Gradient symbolic representations in Harmonic Grammar

Brian Hsu
University of North Carolina at Chapel Hill

Author's note: This is the initial submitted version of this article. Comments are welcome!

Key words: Harmonic Grammar, gradient symbolic representations, optimality, weighted constraints, idiosyncrasy, exceptionality, blends

Abstract

This paper presents an overview of Gradient Harmonic Grammar, a weighted-constraint model of grammatical computation in which language structures are represented with numerically continuous levels of activity, or degree of presence. In this system, the penalty of each constraint violation is proportional to the activity of the structure that incurs it. The adoption of gradient activity permits unique advances in generative approaches to key types of idiosyncratic patterns in language. I review the main proposals that have been made in the framework, and outstanding issues such as the existence of gradient activity in output structures.

1. Introduction

Within generative linguistic theory, there are two key components to the success of Optimality Theory (Prince and Smolensky, 1993/2004) and its related constraint-based theories of grammatical computation. First, computation with violable constraints allows insightful explanations for many natural language patterns, particularly in relation to typology. Second, constraint-based frameworks offer promising explanations for how fundamental properties of natural language, characterized by discrete structures and well-formedness principles, can emerge from a connectionist cognitive architecture. At the same time, the closer investigation of language patterns within constraint-based frameworks can lead to a deeper understanding of how candidate optimality is computed. In the past two decades, numerous arguments have been advanced in favor of Harmonic Grammar (Legendre, Miyata, & Smolensky, 1990; Smolensky & Legendre, 2006), a model in which constraints are associated with numeric weights, rather than strict rankings (Farris-Trimble, 2008; Pater, 2009, 2016; Zuraw & Hayes, 2017; among many others).

This article is an overview of Gradient Harmonic Grammar (Smolensky & Goldrick, 2016; Smolensky, Goldrick, & Mathis, 2014), a weighted-constraint model in which language structures are represented with numerically continuous levels of *activity*, which emerges from activation levels of units in a neural network. The promise of Gradient Harmonic Grammar in linguistics lies in its power to generate patterns that have posed longstanding challenges in prior frameworks, primarily related to several types of idiosyncrasy. The use of gradient activity permits theoretically parsimonious approaches to these phenomena, while maintaining a sufficient level of typological restrictiveness. Although the framework has been most widely adopted in phonological analyses, it has seen analogous applications in constraint-based theories of syntax and semantics.

In Section 2, I illustrate the basic architecture of Gradient Harmonic Grammar, and the effects of gradient activity on constraint penalties. Section 3 introduces applications of gradient activity in the analysis of idiosyncratic patterns in language. Section 4 discusses the use of blends of partially active symbols in generating patterns of allomorph selection, and idiosyncratic processes conditioned by combinations of morphemes. Throughout these sections, I summarize the empirical and theoretical advantages of Gradient Harmonic Grammar analyses over prior approaches to the same phenomena. Section 5 turns to outstanding questions around whether gradient activity is possible in output representations. Section 6 concludes.

2. Constraint interaction in Gradient Harmonic Grammar

Informally, Gradient Harmonic Grammar is a version of Harmonic Grammar with an enriched representational system. As in Harmonic Grammar, grammatical computation involves the generation of output candidates from an input structure, and the selection of the candidate with the highest *harmony*, a numerical measure of well-formedness. Each constraint is associated with a numeric weight, and the harmony of a candidate is the sum total of each of its constraint violations multiplied by its weight. Gradient Harmonic Grammar further integrates two tenets of Parallel Distributed Processing connectionism (McClelland, Rumelhart, & The PDP Research Group, 1986): (i) Representational units have continuous degrees of activation that can change during computation. (ii) Mental representations can contain the simultaneous activation of multiple alternative structures.

Claim (i) is implemented as continuous gradience in the activity of structures in input representations. Specifically, structures in the input can be represented with any level of activity, or degree of presence, between 0 and 1. Structures with partial activity are often referred to as *gradient symbolic representations* or *gradient symbols*.[1] This contrasts with "categorical" Harmonic Grammar, in which all structures are either fully present or absent. Claim (ii) is implemented in the form of *conjunctive blends*; single positions in an input structure can be associated with multiple, partially active symbols. Following Smolensky Goldrick, & Mathis, (2014) and Smolensky & Goldrick (2016), many analyses in the framework adopt the working assumptions that output representations cannot contain blends or non-discrete structures; input structures with partial activity must must be either deleted or realized with full activity (1.0) in the output. I discuss questions around gradience in output structures in Section 5.

Crucially, the penalty of each constraint violation is proportional to the activity of the structure that incurs it. I illustratrate the effects of gradient activity on the evaluation of faithfulness constraints in tableau (1). The input contains a final consonant represented with an activity value of 0.75. The penalty of each violation of MAX, which penalizes deletion, is proportional to the amount of activity that is deleted in the input-output mapping.[2] The MAX penalty incurred by deleting the final consonant is the constraint weight times the amount of deleted activity, i.e. 8 × (0.75) = 6. The penalty of each violation of DEP, which penalizes epenthesis, is proportional to the

---

[1] The use of "gradience" in this context should not be conflated with other uses of the term to describe degrees of acceptability, continuous acoustic or articulatory states, or relative probability of occurrence (frequency). The use of gradient symbols is in principle compatible with any constraint-based model that uses weighted constraints, whether it produces categorical input-out mappings, a frequency distribution over output candidates (Maximum Entropy Harmonic Grammar; Goldwater & Johnson, 2003) or gradient acceptability scores (Linear Optimality Theory; Keller, 2010).

[2] Smolensky & Goldrick (2016) and several other works (Rosen 2016; 2018; 2019a; 2019b; Smolensky, Rosen & Goldrick, 2020) posit that MAX assigns a positive harmony reward to activity that persists in the output, rather than a penalty for the deletion of input activity. This distinction is not crucial to the patterns and analyses discussed in this paper, because it does not alter the relative harmony difference between output candidates that delete structure and those that maintain it.

amount of activity that is added to bring the violating segment's activity to 1, i.e. the difference between 1 and the input activity value. The Dep penalty of realizing the final consonant as a discrete segment in the output is the constraint weight times (1-0.75), i.e. 4 × (1-0.75) = 1. In a candidate set where all outputs must be discrete, underlying activity differences affect only the violations of faithfulness constraints. The output [dat] incurs one "full" violation of NoCoda because the structure that violates it, [$t_{1.0}$] in the syllable coda, is discrete.

(1) *Effects of gradience on penalties of Max, Dep*

/d a $t_{0.75}$/	Dep *w*=4	Max *w*=8	NoCoda w=2	H
☞ a. d a $t_{1.0}$	-0.25		-1	-3
b. d a		-0.75		-6

The important consequence for linguistic analyses is that two input representations that contain identical structures, and differ uniquely in the activity associated with those structures, can be associated with different optimal outputs, given a single set of constraints and constraint weights. The remainder of the paper examines key empirical support for Gradient Harmonic Grammar in language patterns that can be captured in terms of systematic effects of a grammatical property on gradient activity. I discuss how these analyses satisfy theoretical desiderata related to typological restrictiveness, representational parsimony (possible linguistic structures), computational parsimony (possible grammatical operations), and the psychological grounding of representations.

While the illustration above applies gradient activity to a consonant segment, the default assumption in Gradient Harmonic Grammar is that gradience can be a property of any structure in linguistic representations. In phonological analyses, gradient activity has been proposed to characterize featural dependents of segments (Hsu, 2019; Revithiadou & Markopoulos, 2021; Rosen, 2016), tonal autosegments (Kushnir, 2019; Rosen, 2019a; Zimmermann, 2018a), feet (Zimmermann, 2018b), whole morphemes or allomorphs (Faust & Smolensky, 2016; Rosen, 2019b), and prosodic boundary symbols (Kawahara & Tanaka, 2021). In syntax, gradient activity has been posited on movement-triggering features and the features that they target (Hsu, 2021; Lee & Müller, 2018; Müller, 2019; Müller, Englisch, & Opitz, to appear), and structural relations between verbs and prepositional phrase dependents (Kim, Rawlins, & Smolensky, 2019).

3. Gradient symbols and idiosyncrasy

The domain that has seen the most applications of Gradient Harmonic Grammar is the analysis of patterns that are described as *idiosyncratic* and/or *exceptional*, i.e. not predictable from more general or contextual properties (e.g. loanword status, grammatical category, position within morpho-syntactic or prosodic structure). The essential claim is that morphemes that contain with the same set of underlying symbols can contrast in the activity values that those symbols are represented with, leading them to pattern differently with respect to grammatical processes or restrictions (Faust & Smolensky, 2016; Guekguezian & Jesney, 2021; Hsu, 2019; Revithiadou, 2020; Revithiadou & Markopoulos, 2021; Zimmermann, 2018a, 2019).

The basic system is illustrated with an example of idiosyncrasy in vowel reduction in Palauan (Flora, 1974; Wilson, 1972; Zuraw, 2003). In this language, stem vowels occur in their full forms in bare stems, but can be deleted in the presence of suffixes that attract stress. Deletion does not apply equally to all stems in the language; some stem vowels undergo deletion in these contexts (2a), while others do not (2b).[3]

(2)	a.	ð**í**ŋəs	‘satisfaction’	θŋəs-ɛ́l	‘his/her satisfaction’
		ð**í**k	‘wedge’	θk-ɛ́l	‘his/her wedge’
	b.	ð**í**ŋ	‘ear’	ð**i**ŋá-l	‘his/her ear’ (/a/ is part of the stem UR)
		ð**ó**ko	‘swim bladder’	ð**o**k-ɛ́l	‘his swim bladder’

For the constraint-based analysis, I use the faithfulness constraints Max and Dep, and a markedness constraint *UnstressedV that penalizes any vowel that does not bear stress. We restrict our attention to two candidate types: one that maintains the stem vowel (rendered as a discrete segment with 1.0 activity), and one that deletes it. I abstract away from representations and constraints that determine stress placement and consonant voicing changes.

The tableaux in (3) and (4) show one set of constraint weights for which the two types of pattern result from a contrast in the input activity values associated with deleting versus non-deleting vowels. The input structure in both instances includes an underlying stem vowel /i/; however, they differ in the activity associated with the corresponding root node (0.75 vs. 0.25). For a stem vowel with 0.75 input activity such as /ð $i_{0.75}$ ŋa/ in (3), deletion of the vowel incurs a sufficiently high violation of Max (16 × 0.75 = 12), which exceeds the combined penalties of *UnstressedV (4 × 1 = 4) and Dep (4 × 0.25 = 1) incurred by producing a full vowel. However, for a stem vowel with 0.25 input activity such as /ð $i_{0.25}$ ŋəs/ in (4), the penalty of Max (16 × 0.25 = 4) is lower than the combined penalties of *UnstressedV (4 × 1 = 4) and Dep (4 × 0.75 = 3). Here, deletion is the optimal output.

(3) *No reduction of stem vowels with relatively high activity*

/ð **i**$_{0.75}$ ŋa-ɛl/	*UnstressedV *w*=4	Dep *w*=4	Max *w*=16	H
☞ a. ðiŋ-ál	-1	-0.25		-5
b. θŋ-ál			-0.75	-12

(4) *Reduction of stem vowels with relatively low activity*

/ð **i**$_{0.25}$ ŋəs-ɛl/	*UnstressedV *w*=4	Dep *w*=4	Max *w*=16	H
a. ðiŋəs-ɛ́l	-1	-0.75		-7
☞ b. θŋəs-ɛ́l			-0.25	-4

It is important to note that although Gradient Harmonic Grammar permits input structures to have a continuous range of activity states, the number of possible output forms does not differ from that of categorical Harmonic Grammar (assuming that outputs contain only discrete structures). As in Harmonic Grammar, the inventory of outputs is restricted by the operations of Gen, the set of constraints, and their associated weights. For example, given the set of constraints and weights in the tableaux above, there is a unique *threshold* activity value 0.6 on stem vowels

---

[3] The example is used here only to illustrate constraint evaluation with gradient symbols. It is not intended as a full analysis of vowel reduction in Palauan, whose conditioning factors are substantially more complex (see Zuraw, 2003 for an overview).

that determines whether they will surface; this is the activity value for which the output candidates associated with deletion and non-deletion have the same harmony. A stem vowel with any amount of input activity above the threshold value surfaces faithfully, while a stem vowel with any amount of input activity below that threshold is deleted (see Smolensky & Goldrick, 2016 for further discussion of thresholds). Crucially, input activity thresholds are emergent from the set of constraints and their weights, and arise only in a subset of possible weighting conditions. For instance, in a hypothetical grammar in which the weight of *UNSTRESSEDV exceeds *w*=12, but maintains the same weights for MAX and DEP, vowel deletion is the optimal output for all possible values of underlying activity on stem vowels. Output selection is determined only by *relative* activity values and constraint weights, rather than particular numeric values. The upshot is that it is easy to prevent the overgeneration of rampant idiosyncratic patterns.

This approach to idiosyncrasy as the result of input contrasts in gradient activity joins a long list of generative treatments of these patterns, including morpheme-specific rules (Chomsky & Halle, 1968), idiosyncratic underlying structures (Hyman, 1970), idiosyncratic underspecification of features (Inkelas, 1994; Kiparsky, 1993), morpheme-specific constraints (Pater, 2000), and morpheme-specific constraint scaling (Coetzee & Kawahara, 2013; Coetzee & Pater, 2011). As space constraints prohibit an exhaustive comparison with prior accounts, I focus on key strengths of the Gradient Harmonic Grammar approach in terms of empirical coverage, parsimony of analysis, and the plausibility of its psychological grounding.

3.1 *Clustering and implicational relations among idiosyncratic processes*

The analysis of idiosyncrasy in terms of gradient activity contrasts permits a parsimonious explanation of two key generalizations involving idiosyncratic patterns. First, it predicts *clustering* relations: the tendency for structures that show idiosyncrasy for one process to pattern idiosyncratically for other processes that are sensitive to the same type of structure (Zimmermann 2018b; 2020). Second, it predicts *implicational* relations, in which all forms that undergo one idiosyncratic process are subject to a second one, but not vice versa (Hsu, 2019; Müller, 2019; Revithiadou & Markopoulos, 2021; Zimmermann, 2020).

I illustrate clustering with two idiosyncratic processes in Barrow Iñupiaq, involving vowels that surface as [i] is some contexts (Archangeli & Pulleyblank, 1994; Kaplan, 1981). On some lexical items, underlying /i/ is subject to a process of labial assimilation in the presence of relativizer *-m*, which alters it to the round back vowel [u]. On the same forms, underlying /i/ undergoes a separate process of dorsal assimilation, which lowers the vowel to [a] before velar and uvular consonants.

(5) *Correlation among stems for labial assimilation and dorsal assimilation of /i/*

Root	*Labial assimilation*	*Dorsal assimilation*
/kamik/ ‘boot’	Yes [kamŋ**u**m] ‘boot-REL’	Yes [kamm**a**k] ‘boot-DUAL’
/amiq/ ‘skin’	No [ammim] ‘skin-REL’	No [ammik] ‘skin-DUAL’

I assume that labial assimilation is driven by a markedness constraint against unround vowels before [m], $*V_{[\text{-RD}]}$m, and that dorsal assimilation is driven by a markedness constraint against dorsal consonants preceded by a high vowel, $*V_{[\text{HI}]}C_{[\text{DOR}]}$. Vowel quality changes incur violations of faithfulness constraints DEPV$_{[\text{PLACE}]}$ and MAXV$_{[\text{PLACE}]}$. Furthermore, I posit that if a vowel root node is represented with an activity below 1, its featural dependents, such as its place features, cannot exceed this activity level. This accounts for the fact that features related to height and rounding are equally susceptible to change on a given vowel.

The tableaux below show one set of activity values and constraint weights that generates the Iñupiaq pattern. A vowel with a relatively high input activity (above the threshold of 0.667) does not undergo either process (6), whereas vowels with input activity below that threshold undergo both labial assimilation and dorsal assimilation (7).[4] For the calculation of constraint penalties, the surfacing of the vowel as [i] incurs a DEPV$_{[\text{PLACE}]}$ violation of the difference between 1 and the input activity of its high place feature; in (6a), faithful realization of input /i/ incurs a DEP penalty of 2 × (1-0.75) = 0.5. An output candidate that contains a different vowel in its place requires both deletion of the gradient place feature in the input (penalized by MAXV$_{[\text{PLACE}]}$) and the epenthesis of a distinct place feature (incurring a full violation of DEPV$_{[\text{PLACE}]}$); in (6b), the output candidate with [u] incurs MAXV$_{[\text{PLACE}]}$ penalty of 4 × (0.75) = 3 and a full violation of DEPV$_{[\text{PLACE}]}$.

(6) a. *Vowels with relatively high activity do not undergo labial assimilation*

/ k a m $i_{0.75}$ k - m/	DEPV$_{[\text{PLACE}]}$ *w*=2	MAXV$_{[\text{PLACE}]}$ *w*=4	$*V_{[\text{-RD}]}m$ *w*=4	H
☞ a. kamŋim	-0.25		-1	-4.5
b. kamŋum	-1	-0.75		-5

b. *Vowels with relatively high activity do not undergo dorsal assimilation*

/k a m $i_{0.75}$ k - k/	DEPV$_{[\text{PLACE}]}$ *w*=2	MAXV$_{[\text{PLACE}]}$ *w*=4	$*V_{[\text{HI}]}C_{[\text{DOR}]}$ *w*=4	H
☞ a. kammik	-0.25		-1	-4.5
b. kammak	-1	-0.75		-5

(7) a. *Vowels with relatively low activity undergo labial assimilation*

/ k a m $i_{0.5}$ k - m/	DEPV$_{[\text{PLACE}]}$ *w*=2	MAXV$_{[\text{PLACE}]}$ *w*=4	$*V_{[\text{-RD}]}m$ *w*=4	H
a. kamŋim	-0.5		-1	-5
☞ b. kamŋum	-1	-0.5		-4

b. *Vowels with relatively low activity undergo dorsal assimilation*

/k a m $i_{0.5}$ k - k/	DEPV$_{[\text{PLACE}]}$ *w*=2	MAXV$_{[\text{PLACE}]}$ *w*=4	$*V_{[\text{HI}]}C_{[\text{DOR}]}$ *w*=4	H
a. kammik	-0.5		-1	-5
☞ b. kammak	-1	-0.5		-4

Crucially, the fact that vowels that undergo dorsal assimilation also undergo labial assimilation (and vice versa) follows from a fundamenal property of how harmony is calculated in Gradient

---

[4] This threshold is the input activity value for which the candidate harmony of realizing the input vowel as a discrete [i] (violating DEPV[PLACE] and *V[-RD]m) equals the harmony of place assimilation (violating DEPV[PLACE] and MAXV$_{[\text{PLACE}]}$), given this set of constraint weights.

Harmonic Grammar: any change in the activity of a representational unit (here, a vowel and its place features) proportionally affects the penalties of *all and only* the faithfulness constraints that refer to that structure.

This analysis of Barrow Iñupiaq requires the relative penalties of repair vs. non-repair for each assimilation process to be identical for all levels of input activity. Using a different set of constraint weights, one can also generate systems where two processes are subject to distinct activity thresholds. For example, a hypothetical language that differs from Barrow Iñupiaq only by having a higher weight w = 4.5 for $\text{*V}_{[\text{HI}]}\text{C}_{[\text{DOR}]}$ would show an *implicational threshold pattern*: Input vowels with activity below 0.667 would undergo both processes, those with activity between 0.667 and 0.75 would undergo only labial assimilation, and those with activity above 0.75 would resist both processes. Although this does not characterize Barrow Iñupiaq, patterns of precisely this kind are attested. For example, in Northern Standard Greek, clitic-final nasals vary in whether they undergo two types of processes: deletion before fricatives and coalescence with a following voiceless obstruent (Revithiadou & Markopoulos, 2021). The two processes show an implicational relation: all clitic nasals that undergo deletion also undergo coalescence, while the converse relation does not hold (8). Hsu (2019) analyzes a similar implicational relation between [n]-liaison and vowel denasalization in French.

(8) *Implicational relations between deletion and coalescence* (Revithiadou and Markopoulos 2021; 62)

	Deletion before fricatives?	*Coalescence before voiceless obstruents?*
Class 1 ex. DET.M.PL.GEN	No – ex. [to**n** θiɾon]	No – ex. [to**n** piston]
Class 2 Ex. CLITIC.M	No – ex. [to**n** θelo]	Yes – ex. [to ^{m}bistevo]
Class 3 Ex. CLITIC.F	Yes – ex. [ti θelo]	Yes – ex. [ti ^{m}bistevo]

Such patterns in which structures in a language exhibit “degrees” of idiosyncratic patterning are readily captured in Gradient Harmonic Grammar because the dimension of possible activity contrasts is numerically continuous. In contrast, many such patterns would pose significant difficulties for alternative analyses relying on featural underspecification alone (for discussion of specific cases, see Guekguezian & Jesney, 2021; Revithiadou & Markopoulos, 2021; Zimmermann, 2019, 2020).

Clustering and implicational relations among idiosyncratic patterns are not as straightforwardly predicted by approaches where the locus of idiosyncrasy lies not in representational properties of idiosyncratic/exceptional structures, but purely in item-specific effects on the evaluation of constraint penalties. There are several proposals of this kind. First, one may posit lexically indexed constraints (Pater, 2000, 2010), which are only violated if a penalized structure contains part of a morpheme bearing the relevant index. Within weighted constraint frameworks, item-specific effects can be generated by lexical constraint scaling (Coetzee & Kawahara, 2013; Coetzee & Pater, 2011; Linzen, Kasyanenko, & Gouskova, 2013); the penalty of a markedness or faithfulness constraint is adjusted if the violating structure contains a particular morpheme. A similar reweighting approach in Sande, Jenks, & Inkelas (2020) posits that individual morphemes can impose specific numeric weights on one or more constraints, which affect constraint evaluation on all structures within a phasal domain. The issue for these approaches arises in situations where two or more idiosyncratic processes in a language each require the adjustment (indexing, scaling, or reweighting) of a different set of constraints. In these cases, there is no expectation that the two types of adjustments should each be associated with the same set of lexical items, as needed to account for clustering patterns. Similarly, there is no expectation for languages to show a preference for patterns where the set of items associated with one adjustment is a proper subset of the set of items associated with the second adjustment, as needed to account for implicational patterns.

3.2 *Interaction with other continuous influences on harmony*

A key property of weighted constraint systems is that large numbers of constraints, even those related to different types of principles and structures, interact in a single component of grammar (Zuraw & Hayes, 2017). We can thus find support for the inclusion of gradient symbols in Harmonic Grammar from cases where activity contrasts that account for idiosyncrasy interact with constraint families that are not directly sensitive to these contrasts. For example, this arises in language patterns where an idiosyncratic pattern replicates a regular pattern that applies in a distinct morphological or prosodic context (Hsu, 2019; Revithiadou, 2020; Revithiadou & Markopoulos, 2021). These patterns arise from a predicted type of interaction between gradient activity contrasts and positionally scaled markedness constraints (Hsu, 2019; Hsu & Jesney, 2016), proposed independently in categorical Harmonic Grammar.

I illustrate this with an abridged version of Revithiadou and Markopoulos' (2021) Gradient Harmonic Grammar analysis of the patterning of coronal nasals in northern varieties of Standard Greek. Within words, coronal nasals cannot precede voiceless stops; underlying NC̥ sequences coalesce into a prenasalized voiced stop.

(9) *Obligatory coalescence within words*
/sin-paθ-i-a/ > $_{\omega\text{min/max}}$(siˈ$^{\text{m}}$baθia) ‘compassion’
/sin-tak-s-i/ > $_{\omega\text{min/max}}$(ˈsi$^{\text{n}}$daksi) ‘syntax’

In contrast, coronal nasals associated with clitics and other functional elements show item-specific variation when they precede a stem-initial voiceless stop. Some morphemes surface with [n] (10), while others undergo coalescence (11), replicating the stem-internal distribution of nasals. Revithiadou and Markopoulos present independent evidence (based on glide strengthening, syllabification, secondary stress placement) that clitics are prosodified in a recursive prosodic word (ω) structure, regardless of whether coalescence applies: all stem material is contained both a minimal ω (ωmin) and maximal ω (ωmax), while affixes are parsed within ωmax only. This rules out an alternative analysis of idiosyncrasy in terms of prosodic prespecification (Bennett, Harizanov, & Henderson, 2018; Inkelas, 1989).

(10) *Item-specific: non-coalescence across ω boundary*
/t-on pist-on/ > $_{\omega\text{max}}$(to**n** $_{\omega\text{min}}$(**p**i'ston))
the-PL.GEN believer-PL.GEN/

(11) *Item-specific: coalescence across ω boundary*
/t-on pist-o/ > $_{\omega\text{max}}$(to $_{\omega\text{min}}$($^{\mathbf{m}}$**b**i'sto))
the-M.SG.ACC believer-M.SG.ACC

Revithiadou and Markopoulos posit that the violations of the constraint against nasal consonants followed by voiceless stops, *NC̥ (Pater, 1999), is *scaled* according to the prosodic domains in which the structure occurs (Hsu, 2019). Specifically, each output NC̥ sequence incurs a penalty of $w + (s \times d)$, where w is the basic constraint weight, s is a constraint-specific scaling factor, and d is a value along a prominence scale (0, 1, … n), determined by the smallest prosodic constituent that contains the NC̥ sequence. Concretely, I assume that NC̥ sequences receive a scaled penalty of $s \times 1$ if the sequence is fully contained in a ωmax, and $s \times 2$ if it is fully contained in a ωmin. Output candidates generated by coalescence incur a full violation of UNIFORMITY; following Smolensky & Goldrick (2016; 17-18), this constraint is not sensitive to gradient activity contrasts because it only evaluates correspondence relationships between input and output segments. For input structures with partially active /n/, realization of [n] in the output incurs a DEP violation proportional to amount of added activity. The following tableaux show one set of weights, scaling factors, and gradient activity levels that generate the attested pattern.

Because hypothetical stem-internal NC̥ sequences are fully contained within the most embedded prosodic domain ωmin, *NC̥ violations are subject to a relatively high scaling factor; an output NC̥ sequence incurs a penalty of $1_w + (2_s \times 2_d) = 5$. Given the set of weights and scaling factors in (12), there is no level of input activity on the nasal that prevents the candidate with coalescence from being optimal.

(12) *All stem-internal /NC̥/ sequences undergo coalescence*

/sin$_{0.8}$-tak-s-i/	DEP *w*=4	UNIFORMITY *w*=4	*NC̥ *w*=1, *s*=2	H
a. $_{\omega\text{max/min}}$(si**nt**aksi)	-0.2		$-1(1_w+(2_s\times 2_d))$	-5.8
☞ b. $_{\omega\text{max/min}}$(si$^{\mathbf{n}}$**d**aksi)		-1		-4

/sin$_{0.5}$-tak-s-i/	DEP *w*=4	UNIFORMITY *w*=4	*NC̥ *w*=1, *s*=2	H
a. $_{\omega\text{max/min}}$(si**nt**aksi)	-0.5		$-1(1_w+(2_s\times 2_d))$	-7
☞ b. $_{\omega\text{max/min}}$(si$^{\mathbf{n}}$**d**aksi)		-1		-4

In contrast, because clitic nasals are parsed outside of ωmin, the output candidate that faithfully realizes the nasal incurs a smaller scaled penalty of *NC̥: $1_w + (2_s \times 1_d) = 3$. As a result, the optimal treatment of an underlying nasal is subject to an activity threshold of 0.75. Clitic nasals with input activity above this threshold surface because the UNIFORMITY penalty incurred by coalescence exceeds the combined penalty of scaled *NC̥ and DEP (13). Clitic nasals with input activity below the threshold undergo coalescence due to the proportionally higher DEP penalty (14). The latter pattern in effect replicates the regular stem-internal distribution of nasals.

(13) *Clitic nasal with relatively high activity – no coalescence*

/t-on$_{0.8}$ piston/	DEP w=4	UNIFORMITY w=4	*NC̥ w=1, s=2	H
☞ a. $_{\omega\text{max}}$(to**n** $_{\omega\text{min}}$ (**p**i'ston))	-0.2		-1(1_w+(2_s×1_d))	-3.8
b. $_{\omega\text{max}}$(to $_{\omega\text{min}}$ ($^{\mathbf{m}}$**b**i'ston))		-1		-4

(14) *Clitic nasal with relatively low activity – coalescence optimal*

/t-on$_{0.5}$ pisto/	DEP w=4	UNIFORMITY w=4	*NC̥ w=1, s=2	H
a. $_{\omega\text{max}}$(to**n** $_{\omega\text{min}}$ (**p**i'sto))	-0.5		-1(1_w+(2_s×1_d))	-5
☞ b. $_{\omega\text{max}}$(to $_{\omega\text{min}}$ ($^{\mathbf{m}}$**b**i'sto))		-1		-4

Such patterns result from a predicted type of interaction between two continuous influences on harmony in Gradient Harmonic Grammar: constraint strength (numeric weighting, scaling) and gradient activity. Crucially, this kind of resemblance between exceptional patterns in one morpho-syntactic domain and regular patterns in another one would be coincidental in approaches that account for lexical idiosyncrasy using special rules or special underlying structures (see Hsu 2019 for further discussion).

3.3 *Psychological grounding of gradient activity contrasts*

To the extent that contrasts in gradient activity can successfully model the linguistic representation of item-specific behavior, Gradient Harmonic Grammar offers the potential to develop explicit models for how these representations are linked to their psychological grounding factors (Smolensky & Goldrick, 2016; Smolensky, Rosen, & Goldrick, 2020). This is because gradient activity and many substantive factors known to influence item-specific behavior, such as item frequency, collocation frequency, perceptual cue strength, and processing difficulty, can be stated in terms of numeric quantities. It may thus be possible to build a mathematically precise model of how these various factors determine the activity levels of language structures.

Though not explicitly couched in Gradient Harmonic Grammar, Inkelas (2015) illustrates how certain patterns can arise from activity contrasts that are systematically grounded in this way.[5] In particular, she focuses on patterns often referred to as morphologically derived environment effects, in which a process targets only segments at the edges of stems, but not in stem-internal contexts where its conditioning environment is apparently met (Kean, 1974; Kiparsky, 1973, et seq.; Łubowicz, 2002, among may others). Consider a well-known case in Turkish, in which stem-final velars are deleted before certain suffixes that would make them surface between two vowels (Inkelas & Orgun, 1995; Lees, 1961). However, intervocalic velar consonants are permitted in stem-internal positions.

---

[5] Inkelas posits that structures in input representations vary in their "degree of confidence," reflecting their robustness of storage in memory, and that structures represented with higher confidence incur greater penalties of faithfulness constraints. Based on this usage, we can consider this notion to be equivalent to activity in Gradient Harmonic Grammar.

(15) a. *Deletion of intervocalic stem-final velars*

/oluk-u/	>	[oluu]	‘gutter-ABL’
/salak-ɨm/	>	[sala-ɨm]	‘stupid-1.SG’

b. *No deletion of stem-internal velars*

[sekiz]	(*[seiz])	‘eight’
[oku]	(*[ou])	‘read’

Inkelas posits that stem-final velars are uniquely susceptible to deletion because they are systematically less robustly represented in memory than their stem-internal equivalents, based on several criteria. First, stem-internal velars occur in a more predictable phonological context; medial *k* in *sekiz* is always flanked by vowels, whereas stem-final *k* can be followed by a greater variety of segments. Second, stem-final velars are often perceptually less robust; stem-internal *k* is uniformly in an onset, while stem-final *k* in unaffixed words occurs in a coda (where contrasts are more difficult to perceive). If these differences lead to stem-final consonants being systematically represented with lower activity values than their stem-internal equivalents, they can become uniquely susceptible to deletion under some weighting conditions, as shown in (16).

(16) a. *Relatively high activity on stem-internal velars – no deletion*

/s e $k_{1.0}$ i z/	*VGV *w*=6	DEP *w*=4	MAX *w*=8	H
☞ a. s e $k_{1.0}$ i z	-1			-6
b. s e i z			-1	-8

b. *Relatively low activity on stem-final velars – deletion applies*

/o l u $k_{0.5}$ - u/	*VGV *w*=6	DEP *w*=4	MAX *w*=8	H
a. o l u k_{1} u	-1	-0.5		-8
☞ b. o l u u			-0.5	-4

Much work remains to be done in the creation and testing of models that explicitly link gradient activity in mental representations to quantifiable grounding factors. However, in some promising early work in this area, Müller, Englisch, & Opitz, (to appear) note that collocation frequencies between noun-verb pairs in a corpus can be used to predict gradient activity values that influence restrictions on prepositional phrase extraction from noun phrases in German. More broadly, there are many potential avenues of research in identifying the factors that systematically influence activity in linguistic representations, including the aforementioned psychological and cognitive factors, specific language structures or processes (ex. proposed effects of reduplication on activity; Zimmermann 2021), and patterns of diachronic change (Hsu, 2019: 241). One key potential benefit of such work is the creation of rigorous criteria that can be used to evaluate proposed uses of gradient symbols in linguistic analyses.

4. Blended representations

The second key innovation of Gradient Harmonic Grammar lies in its use of conjunctive blends in linguistic representations. In practice, the term is used to describe two types of pattern analyzed by Smolensky & Goldrick (2016), which I discuss in turn. Although blending is a distinct theoretical notion from gradient activity, gradient activity contrasts play a critical role in blending analyses.

The first type of blend involves cases in which the optimal output is determined by the combined activity of two or more partially active symbols in the input, each associated with a distinct underlying position. This type of analysis is well suited to idiosyncratic patterns that must be described in terms of specific combinations of lexical items, rather than unique sets of triggering or undergoing items.[6] I illustrate this with an adapted version of Rosen's (2016) analysis of Japanese *rendaku*, the occurrence of voicing on the first consonant of the second member of a compound, which would not be voiced in isolation (17a, 17c).

(17)	W1	W2		Compound	
a.	[kɯma] 'bear'	[te] 'hand'	>	[kɯma-**d**e]	'rake'
b.	[yama] 'mountain'	[te] 'hand'	>	[yama-te]	'mountainside'
c.	[yama] 'mountain'	[tori] 'bird'	>	[yama-**d**ori]	'mountain-bird'

The challenge of rendaku for phonological analyses is that in some cases, whether rendaku occurs in a particular compound cannot be predicted by the identity of its first member or second member alone, putting aside all phonological restrictions on the process (see Vance 2015 for an overview). For example, *te* 'hand' as the second member of a compound undergoes rendaku in *kuma-de* 'rake' (17a) but not *yama-te* 'mountainside' (17b). Neither does *yama* as the first member of a compound predict whether rendaku will occur on the second item (compare 17b, 17c).

At the same time, the pattern is not fully unpredictable; individual stems vary in their overall propensity to occur with rendaku across attested compound forms that they occur in. Rosen (2016) posits that all morphemes occur with a [VOI] feature at a compound juncture, each with a lexically specified activity value. In compounds, rendaku occurs if the additive activity of the two [VOI] features of each stem exceeds a threshold that makes deletion of these features costlier than producing a voiced consonant. One set of constraints, weights and activity values that generates the pattern in (17) is shown below. In the first tableau in (18), rendaku (candidate a) incurs a DEP[VOI] penalty for the amount of remaining activity that must be epenthesized to produce a fully voiced output segment (1 − (0.4 + 0.225) = 0.375). The absence of rendaku (candidate b) incurs a MAX[VOI] penalty for the deletion of each input feature (0.4 + 0.225 = 0.625).

(18)

/kuma([voi]$_{0.4}$)+([voi]$_{0.225}$)te/	MAX[VOI] *w*=1	DEP[VOI] *w*=1	H
☞a. kuma-de		-0.375	-0.375
b. kuma-te	-0.625		-0.625

---

[6] Sande (2019; 2020) refers to these as *doubly conditioned* or *doubly triggered* processes, in the context of proposed analyses using cophonologies by phase (Sande et al., 2020). The extent to which this approach and a gradient symbols analysis make different predictions remains an open question; see Zimmermann (2021b) for one recent discussion.

/yama($[\text{voi}]_{0.225}$)+($[\text{voi}]_{0.225}$)te/	Max[voi] *w*=1	Dep[voi] *w*=1	H
a. yama-de		-0.55	-0.55
☞b. yama-te	-0.45		-0.45

/yama($[\text{voi}]_{0.225}$)+($[\text{voi}]_{0.4}$)tori/	Max[voi] *w*=1	Dep[voi] *w*=1	H
☞ a. yama-dori		-0.375	-0.375
b. yama-tori	-0.625		-0.625

Rosen identifies a set of seven activity levels that can be associated with each inconsistent rendaku undergoer, which generate their attested forms without contradictions across possible combinations of forms. This analysis is more parsimonious than alternative grammatical descriptions that require lexical listing of exceptional rendaku blockers, or listed allomorph sets for all morphemes in compounds, which would not also not directly encode information about lexical propensities in the phonological representation of each morpheme. Rosen (2018; 2019a) extends a similar account to variability in pitch accent patterns among compounds in Japanese.

In the second type of blending analysis, a single structural position in an input representation is associated with multiple, simultaneously present, partially active symbols. Broadly, this accounts for patterns that involve a choice in the form that is associated with a unique structural position, which is determined by both contextual and lexically-specific factors. I illustrate this with the pattern of allomorphy of the genitive marker in Djabugay (Mascaró, 2007; Patz, 1991), which occurs as *-ŋun* after consonant-final stems, and as *-n* after vowel-final stems (19). The choice of allomorph is partly phonologically determined; the occurrence of *-ŋun* with consonant-final stems appears to be motivated by the general ban against complex codas in the language. However, the occurrence of *-n* after vowel-final stems is not similarly predictable; both attested *guludu-n* and hypothetical **guludu-ŋun* satisfy the language's general syllable structure restrictions. Thus, *-n* is the default allomorph that occurs in contexts where neither allomorph is preferred by active phonological restrictions.

(19) gaŋal ‘goanna lizard’ gaŋal-ŋun (*gaŋal-n)
guludu ‘dove’ guludu-n (*guludu-ŋun)

The pattern can be generated in Gradient Harmonic Grammar as follows, assuming that morphemes themselves can be represented with partial activity (Faust & Smolensky, 2016). The genitive suffix is represented in the input as a single position in structure with a conjunctive blend of both allomorphs, each represented with partial activity {n $_{0.5}$, ŋun $_{0.25}$}. We further assume that Gen only produces output structures where each position is associated with exactly one fully active symbol. Allomorph selection is determined in part by versions of the constraints Max and Dep that are sensitive to the activity of each realized or deleted allomorph (rather than activity values of individual segments). The realization of *-n* in this position in an output incurs a Dep violation proportional to the remaining activity needed to make this structure discrete (1-0.5=0.5), and a Max violation proportional to the activity of the deleted allomorph *-ŋun* (0.25). Given the activity values and constraint weights below, *-n* is selected as the optimal output when no other constraint violations are at stake, simply because it has a higher input activity than *-ŋun*, making its deletion more costly (20a). In the presence of a consonant-final stem, however, selection of *-n* incurs a sufficiently large penalty of *Complex, making *-ŋun* the optimal selected allomorph (20b).

(20) a. *Selection of default allomorph with higher input activity*

/guludu-{n $_{0.5}$, ŋun $_{0.25}$}	*Complex *w*=3	Dep *w*=4	Max *w*=4	H
☞ a. gu.lu.du**n**		-0.5(n)	-0.25(ŋun)	-3
b. gu.lu.du.**ŋun**		-0.75 (ŋun)	-0.5(n)	-5

b. *Selection of allomorph with lower input activity to avoid a markedness penalty*

/gaŋal-{n $_{0.5}$, ŋun $_{0.25}$}	*Complex *w*=3	Dep *w*=4	Max *w*=4	H
a. ga.ŋal**n**	-1	-0.5(n)	-0.25(ŋun)	-6
☞ b. ga.ŋal.**ŋun**		-0.75 (ŋun)	-0.5(n)	-5

Key aspects of this analysis resemble the proposal by Bonet, Lloret, & Mascaró (2007) and Mascaró (2007) couched in Optimality Theory. In this view, all allomorphs of a morpheme are present in the input to the phonological grammar, listed as a (partially) ordered set. The preference for the default allomorph is driven by a constraint Priority, violated by deviations from the listed order in the lexical entry. This preference is overridden if there is a violation of a higher ranked markedness constraint, such as *Complex, that can be avoided by selecting a less prioritized allomorph. Gradient Harmonic Grammar allows the insights of this approach to be captured in a more theoretically parsimonious way. Apparent preferences among allomorphs result from contrasts in gradient activity, which independently explain other phenomena, and the interaction of well-known families of constraints (Max and Dep). This obviates the need for phonological representations to contain ordered lists, a novel type of object, and special constraint types that refer to them (Priority).

This use of structural positions that contain blends of partially active symbols is a key component of Smolensky & Goldrick's (2016) analysis of French liaison. In the domain of allomorphy, Zimmermann (2019) analyzes a pattern in Welsh in which two default allomorphs conflict, for which an approach using Priority would require a more complex constraint definition or lexically indexed forms of this constraint (Brinkerhoff, 2019). Rosen (2019) and Rosen (2021) apply similar analyses to relatively complex patterns of inflectional allomorphy in Mazatec and Estonian, respectively. In a separate vein, blends provide a powerful way to model intersentential code-mixing and co-production by multilingual speakers (Goldrick, Putnam, & Schwarz, 2016), and a range of performance errors in language production (Goldrick & Chu, 2014; Smolensky & Goldrick, 2016; Smolensky, Goldrick, & Mathis, 2014), further supporting the psychological reality of such structures.

5. Output gradience

Arguably the most significant unresolved questions in Gradient Harmonic Grammar involve whether gradient activity can be a property of output representations, in addition to being present in inputs. While there are several types of patterns proposed to involve output gradience, I illustrate with one such application from Zimmermann's (2020) analysis of an idiosyncratic pattern in Finnish. In this language, some suffixes that begin with [i] trigger deletion of preceding stem-final /a/, while others do not (Anttila, 2002), as shown in (21).[7] The key difference between this pattern and the aforementioned examples of idiosyncrasy is that this contrast does not involve whether a segments of a morpheme undergo a process; rather, it affects whether they trigger a process that occurs on segments of a separate morpheme.

(21) *Non-triggering suffix*: /ant**a**-isi/ > [ant**a**isi] 'give'-COND
Deletion-triggering suffix: /ott**a**-i/ > [otti] 'take'-PST

The intuition behind Zimmermann's analysis is that [ai] sequences that include a deletion-triggering suffix are more ill-formed than those that include a non-triggering suffix, such that only the former compel a repair. Specifically, Zimmermann proposes that /i/ on deletion-triggering suffixes is represented with a higher level of activity, and that this activity difference persists in the output representation. Furthermore, the penalty of the markedness constraint against [ai] sequences, *[ai], is proportional to the highest amount of activity in the structure that violates it. This is shown in the tableaux below. For a non-triggering suffix in (22a), the markedness penalty incurred by the faithful output with [ai] is lower than the penalty of deleting either of its two segments. For a triggering suffix (22b), higher activity on /$i_{1.0}$/ results in a *[ai] penalty that exceeds the MAX penalty of deleting the preceding vowel.[8] The activity difference between the two vowels also ensures that deletion of /$i_{1.0}$/ incurs a greater MAX penalty than deletion of /$a_{0.75}$/.

(22) a. *Suffix [i] with relatively low output activity – no deletion of stem-final /a/*

/ant $\mathbf{a_{0.75}}$-$\mathbf{i_{0.75}}$ si/	MAX *w*=5	*[ai] *w*=4	H
☞ a. ant $\mathbf{a_{0.75}}$ $\mathbf{i_{0.75}}$ si		-0.75	-3
b. ant $\mathbf{i_{0.75}}$ si	-0.75		-3.75
c. ant $\mathbf{a_{0.75}}$ si	-0.75		-3.75

b. *Suffix [i] with relatively high output activity – deletion of stem-final /a/*

/ott$\mathbf{a_{0.75}}$-$\mathbf{i_{1.0}}$/	MAX *w*=5	*[ai] *w*=4	H
a. ott $\mathbf{a_{0.75}}$ $\mathbf{i_{1.0}}$		-1.0	-4
☞b. ott $\mathbf{i_{1.0}}$	-0.75		-3.75
c. ott $\mathbf{a_{0.75}}$	-1.0		-5

In contrast, the difference between suffixes that trigger stem-final [a] deletion and those that do not cannot be accounted for uniquely in terms of a difference in input activity alone across suffixes. This is because any activity difference would only affect the susceptibility of suffix vowels themselves to change, and not their ability to trigger processes that affect stem vowels.

---

[7] This illustration abstracts away from many details of the full Finnish pattern to focus on fundamental properties of constraint evaluation with output gradience. Suffixes that trigger changes to stem-final /a/ do not affect all stems. Furthermore, stems that undergo changes to final /a/ vary in whether they show deletion or raising to [o].

[8] In Zimmermann's analysis, the high front vowel on triggering suffixes is represented with an output activity value of 1.5, while the vowel on non-triggering suffixes has 1.0 activity; because output selection is determined by relative constraint weights and activity values, the pattern is equally compatible with the more commonly assumed activity range between 0 and 1.0. It remains to be determined whether there are empirical patterns to support a distinction between "underactive" structures (e.g. activity below 1.0) and "overactive" structures (e.g. activity above 1.0).

The adoption of output gradience has implications for how grammatical computation proceeds in Gradient Harmonic Grammar, and raises a number of new theoretical questions. First, if structures can be represented with non-discrete levels of activity in outputs, GEN and EVAL must allow significantly larger candidate sets than a theory with input gradience only. Second, because the faithful realization of non-discrete activity in an output incurs fewer MAX or DEP violations than candidates that add or delete activity, one must appeal to new mechanisms to explain patterns where gradient input symbols are subject to thresholds that determine the surfacing versus non-surfacing of structure. One approach to would be to posit new types of constraints that penalize output structures with non-discrete activity (ex. FULL!: Zimmermann, 2019), output activity between 0 and 1.0 or greater than 1.0 (ex. *WEAK and *STRONG: Zimmermann, 2020), or outputs that contain structural blends (ex. QUANTIZATION: Goldrick, Putnam, & Schwarz, 2016). The weight of such constraints relative to other constraints in a language determines the contexts in which discrete and non-discrete output symbols may surface.

In addition, one must ask whether distinctions in output gradience are *substantive*, corresponding to a detectable phonetic or psychological property. Moreover, if output gradience is substantive, which properties can it correspond with? A range of views have been advanced on these questions. In Zimmermann's (2017, 2018a, 2018b, 2020) phonological analyses, contrasts in output gradience are not proposed to result in phonetic differences. Non-substantive gradience also characterizes syntactic applications where gradient activity distinctions on movement-triggering features only reflect their propensity to trigger movement (Lee & Müller, 2018; Müller, 2019). In contrast, Amato (2018) posits a correspondence between output activity on stops with degree of occlusion; lenition corresponds with activity reduction in the input-output mapping and gemination with activity increase. In yet another vein, Walker (2020) proposes that consonant place features are inherently specified with different levels of activity, and that these distinctions are uniform across inputs and outputs, i.e. not modifiable by GEN. Hsu (2021) makes a similar proposal in syntax, claiming that the hierarchy of pronominal strength (Cardinaletti & Starke, 1994) corresponds to relative activity differences across pronoun types.

While it seems implausible that grammars could permit all of these types of substantive and non-substantive output gradience, there is not at present a clear path towards a unified approach to these phenomena. Progress on this issue likely requires a careful comparison of alternative analysis of patterns that output gradience have been proposed for. This includes, but is not limited to, the use of constraint scaling rather than gradient activity contrasts, and alternative input-gradience-only analyses. Finally, work in this area must address the compatibility between various types of output gradience and quantization mechanics that have been proposed at the neural network level to disprefer non-discrete output states (Smolensky, Goldrick, & Mathis, 2014). A careful assessment of the theoretical benefits of various alternatives remains a major challenge.

5 Conclusion

Gradient Harmonic Grammar offers exciting solutions to a variety of analytical problems in linguistics. The descriptive and explanatory success of these analyses provides novel evidence in favor of numerically weighted constraints as the basis of computation in natural language, building upon the many insights of recent work on Harmonic Grammar. More generally, it suggests that the close consideration of language patterns in the context of broader theories of cognitive computation holds great potential value for generative linguistic theory.

Primary goals for future investigation in Gradient Harmonic Grammar include the resolution of questions on the nature of output gradience, and the quantitative testing of grounding factors posited to influence gradient activity contrasts. In addition, analyses in the framework present novel questions for theories of learning and acquisition. Broadly, what are the conditions that lead learners to posit representations with blends or gradient activity contrasts (Rosen, 2016, et seq.; Smolensky, Rosen, & Goldrick, 2020; Tessier & Jesney, 2021)? Finally, there remain many unexplored empirical domains in syntax, phonology, and beyond that may futher inform our understanding and evaluation of the framework.

Acknowledgements

For their comments and discussion on various stages of this work, I am grateful to Matt Goldrick, Elliott Moreton, Eric Rosen, Caitlin Smith, and Jennifer L. Smith.

References

Amato, I. (2018). *A gradient view of Raddoppiamento Fonosintattico*. Ms: University of Leipzig.

Anttila, A. (2002). Morphologically conditioned phonological alternations. *Natural Language & Linguistic Theory*, *20*(1), 1–42.

Archangeli, D., & Pulleyblank, D. (1994). *Grounded phonology*. Cambridge, MA: MIT Press.

Bennett, R., Harizanov, B., & Henderson, R. (2018). Prosodic smothering in Macedonian and Kaqchikel. *Linguistic Inquiry*, *49*(2), 195–246.

Bonet, E., Lloret, M.-R., & Mascaró, J. (2007). Allomorph selection and lexical preferences: two case studies. *Lingua*, *117*(6), 903–927.

Cardinaletti, A., & Starke, M. (1994). The typology of structural deficiency on three grammatical classes. *University of Venice Working Papers in Linguistics*, *4*, 41–109.

Chomsky, N., & Halle, M. (1968). *The sound pattern of English*. New York: Harper and Row.

Coetzee, A., & Kawahara, S. (2013). Frequency biases in phonological variation. *Natural Language and Linguistic Theory*, *31*(1), 47–89. https://doi.org/10.1007/s11049-012-9179-z

Coetzee, A., & Pater, J. (2011). The place of variation in phonological theory. In J. A. Goldsmith, J. Riggle, & A. C. L. Yu (Eds.), *The Handbook of Phonological Theory, 2nd edition* (pp. 401–434).

Farris-Trimble, A. W. (2008). *Cumulative faithfulness effects in phonology*. Doctoral dissertation, Indiana University.

Faust, N., & Smolensky, P. (2016). *Activity as an alternative to autosegmental association*. Ms: Université Paris 8 and Johns Hopkins University.

Flora, M. (1974). *Palauan Phonology and Morphology*. Doctoral dissertation, University of California San Diego.

Goldrick, M., & Chu, K. (2014). Gradient co-activation and speech error articulation : comment on Pouplier and Goldstein (2010). *Language, Cognition and Neuroscience*, *29*, 452–458. https://doi.org/10.1080/01690965.2013.807347

Goldrick, M., Putnam, M., & Schwarz, L. (2016). Coactivation in bilingual grammars : A computational account of code mixing. *Bilingualism: Language and Cognition*, *19*(5), 857–876. https://doi.org/10.1017/S1366728915000802

Goldwater, S., & Johnson, M. (2003). Learning OT constraint rankings using a maximum entropy model. In J. Spenader, A. Eriksson, & O. Dahl (Eds.), *Proceedings of the Workshop* *on Variation within Optimality Theory* (pp. 111–120). Stockholm: Stockholm University.

Guekguezian, P., & Jesney, K. (2021). “Epenthetic” vowels are not all equal: gradient representation in Yokuts roots and suffixes. In R. Soo, U. Y. Chow, & S. Nederveen (Eds.), *Proceedings of the 38th West Coast Conference on Formal Linguistics* (pp. 221–230).

Hsu, B. (2019). Exceptional prosodification effects revisited in Gradient Harmonic Grammar. *Phonology*, *36*(2), 225–263. https://doi.org/10.1017/S0952675719000125

Hsu, B. (2021). Harmonic Grammar in phrasal movement: an account of probe competition and blocking. In A. Farinella & A. Hill (Eds.), *NELS 51: Proceedings of the 51st Annual Meeting of the North East Linguistic Society* (pp. 237–250). Amherst, MA: GLSA.

Hsu, B., & Jesney, K. (2016). Scalar Positional Markedness and Faithfulness in Harmonic Grammar. In *Proceedings of the 51st Annual Meeting of the Chicago Linguistic Society* (pp. 241–255). Chicago: Chicago Linguistic Society.

Hyman, L. M. (1970). How concrete Is phonology? *Language*, *46*(1), 58–76.

Inkelas, S. (1989). *Prosodic constituency in the lexicon*. Doctoral dissertation, Stanford University.

Inkelas, S. (1994). The consequences of optimization for underspecification. In *Proceedings of the North East Linguistic Society 25* (pp. 287–302). Amherst, MA: GLSA.

Inkelas, S. (2015). Confidence scales : a new approach to derived environment effects. In Y. E. Hsiao & L.-H. Wee (Eds.), *Capturing Phonological Shades* (pp. 45–65). Newcastle upon Tyne: Cambridge Scholars Publishing.

Inkelas, S., & Orgun, C. O. (1995). Level ordering and economy in the lexical phonology of Turkish. *Language*, *71*(4), 763–793.

Kaplan, L. (1981). *Phonological issues in North Alaskan iñupiaq*. Fairbanks, AK: Alaska Native Language Center.

Kawahara, S., & Tanaka, Y. (2021). Partially activated morpheme boundaries in Japanese surnames. In A. Farinella & A. Hill (Eds.), *NELS 51: Proceedings of the Fifty-First Annual Meeting of the North East Linguistic Society, Volume 2* (pp. 25–38). Amherst, MA: GLSA.

Kean, M.-L. (1974). The strict cycle in phonology. *Linguistic Inquiry*, *5*, 179–203.

Kim, N., Rawlins, K., & Smolensky, P. (2019). *The complement-adjunct distinction as gradient blends: the case of English prepositional phrases*. Johns Hopkins University.

Kiparsky, P. (1973). Phonological representations. In O. Fujimura (Ed.), *Three dimensions of linguistic theory* (pp. 3–135). Tokyo: Taikusha.

Kiparsky, P. (1993). Blocking in non-derived environments. In S. Hargus & E. Kaisse (Eds.), *Studies in Lexical Phonology (Phonetics and Phonology 4)* (pp. 277–313). San Diego: Academic Press.

Kushnir, Y. (2019). *Prosodic Patterns in Lithuanian Morphology*. Doctoral dissertation, Universität Leipzig.

Lee, H., & Müller, G. (2018). Asymmetries in long-distance dependencies: a view from Gradient Harmonic Grammar. Paper presented at the Workshop on Long-Distance Dependencies, Humboldt-Universität zu Berlin, October 4-5.

Lees, R. (1961). *The phonology of Modern Standard Turkish.* Bloomington: Indiana University Publications.

Legendre, G., Miyata, Y., & Smolensky, P. (1990). Can connectionism contribute to syntax?: Harmonic Grammar, with an application. In K. Deaton, M. Noske, & M. Ziolkowski (Eds.), *Proceedings of the 26th regional meeting of the Chicago Linguistic Society* (pp. 237–252). Chicago: Chicago Linguistic Society.

Linzen, T., Kasyanenko, S., & Gouskova, M. (2013). Lexical and phonological variation in Russian prepositions. *Phonology*, *30*, 453–515.

Łubowicz, A. (2002). Derived Environment Effects in Optimality Theory. *Lingua*, *112*, 243–280.

Mascaró, J. (2007). External allomorphy and lexical representation. *Linguistic Inquiry*, *38*(4), 715–735.

McClelland, J. L., Rumelhart, D. E., & Group, T. P. R. (1986). *Parallel distributed processing: Explorations in the microstructure of cognition. Vol. 2: Psychological and biological models.* Cambridge, MA: MIT Press.

Müller, G. (2019). The third construction and strength of C: a Gradient Harmonic Grammar approach. In K. R. Christensen, H. Jørgensen, & J. L. Wood (Eds.), *The sign of the V - Papers in honour of Sten Vikner* (pp. 419–448). https://doi.org/10.7146/aul.348.107.

Müller, G., Englisch, J., & Opitz, A. (to appear). Extraction from NP, frequency, and Minimalist Gradient Harmonic Grammar. *Linguistics*.

Pater, J. (1999). Austronesian Nasal Substitution and other NC effects. In R. Kager, H. van der Hulst, & W. Zonneveld (Eds.), *The Prosody Morphology Interface* (pp. 310–343). Cambridge: Cambridge University Press.

Pater, J. (2000). Non-uniformity in English secondary stress : the role of ranked and lexically specific constraints. *Phonology*, *17*, 237–274.

Pater, J. (2009). Weighted constraints in generative linguistics. *Cognitive Science*, *33*, 999–1035. https://doi.org/10.1111/j.1551-6709.2009.01047.x

Pater, J. (2010). Morpheme-specific phonology: constraint indexation and inconsistency resolution. In S. Parker (Ed.), *Phonological argumentation: essays on evidence and motivation* (pp. 123–154). London: Equinox.

Pater, J. (2016). Universal grammar with weighted constraints. In J. J. McCarthy & J. Pater (Eds.), *Harmonic Grammar and Harmonic Serialism* (pp. 1–46). London: Equinox Press.

Patz, E. (1991). Djabugay. In R. M. W. Dixon & B. J. Blake (Eds.), *The handbook of Australian Languages, volume 4* (pp. 245–347). Melbourne: Oxford University Press Australia.

Prince, A., & Smolensky, P. (1993). *Optimality Theory: Constraint Interaction in Generative Grammar*. Ms: Rutgers University & University of Colorado, Boulder. Published 2004, Malden, Mass. & Oxford: Blackwell.

Revithiadou, A. (2020). Same exponent, different strength: a Gradient Harmonic account of allomorphy in Greek. In A. Gürer, D. Uygun-Gökmen, & B. Öztürk (Eds.), *Morphological complexity within and across boundaries: in honour of Aslı Göksel* (pp. 264–283). Amsterdam/Philadelphia: John Benjamins.

Revithiadou, A., & Markopoulos, G. (2021). A Gradient Harmonic Grammar account of nasals in extended phonological words. *Catalan Journal of Linguistics*, *20*, 57–75.

Rosen, E. (2016). Predicting the unpredictable: Capturing the apparent semi-regularity of rendaku voicing in Japanese through Harmonic Grammar. In E. Clem, G. Bacon, A. Chang, V. Dawson, E. H. Maier, A. Shen, & A. H. Skilton (Eds.), *Proceedings of the 42nd Annual Meeting of the Berkeley Linguistics Society* (pp. 235–250). Berkeley: Berkeley Linguistics Society.

Rosen, E. (2018). Predicting semi-regular patterns in morphologically complex words. *Linguistics Vanguard*, *4*(1), 1–15.

Rosen, E. (2019a). Evidence for Gradient Input Features from Sino-Japanese Compound Accent. In K. Hout, A. Mai, A. McCollum, S. Rose, & M. Zaslansky (Eds.), *Proceedings of the 2018 Annual Meeting on Phonology* (pp. 1–11). Washington, D.C.: Linguistic Society of

America.

Rosen, E. (2019b). Learning complex inflectional paradigms through blended gradient inputs. *Proceedings of the Society for Computation in Linguistics*, *2*, 102–112.

Rosen, E. (2021). Inflectional paradigms as interacting systems. *Proceedings of the Society for Computation in Linguistics*, *4*, 136–147.

Sande, H. (2019). A unified account of conditioned phonological alternations : Evidence from Guébie. *Language*, *95*(3), 456–497.

Sande, H. (2020). Morphologically conditioned phonology with two triggers. *Phonology*, *37*, 453–493. https://doi.org/10.1017/S0952675720000238

Sande, H., Jenks, P., & Inkelas, S. (2020). Cophonologies by Ph(r)ase. *Natural Language & Linguistic Theory*, *38*, 1211–1261. https://doi.org/10.1007/s11049-020-09467-x

Smolensky, P., & Goldrick, M. (2016). Gradient symbolic representations in grammar: the case of French Liaison. Ms: Johns Hopkins University & Northwestern University. Available as ROA-1286 from the Rutgers Optimality Archive.

Smolensky, P., Goldrick, M., & Mathis, D. (2014). Optimization and quantization in gradient symbol systems: A framework for integrating the continuous and the discrete in cognition. *Cognitive Science*, *38*(6), 1102–1138. https://doi.org/10.1111/cogs.12047

Smolensky, P., & Legendre, G. (2006). *The harmonic mind: from neural computation to optimality-theoretic grammar*. Cambridge, MA: MIT Press.

Smolensky, P., Rosen, E., & Goldrick, M. (2020). Learning a gradient grammar of French liaison. In A. H.-L. Y. Hyunah Baek, Chikako Takahashi (Ed.), *Proceedings of the 2019 Annual Meeting on Phonology*. Washington, D.C.: Linguistic Society of America.

Tessier, A.-M., & Jesney, K. (2021). Learning French liaison with gradient symbolic representations: errors, predictions, consequences. In R. Bennett, R. Bibbs, M. L. Brinkerhoff, M. J. Kaplan, S. Rich, A. Rysling, … M. W. Cavallaro (Eds.), *Proceedings of the 2020 Annual Meeting on Phonology*. Washington, D.C.: Linguistic Society of America.

Vance, T. (2015). Rendaku. In H. Kubozono (Ed.), *Handbook of Japanese phonetics and phonology* (pp. 397–441). Berlin: Mouton de Gruyter.

Walker, R. (2020). Gradient feature activity in Korean place assimilation. In M. Asatryan, Y. Song, & A. Whitmal (Eds.), *NELS 50: Proceedings of the 50th Annual Meeting of the North East Linguistic Society, Volume 3* (pp. 207–220). Amherst, MA: GLSA.

Wilson, H. (1972). *The phonology and syntax of Palauan verb affixes*. Doctoral dissertation, University of Hawaii.

Zimmermann, E. (2017). Gradient symbolic representations in the output: A typology of lexical exceptions. Paper presented at the 48th Meeting of the North East Linguistic Society. Reykyavik: University of Iceland.

Zimmermann, E. (2018a). Being Exceptional is Being Weak : Tonal Exceptions in San Miguel el Grande Mixtec. In S. H. Y. Gillian Gallagher, Maria Gouskova (Ed.), *Proceedings of the 2017 Annual Meeting on Phonology* (pp. 1–12). Washington, D.C.: Linguistic Society of America.

Zimmermann, E. (2018b). Gradient symbolic representations in the output: a case study from Moses Columbian Salishan stress. In S. Hucklebridge & M. Nelson (Eds.), *NELS 48: Proceedings of the 48th Meeting of the North East Linguistic Society* (pp. 275–288.).

Zimmermann, E. (2019). Gradient symbolic representations and the typology of ghost segments. In K. Hout, A. Mai, A. McCollum, S. Rose, & M. Zaslansky (Eds.), *Proceedings of the 2018 Annual Meeting on Phonology* (pp. 1–12). Washington, D.C.: Linguistic Society of

America.

Zimmermann, E. (2020). Gradient symbolic representations and the typology of phonological exceptions. Paper presented at MIT colloquium, February 28, 2020.

Zimmermann, E. (2021a). Faded copies: reduplication as distribution of activity. *Glossa: A Journal of General Linguistics*, *6*(1), 58. 1-38. https://doi.org/10.5334/gjgl.1117

Zimmermann, E. (2021b). Local cooperation between phonological elements instead of non-local grammar adjustment: An argument against Cophonology by Phase theory. Paper presented at the 28th Manchester Phonology Meeting, Manchester, May 28, 2021.

Zuraw, K. (2003). Vowel Reduction in Palauan Reduplicants. In A. Rakowski & N. Richards (Eds.), *Proceedings of the Eighth Meeting of the Austronesian Formal Linguistics Association (AFLA 8)* (pp. 385–398). Cambridge, MA: MIT Working Papers in Linguistics.

Zuraw, K., & Hayes, B. (2017). Intersecting constraint families: An argument for harmonic grammar. *Language*, *93*(3), 497–548. https://doi.org/10.1353/lan.2017.0035