à Vincennes*, chapter Towards a New Empiricism, 9–36.

Goldsmith, John, and Jason Riggle. 2012. Information theoretic approaches to phonology: the case of Finnish vowel harmony. *Natural Language and Linguistic Theory* 30:859–96.

Irwin, Mark. 2015. Rendaku Lovers, Rendaku Haters, and the Logistic Curve. In *Japanese/Korean Linguistics*, ed. Mikio Giriko, Naonori Nagaya, Akiko Takemura, and Timothy Vance. CSLI Publications.

Itô, Junko, and Armin Mester. 1995. The core-periphery structure of the lexicon and constraints on reranking. *University of Massachusetts Occasional Papers in Linguistics* 18:181–209.

Itô, Junko, and Armin Mester. 1998. Markedness and Word Structure: OCP Effects in Japanese. Technical Report 255, Rutgers Optimality Archive.

Kawahara, Shigeto. 2015. Can we use rendaku for phonological argumentation? *Linguistics Vanguard* .

Kawahara, Shigeto, and Shin-ichiro Sano. 2014. Testing Rosen's Rule and Strong Lymans Law. Technical Report 7, National Institute for Japanese Language and Linguistics.

Kubozono, Haruo. 2005. *Voicing in Japanese*, chapter Rendaku: Its domain and linguistic conditions, 15–20. Berlin: Mouton de Gruyter.

[8]For example, Smolensky et al. (2004, p. 1) write, regarding the strong output-oriented nature of OT: 'The strongest hypothesis is that all systematic, language-particular patterns are the result of output constraints – that there is no other locus from which such patterns can derive. In particular, the input is not such a locus'.

Legendre, Géraldine, Yoshiro Miyata, and Paul Smolensky. 1990a. Can connectionism contribute to syntax?: Harmonic grammar, with an application. In *Proceedings of the 26th regional meeting of the Chicago Linguistic Society*, ed. M. Ziolkowski, M. Noske, and K. Deaton, 237–252. The Chicago Linguistic Society.

Legendre, Géraldine, Yoshiro Miyata, and Paul Smolensky. 1990b. Harmonic Grammar: A formal multi-level connectionist theory of linguistic well-formedness: Theoretical foundations. In *Proceedings of the twelfth annual conference of the Cognitive Science Society*, ed. Lawrence Erlbaum, 388–395. Cambridge MA: Cognitive Science Society.

Martin, Samuel. 1987. *The Japanese Language Through Time*. Yale University Press.

McCawley, James. 1968. *The phonological component of a grammar of Japanese*. Mouton.

Nippon Hoosoo Kyookai (Japanese Broadcasting Corporation). 1998. *Nihongo Hatsuon Akusento Jiten (Japanese pronunciation and accent dictionary)*. NHK Bunka Kenkyuujo (NHK Cultural Research Institute).

Ohno, Kazutoshi. 2000. The lexical nature of Rendaku in Japanese. In *Japanese/Korean Linguistics*, ed. Mineharu Nakayama and Charles Quinn Jr., volume 9, 151–164. Stanford CA: CSLI Publications.

Pater, Joe. 2009. Weighted Constraints in Generative Linguistics. *Cognitive Science* 33:999–1035.

Prince, Alan, and Paul Smolensky. 1993. Optimality Theory: Constraint Interaction in Generative Grammar. Technical report, Rutgers University Center for Cognitive Science and Computer Science Department, University of Colorado at Boulder.

Rissanen, Jorma. 1989. *Stochastic complexity in statistical inquiry*, volume 15 of *World Scientific Series in Computer Science*. Singapore; Teaneck NJ: World Scientific.

Rosen, Eric. 2001. Phonological processes interacting with the lexicon: Variable and non-regular effects in Japanese phonology. Doctoral dissertation, Univeristy of British Columbia.

Rosen, Eric. 2003. Systematic irregularity in Japanese rendaku : How the grammar mediates patterned lexical exceptions. *Canadian Journal of Linguistics* 48:1–37.

Smolensky, Paul. 1986. *In parallel distributed processing*, chapter Information Processing in Dynamical Systems: Foundations of Harmony Theory, 194–281. Cambridge MA: MIT Press.

Smolensky, Paul. 2015. Grammar with Gradient Symbolic Structure. Sapir Lecture, Linguistic Society of America Summer Institute, University of Chicago.

Smolensky, Paul, Lisa Davidson, and Peter Jusczyk. 2004. The Initial and Final States: Theoretical Implications and Experimental Explorations of Richness of the Base. Technical Report 428, Rutgers University Center for Cognitive Science and Computer Science Department, University of Colorado at Boulder.

Smolensky, Paul, and Matthew Goldrick. 2015. Gradient Symbolic Computation. Class notes, LSA Summer Institute, Chicago.

Vance, Timothy. 2014. *Usage-based approaches to Japanese grammar: Towards the understanding of human language*, chapter If rendaku isn't a rule, what in the world is it?, 137–152. Amsterdam: John Benjamins.