# Learning a gradient grammar of French liaison\*

Paul Smolensky, a,b Eric Rosen, a and Matthew Goldrick Johns Hopkins University, Microsoft Research, Northwestern University

## 1 Introduction

In French, liaison consonants (which we represent here as  $\mathcal{L}$ ) are weak elements that only appear in inter-word contexts ( $W_1 \_ W_2$ ) when  $W_1$  is one of a specific set of words (typically written as ending with a specific liaison consonant) and  $W_2$  is vowel-initial (Côté, 2011, 2008, 2001). However, liaison consonants do not surface when  $W_2$  is an "h-aspiré" word like *héros* 'hero.' These words appear to be vowel-initial but resist liaison and block  $\mathfrak{p}$ -deletion when following *le/la* 'the' or (optionally) *une* 'a' (fem.). h-aspiré words also block resyllabification of  $W_1$  final consonants (e.g., *quel hibou* (kɛl.i.bu)) 'what an owl'). Furthermore, such effects vary across morphemes (Zuraw & Hayes, 2017) and interact with the appearance of feminine and plural morphemes at the  $W_1 \_ W_2$  juncture.

Learning to account for the full range of interacting phenomena presents a significant challenge. Here, we propose the Error-Driven Gradient Activation Readjustment (EDGAR), an error-driven algorithm that can successfully acquire a grammar that accounts for the full range of phenomena reviewed above. This grammar is stated in the Gradient Symbolic Computation (GSC) framework, a type of Harmonic Grammar incorporating partially-activated representations (in addition to real-valued constraint weights). Smolensky & Goldrick (2016) (henceforth S&G) developed a GSC analysis deriving liaison patterns from the coalescence of partially-activated input consonants  $\mathcal{L}_1$  and  $\mathcal{L}_2$  that occur finally in  $W_1$  ( $\mathcal{L}_1$ ) and initially in  $W_2$  ( $\mathcal{L}_2$ ). A liaison consonant  $\mathcal{L}$  surfaces iff its aggregate activation surpasses an epiphenomenal threshold determined by weighted MAX and DEP constraints. The work reported here goes beyond S&G's hand-calculated activation and weight values. EDGAR learns both constraint weights and the activation level of representations, accounting for liaison and its interactions with other processes.

The paper is structured as follows. Following a review of the core empirical patterns in French liaison, we discuss the two primary analysis strategies in the literature (viewing liaison consonants as belonging exclusively to  $W_1$  or  $W_2$ ). We then present the basic structure of the GSC analysis which blends together these two accounts (viewing liaison consonants as associated to both words). We introduce our learning model, EDGAR, and show how the acquired grammar accounts for the empirical data reviewed above.

# 2 Properties of French liaison

Côté (2011) describes French liaison as "a type of external sandhi ...occurring at word boundaries in French ... [which determines ] the pronunciation of a consonant ... between two words ... in certain liaison-triggering contexts." The following examples from Côté illustrate all productive liaison consonants and several syntactic contexts where a liaison consonant (LC) associated with  $W_1$  surfaces:

				Word1	LC	Word2	Context
	a.	un homme	'a/one- MASC man'	[	n	om]	det+noun
(1)	b.	vous allez	'you- PL go-2 PL'	[vu	Z	ale]	pronoun+verb
	c.	grand ami	'big- MASC friend'	[gĸā	t	ami]	adj+noun
	d.	très utile	'very useful'	[tre	Z	ytil]	adv+adj

<sup>\*</sup> Supported in part by NSF BCS1344269 and a sponsored research agreement between Microsoft and Johns Hopkins. Author contributions: PS developed the core analysis in collaboration with MG; ER extended the analysis, and developed and tested EDGAR; ER drafted the manuscript, which was edited by MG & PS. We are grateful for helpful comments from audiences at AMP 2019 and the University of Massachusetts Amherst, the anonymous reviewers, and the Neurosymbolic Computation Group at JHU.

But in addition to syntactic constraints on the occurrence of liaison, we find that it occurs only in certain phonological contexts, as shown in (2):

Liaison-asso	ociated $\mathbf{W}_1$ followed by vowel-initial word	
petit ami:	Liaison consonant surfaces in onset of W	/pətit/ + /ami/ $\rightarrow$ [pə.ti. <b>t</b> a.mi]
Non-liaison	vowel-final $\mathbf{W}_1$ followed by vowel-initial w	ord
joli ami:	No liaison	/ʒoli/ + /ami/ $\rightarrow$ [ʒo.li.a.mi]
Liaison-asso	ociated $\mathbf{W}_1$ followed by h-aspiré $\mathbf{W}_2$	
petit héros:	Liaison consonant does not surface	$\texttt{/patit/+/eso/} \rightarrow [\texttt{pa.ti.eso}]$
Liaison-asso	ociated $\mathbf{W}_1$ followed by consonant-initial $\mathbf{W}$	V <sub>2</sub> .
petit copain:	Liaison consonant does not surface	/pətit/ + /kop $\tilde{\epsilon}$ / $ ightarrow$ [pə.ti.ko.p $\tilde{\epsilon}$ ]

# 3 Two competing analyses of liaison

In the literature we find two competing analyses of liaison: a final- $\mathcal{L}$  analysis (e.g. Tranel (1994)) and a  $\mathcal{L}$ -initial analysis (e.g., Morin (2005)). In the former, a liaison consonant, which occurs underlyingly at the right edge of  $W_1$ , surfaces when, and only when, it is necessary to provide an onset for a following vowel-initial word (satisfying a constraint such as ONSET). This occurs *except* before h-aspiré words.

In a  $\mathcal{L}$ -initial analysis, the liaison consonant is morphologically affiliated with  $W_2$  rather than  $W_1$ . For example, the UR for *ami* 'friend', contains multiple allomorphs: /tami/, /zami/, /nami/, /ami/.  $W_1$  selects the correct allomorph of  $W_2$ .  $W_1$  = petit selects  $/W_2/$  = /tami/.  $W_1$  = les, un, joli respectively select  $/W_2/$  = /zami/, /nami/, /ami/.

If we compare the two approaches, we find that each approach has its own advantages and disadvantages: neither covers all of the data. Here, we briefly review the empirical data; see S&G §3.1-§3.10 for further discussion. (The following phenomena are treated in the full account of which the results presented here are a part,)

- **3.1** Challenges for syllabification-motivated final- $\mathcal{L}$  analyses In final- $\mathcal{L}$  analyses, liaison consonants surface to 'repair' onset-less  $W_2$ . To account for h-aspiré words' exceptionality, this constraint must be weakened. But this fails to explain why  $W_1$ s with fixed final consonants regularly exhibit resyllabification with h-aspiré words (e.g., quel hasard [kɛl.a.zaʁ/kɛ.la.zaʁ] (Tranel, 1995)) suggesting a pressure for onsets to occur even in the initial position of h-aspiré words.
- **3.2** Accounting for  $W_1$  effects in an  $\mathcal{L}$ -initial analysis Allomorph selection for  $W_2$ , e.g., among /tami, nami, zami/, depends on idiosyncratic information in  $W_1$ . If there is no underlying final /t/ in petit 'small, masc.' (as proposed in the  $\mathcal{L}$ -initial analysis), it is not clear how it selects allomorph /tami/ as  $W_2$ . It is also unclear why the  $W_2$  allomorph shares a consonant with the related feminine form petite, where a /t/ surfaces word-finally in citation form.
- **3.3** Liaison is not exclusively driven by properties of a single morpheme A further challenge for either a final- $\mathcal{L}$  or  $\mathcal{L}$ -initial analysis is that when liaison is variable or optional, there is a frequency effect in which the probability that liaison occurs increases with the conditional probability of the  $W_2$  given  $W_1$  (See §8.4 below for references and further discussion.)

For example, Tranel (1994) proposes that exceptionally for h-aspiré words, ONSET is ranked below an ALIGN-LEFT constraint rewarding alignment of the left edge of morphemes with syllable boundaries.

# 4 Fundamentals of the Gradient Symbolic Computation Analysis

We show that Gradient Symbolic Computation (Smolensky & Goldrick, 2016) (henceforth GSC) can effectively blend the two approaches to French liaison:  $\mathcal{L}$ -initial and  $\mathcal{L}$ -final, to achieve a greater degree of explanation than either of them alone. GSC is a novel category of computation — a cognitive architecture that unifies symbolic and neural-network computation. Representations are symbol structures whose components are associated with continuously-varying activation values. Knowledge is represented through weighted constraints, specified by a Harmonic Grammar. This formalism is part of a larger research program in which computation derives outputs from gradient representations in phonology, syntax and semantics (Cho et al., 2017; Faust & Smolensky, 2017; Faust, 2017; Goldrick et al., 2016; Hsu, 2018; Müller, 2017; Rosen, 2016, 2018b,a, 2019; Smolensky et al., 2014; Smolensky & Goldrick, 2016; van Hell et al., 2016; Zimmermann, 2017a,b, 2018).

**4.1** Constraints over partially activated representations GSC, like other versions of Harmonic Grammar, has weighted constraints. For example, the constraint MAX (which provides Harmonic reward for an input that surfaces) could have weight 0.6. Novel to GSC is the use of gradiently activated representations. For example, a consonant /t/ in the underlying representation could have activation 0.48. In addition to consonants, we discuss below gradient morpheme boundaries in the underlying form as well as underlying elements that are pure activation, with no underlying segmental/featural material. (N.b. As the patterns we seek to explain comprise discrete representations, we consider only discrete possible outputs, with activations limited to 1 for present or 0 for absent representational elements.)

Candidates are evaluated for their Harmony, measured through the effect of constraints on possible outputs and on input-output correspondence. These constraints either penalize or reward structure by assigning negative or positive Harmony, respectively.

Constraints considered here that penalize candidates via negative Harmony:

- DEP: assesses a Harmony penalty that is proportional to the activation deficit between output material and its corresponding input material, e.g., input activation of 0.48 on an underlying final liaison /t/ of petit results in a harmonic penalty of  $(1 0.48) \times w_{Dep}$  if the /t/ surfaces.<sup>3</sup>
- UNIFORMITY: penalizes single output segments that correspond to multiple input segments (coalescence). One violation per output segment.<sup>4</sup>
- INTEGRITY: penalizes segments in the input that have multiple correspondences in the output. One violation per input segment.<sup>5</sup>
- ONSET: penalizes onsetless syllables. One violation per syllable.
- \*C.V (Zuraw & Hayes, 2017): penalizes CV sequences that are syllabified as a coda and no onset. One violation per sequence.
- NoCoda: penalizes each syllable with a coda. One violation per syllable.<sup>6</sup>

Constraints considered here that reward candidates via positive Harmony:

<sup>&</sup>lt;sup>2</sup> This reduces to standard DEP in the discrete case where the output segment has activity 1 and there is no corresponding input segment, i.e., input activity is 0.

<sup>&</sup>lt;sup>3</sup> If the surface [t] has only this /t/ as input correspondent. If surface [t] has multiple input correspondents (coalescence), the penalty is  $w_{Dep}$  times 1— the sum of the activations of the input correspondents (up to a maximum sum of 1).

<sup>&</sup>lt;sup>4</sup> We follow S&G in positing that this constraint weakens with increasing frequency of the co-occurrence of  $W_1$  with  $W_2$ . (See S&G §6 for more detailed discussion and an account of 3.3 above.)

<sup>&</sup>lt;sup>5</sup> The violations of UNIFORMITY and INTEGRITY are not scaled by input activation. These constraints refer to the correspondence relations that hold between elements in the input and output, not the input itself. We assume that these correspondence relationships (along with the output) are not gradiently activated. See S&G, §4.2.1 for further discussion.

<sup>&</sup>lt;sup>6</sup> Since a \*C.V violation implies a NOCODA violation but not vice versa, for computational convenience we omit the violation of NOCODA in C.V and include the weight of NOCODA in the learned weight of \*C.V.

- MAX: rewards a candidate in which underlying elements have output correspondents. Reward is proportional to the activity of each underlying element that surfaces; e.g., input activation of 0.48 on an underlying word-final liaison /t/ in *petit* creates a harmonic reward of  $0.48 \times w_{Max}$  when it corresponds to an output segment [t].
- ANCHOR-L/R-C/V-MORPHEME- $\sigma$ : If the left/right edge of a vowel  $V_i$  / consonant  $C_i$  is aligned with the left/right edge of a morpheme  $M_j$  in the input, then the left/right edges of  $V_i/C_i$  and  $M_j$  align with the left/right edge of a syllable in the output. The harmonic reward is  $w_{Anchor}$  times the activation of the underlying material (C/V or morpheme boundaries; as discussed in §8, the strength of each morpheme boundary is determined by the gradient activation of the edge segment(s)).

# 5 The Error-Driven Gradient Activation Readjustment (EDGAR)

To weight these constraints and acquire the activation level of various components of underlying representations, we developed the Error-Driven Gradient Activation Readjustment (EDGAR). This builds on work such as the Gradual Learning Algorithm of Pater & Boersma (2013), but incorporates learning of representation activations in addition to constraint weights. Our dataset comprised a set of examples relevant to the liaison pattern (see below for further discussion, and the Appendix for a comprehensive list of training examples). For each example, the candidate with the greatest Harmony is calculated from input activations and constraint weights. If the wrong winner is chosen, EDGAR strengthens the weights of constraints and activations on segments that favour the desired winner and weakens those favouring the false winner.<sup>7</sup> EDGAR found constraint weights and input activations that derived all the examples correctly. We initialize constraint weights at 0.5 and liaison consonant activations at 0.015. The following pseudocode shows the steps of the algorithm.

```
(3)
         \eta = 0.015
                                                                  ▶ Stepsize for changing activations and weights
         loop
             errors \Leftarrow 0
             for t_k \in training inputs do
                 for t_{kl} \in output candidates for t_k do
                     H_{t_{kl}} = \sum_{c} H_{t_{kl}|c}
                                                          \triangleright Sum the Harmonies resulting from each constraint c.
                 if arg max H_{t_{kl}} has the actual output form then
                      continue
                 else
                      errors \Leftarrow errors + 1
                      Strengthen by stepsize \eta all activations and constraint weights
                           that favour the desired winner.
                      Weaken by stepsize \eta all non-zero activations and constraint weights
                           that favour the false winner.
                  end if
             end for
             if errors == 0 then
                                                                                              ▶ Learning completed.
                  break
             end if
         end loop
```

#### 6 Accounting for basic liaison phenomena

The case of French liaison is an example in which two competing analyses — the liaison consonant belongs exclusively to  $W_1$  vs.  $W_2$  — succeed and fail on distinct sets of phenomena. When a gradient blend

The sign of a weight change depends on the contribution of constraints to Harmony (e.g., strengthening a negative vs. positive constraint means making its weight more negative or more positive, respectively). Activations are treated just like positive constraints. Note as well that EDGAR prevents constraints that reward candidates from having negative weights and constraints that penalize candidates from having positive values.

of structures from the two analyses can capture *all* of the data, we have reason to go beyond classical symbol structures in grammatical theory. As Hankamer  $(1977)^8$  writes: "we must give up the assumption that two or more conflicting analyses cannot be simultaneously correct for a given phenomenon (pp. 583–4) ... such constructions have both analyses at once (in the conjunctive sense)" (p. 592).

**6.1** Analysis overview As shown in the following table (drawing from empirical data reviewed by S&G), we propose here a gradient blend of final- $\mathcal{L}$  and  $\mathcal{L}$ -initial analyses, where  $\mathfrak{t}$  represents a partially activated /t/ at the right edge of  $W_1$  in the input and  $\mathcal{L}_2$  represents a blend of partially-activated /t,z,n/ at the left edge of  $W_2$  in the input.

(4)

		Overview of ir	nput forms			
Word type	Orthog.	Gloss	Input			
		As wor	d 1			
Adj. w. potential liaison petit 'small' (masc.) pətit (t is partially activated)						
Adj. w. fixed final C	juste	ʻjust'	<b>3yst</b> (t is fully activated)			
Vowel-final adj.	joli	'pretty'	30li (no final consonant at all)			
		As wor	rd 2			
Noun w. fixed initial C	copain	'boyfriend'	$kop\tilde{\epsilon}$ (fully activated initial C)			
Noun w. h-aspiré	héros	'hero'	ево (no initial С)			
Noun w. initial vowel	ami	'friend'	$\mathcal{L}_2$ ami ( $\mathcal{L}_2$ = blend of partially-activated /t,z,n/)			

Liaison consonants surface from two underlying positions:  $W_1$ -final and  $W_2$ -initial, with the two partially-activated inputs coalescing in the output. When only one of these underlying positions is present, there is not enough aggregate input activation on that consonant for it to surface.

- (5) petit ami: /pətit/ + /L₂ ami/ → [pə.ti.ta.mi]
   t + L₂ combine their activations through coalescence, surfacing as [t].
  - $\bullet\,$ joli ami: /ʒoli/ + / $\mathcal{L}$ ami/ ightarrow [zo.li.a.mi]
  - petit héros /pətit/ + /eʁo/  $\rightarrow$  [pə.ti.eʁo] (No  $\mathcal{L}_2$  on h-aspiré words.)
    - $\succeq$  no  $\mathcal{L}_2$  The activation of t alone is not enough for a consonant to surface.
  - Similarly with petit copain: /pətit/ + /kop $\tilde{\epsilon}$ /  $\rightarrow$  [pə.ti.ko.p $\tilde{\epsilon}$ ]
- **6.2** Sample optimizations In (6), 9 two liaison consonants on  $W_1$  and  $W_2$  coalesce with enough activation to surface. In candidate (g) the combined activations of coalescing  $\mathcal{L}_1$  and  $\mathcal{L}_2$  result in optimal net Harmony due to faithfulness of 0.34 (MAX reward) 0.25 (DEP penalty). In candidate (a), in which the liaison consonant does not surface, the net Harmonic benefit of 0.06 from Anchor-Left-V-Morpheme- $\sigma$ , combined with –0.72 penalty for an Onset violation, yields lower Harmony than the liaison candidate.

(6)

petit ami	Ĵ-L-V	ONSET	Max	DEP	Ĵ-L-C	*C.V	Ĵ-R-C	Unif	Ĵ-R-V	Н
$\{t, z, n\}$ activation = 0.09, 0.09, 0.09	1.07	-0.72	0.6	-0.58	0.53	-0.47	0.37	-0.04	0.02	
/pəti $(0.48 \cdot t)_1$ / + / $\{t, z, n\}_2$ ami/										
(a) pə.ti.a.mi	0.78	-0.72							0.01	0.07
(b) pə.tit <sub>1</sub> .a.mi	0.78	-0.72	0.29	-0.30		-0.47	0.18			-0.25
(c) pə.ti.t <sub>1</sub> a.mi			0.29	-0.30					0.01	-0.003
(d) pə.tit <sub>2</sub> .a.mi	0.78	-0.72	0.05	-0.53		-0.47				-0.88
(e) pə.ti.t <sub>2</sub> a.mi			0.05	-0.53					0.01	-0.41
(f) pə.tit <sub>12</sub> .a.mi	0.78	-0.72	0.34	-0.25		-0.47	0.18	-0.04		-0.18
(g) 障 pə.ti.t <sub>12</sub> a.mi			0.34	-0.25	0.05			-0.04	0.01	0.11

<sup>&</sup>lt;sup>8</sup> Thanks to Eric Baković for pointing us to this reference.

<sup>&</sup>lt;sup>9</sup> In tableaux, we follow the common practice of omitting the Harmony contributions that are shared by all candidates considered, e.g., the MAX reward from all the stem consonants that surface in all candidates (in (6), /pəti/ and /ami/). Cell entries show Harmony contributions, not degrees of constraint violation/satisfaction.

In (7), a single liaison consonant on  $W_2$  does not have enough activation to surface. Parsing the liaison consonant (in coda as in candidate (b) or onset as in (c)) incurs negative net Harmony with respect to Faithfulness. Because of the partially activated liaison consonant at the left edge of  $/\{t, z, n\}_2$  ami/, there is a gradient morpheme boundary to the left of the /a/ in *ami* which results in candidates (a) and (b) receiving a partial reward from Anchor-L-V. (See §8.2 for further discussion of gradient morpheme boundaries.)

(7)

joli ami	Ĵ-L-V	ONSET	MAX	DEP	Ĵ-L-C	*C.V	Ĵ-R-C	Ĵ-R-V	Н
$\{t, z, n\}$ activation = 0.09, 0.09, 0.09	1.07	-0.72	0.6	-0.58	0.53	-0.47	0.37	0.02	
$/30$ li/ + $/\{t, z, n\}_2$ ami/									
(a) [ 30.li.a.mi	0.78	-0.72						0.02	0.45
(b) ʒo.lit <sub>2</sub> .a.mi	0.78	-0.72	0.05	-0.53		-0.47			-0.75
(c) ʒo.li.t <sub>2</sub> a.mi			0.05	-0.53	0.05			0.02	-0.04

In (8), a single liaison consonant on  $W_1$  does not have enough activation to surface. Parsing it (candidate (b)) would violate NoCoda; this results in a greater Harmonic penalty than the reward from Anchor-Right-C-Morpheme- $\sigma$ , making this less harmonic than the liaison candidate (a).

(8)

petit copain	Ĵ-L-V	ONSET	MAX	DEP	Ĵ-L-C	*C.V	Ĵ-R-C	NoCoda	Ĵ-R-V	Н
No $\{t, z, n\}$ on copain	1.07	-0.72	0.6	-0.58	0.53	-0.47	0.37	-0.34	0.02	
/pəti $(0.48 \cdot t)_1$ / + /kop $\tilde{\epsilon}$ /										
(a) 🎏 pə.ti.ko.pɛ̃					0.53				0.01	0.54
(b) pə.tit <sub>1</sub> .ko.pε̃			0.29	-0.30	0.53		0.18	-0.34	0.01	0.35

In (9), a single liaison consonant on  $W_1$  does not have enough activation to surface as either coda (candidate (b)) or onset (candidate (c)). While parsing the liaison consonant as an onset avoids a penalty from Onset, candidate (a) reaps the full benefit of satisfying constraint Anchor-Left-V-Morpheme- $\sigma$ ; the h-aspiré word has no liaison consonant at its left edge. (See §8 below for further elaboration.)

(9)

petit héros No $\{t, z, n\}$ on héros.	Ĵ-L-V 1.07	ONSET -0.72	MAX 0.6	DEP -0.58	Û-L-C 0.53	*C.V -0.47	Û-R-С 0.37	Ĵ-R-V 0.02	Н
/pəti $(0.48 \cdot t)_1$ / + /ево/									
(a) Derirera	1.07	-0.72						0.01	0.35
(b) pə.tit <sub>1</sub> .e.во	1.07	-0.72	0.29	-0.30		-0.47	0.18		0.04
(c) pə.ti.t <sub>1</sub> ево			0.29	-0.30				0.01	-0.01

In the following sections, we move beyond these core cases to consider liaison's interactions with other aspects of French (morpho)phonology.

# 7 Liaison and other morphemes at the $W_1 = W_2$ juncture

**7.1** Morphemes as pure activation: Masculine/feminine alternations We propose that graded activation of underlying material can exist not only at the level of segments, but also at the level of morphemes (Faust & Smolensky, 2017) and features (Rosen, 2016), so that a morpheme may contain no melodic material at all and yet carry activation. In particular, we propose that in French, the feminine gender morpheme is represented underlyingly by pure activation, whose learned value from the algorithm was 0.35. We use this to account for alternations between masculine and feminine forms such petit/petite, where the liaison /t/ always surfaces in the feminine. The following tableaux show how the learned input activation ensures that the liaison /t/ surfaces in feminine forms even without the support of a following word with an initial  $\mathcal{L}$  consonant.

		Max	DEP	ĴR-С	NoCoda	ĴR-V	Н
		0.6	-0.58	0.37	-0.34	0.02	
	petit 'smal	ll (m.)'	/pəti(0.48	$3 \cdot t)_1$			
(10)	🎏 pə.ti					0.01	0.01
(10)	pə.tit	0.29	30	0.18	-0.34		-0.17
	petite 'smal	ll (f.)'	/pəti(0.48	$\cdot t)_1 / + \phi$	$_{fem}$ $\phi_{fem}$ :	= 0.35	
	pə.ti					0.01	0.01
	🎏 pə.tit	0.50	-0.10	0.18	-0.34		0.24

At the same time, the activation of the feminine morpheme is not sufficient to incorrectly cause an  $\mathcal{L}_2$  consonant to surface in *jolie amie*.

	jolie amie	Ĵ-L-V	ONSET	Max	DEP	Ĵ-L-C	*C.V	Unif	Ĵ-R-V	Н
	$\phi_{fem} = 0.35, t_2 = 0.09$	1.07	-0.72	0.6	-0.58	0.53	-0.47	-0.3	0.02	
(11)	$30\text{li} + \phi_{fem} + \{t, z, n\}_2 \text{ami} + \phi_{fem}$									
(11)	🎏 30.li.a.mi	0.78	-0.72						0.02	0.08
	$30.li.t_{f2}$ a.mi			0.26	-0.32	0.05		-0.3	0.02	-0.29
	$30.lit_{f2}$ .a.mi	0.78	-0.72	0.26	-0.32		-0.47	-0.3		-0.77

**7.2** Competition between plural /z/ and liaison  $\mathcal{L}$  In phrases like petits amis<sup>10</sup>, the plural /z/ surfaces, but not the liaison consonant  $\mathcal{L}$  (Côté, 2011). Even though the algorithm gave the plural morpheme a lower activation than  $\mathcal{L}_1$  liaison consonants, it surfaces in preference to an  $\mathcal{L}_1$  because the plural is rewarded by ANCHOR-LEFT-C-MORPHEME- $\sigma$ . For candidate (a) in (12) below,  $z_{pl}$  is at the left edge of a morpheme in the input and the left edge of a syllable in the output. This cannot be the case for a  $\mathcal{L}_1$  consonant (more strongly activated in the input than an  $\mathcal{L}_2$ ) in candidates (d), (e), or (f), which reap no such reward. Producing both the plural and liaison consonants (candidate (b)) is not optimal as it requires splitting the liaison consonant /t<sub>2</sub>,  $z_2$ ,  $n_2$ / into two output locations <sup>11</sup>.

(12)

petits amis $(z_{pl} = 0.38)$	Ĵ-L-V	ONSET	MAX	DEP	Ĵ-L-C	INTEG	Ĵ-R-C	Unif	Ĵ-R-V	Н
$\{t, z, n\}_2 = 0.09, 0.09, 0.09 \ t_1 = 0.48$	1.07	-0.72	0.6	-0.58	0.53	-0.5	0.37	-0.30	0.02	
$/patit_1/ + /z_{pl}/ + /\{t, z, n\}_2 ami/ + /z_{pl}/$										
(a) <b>I</b> □ pə.ti.z <sub>p,2</sub> a.mi			0.28	-0.30	0.20			0.04	0.01	0.15
(b) pə.tit <sub>1,2</sub> .z <sub>p,2</sub> a.mi			0.62	-0.56	0.20	-0.5	0.18	0.08		-0.85
(c) pə.tit <sub>1,2</sub> .z <sub>p</sub> a.mi			0.57	-0.61	0.20		0.18	0.04		-0.42
(d) pə.ti.t <sub>1,2</sub> a.mi			0.34	-0.25	0.05			0.04	0.01	0.11
(e) pə.ti.t <sub>2</sub> a.mi			0.05	-0.53	0.05				0.01	-0.42
(f) pə.tiz <sub>p,2</sub> .t <sub>2</sub> a.mi			0.34	-0.84	0.20		0.14	0.04	0.01	-0.11
(g) pə.ti.a.mi	0.78	-0.72							0.01	0.07

#### 8 ANCHOR and graded activation

**8.1** *h-aspiré and enchaînement* As seen above, EDGAR found a strong weight to ANCHOR-LEFT-V-MORPHEME- $\sigma$ . By rewarding candidates where the morpheme-initial vowel is also syllable-initial, this constraint accounts for the tendency for h-aspiré words to maintain a syllable boundary at their left edge.

Preservation of syllable boundaries accounts for the lack of "enchaînement" (resyllabification) in word pairs like *quel hasard* 'what chance'. <sup>12</sup> The anchoring constraint rewards maintenance of the syllable boundary at the left edge of *hasard* in spite of the preceding coda consonant.

<sup>&</sup>lt;sup>10</sup> Note that we assume *petit amis* is a frequent phrase and therefore reduce the weight of UNIFORMITY in this tableau. Additionally, for a feminine form *petites amies*, we assume the linear order of morphemes to be /pətit<sub>1</sub>/ + /fem/ + /z<sub>pl</sub>/; the activation of /fem/ combines with t<sub>1</sub> as in the examples reviewed in the preceding section.

<sup>&</sup>lt;sup>11</sup> An additional candidate that would have  $/z_{pl}/$  surface on W<sub>2</sub> is not a contender. While this would earn a 0.14 harmony reward from ANCHOR-RIGHT-C-MORPHEME- $\sigma$ , the net contribution of MAX and DEP would be -0.13, and the violation of NoCoda (not shown in 12), would decrease harmony by -0.34.

<sup>&</sup>lt;sup>12</sup> Côté (2008) claims that enchaînement is optionally possible for these words.

	quel hasard	Ĵ-L-V	ONSET	*C.V	Ĵ-R-C	Н
		1.07	-0.72	-0.47	0.37	
(13)	/kel/ + /azaʁ/					
	l≇ kel.a.zaʁ	1.07	-0.72	-0.47	0.37	0.25
	ke.la.zaʁ					0

**8.2** *h-aspiré* and schwa deletion When a segment at a morpheme edge is partially active (partially present, so to speak), where is the corresponding morpheme boundary? We propose that the boundary is, in effect, split, partially existing on each side of the segment. If the activation of the right-edge segment is a, then there is in effect a (right) boundary with activity a following it, and a (right) boundary with activity a preceding it; mutatis mutandis for the left edge. When, as in the case of liaison, there are a segments with activation a at the left edge, each of these segments is associated with a (left) boundary of activity a to its left; there is then a (left) boundary with activity a to following the a segments; mutatis mutandis for the right edge. Showing the effective boundary activity via superscripts, we have for example:

(14) [pəti].<sup>52</sup>t.<sub>48</sub>].<sup>48</sup> [.09{
$$t, z, n$$
}.<sub>09</sub>[.<sup>73</sup>ami]

The Harmony rewards from ANCHOR constraints are proportional to the activity of the relevant boundary.

The following tableaux illustrate how this accounts for optional schwa deletion for h-aspiré forms such as *une hache*. Similar to the enchaînement cases above, and in contrast to consonant-initial morphemes (see below), ANCHOR-LEFT-V-MORPHEME- $\sigma$  rewards candidates that avoid resyllabification of the preceding consonant (candidates (b) and (c)). The candidate that maintains the schwa (b) is supported by ANCHOR-RIGHT-V-MORPHEME- $\sigma$ , modulated by the input activation of the schwa. Alternatively, if the schwa does not surface (c), there is enough reward from ANCHOR-RIGHT-C-MORPHEME- $\sigma$  (modulated by the gradiently activated morpheme boundary,  $1-a_{schwa}$ ) to syllabify along the morpheme boundary. The extremely close activations of (< 0.03) for these two candidates lead to optional variation between deletion vs. maintenance of the schwa <sup>14</sup> between these forms.

(15)

une hache	Ĵ-L-V	ONSET	MAX	DEP	*C.V	Ĵ-R-C	Unif	Ĵ-R-V	Н
[ $^{M}$ ynə·(0.07)] [ $\phi_{fem} \cdot 0.35$ ] [ $^{M}$ aʃ]	1.07	-0.72	0.6	-0.58	-0.47	0.37	-0.30	0.02	
(a) y.na∫									0
(b) <b>l</b>	1.07	-0.72	0.25	-0.34			04	0.001	0.23
(c) [ <b>宣</b> yn.aʃ	1.07	-0.72			-0.47	0.34			0.22

In contrast, for other lexical items schwa deletion occurs. When  $W_2$  is consonant-initial, the high ranked Anchor-Left-V-Morpheme- $\sigma$  constraint is inactive. This is illustrated in (16); inactive Anchor is omitted from the tableau. Without the \*C.V penalty for (15c), the net negative Harmony contributions from faithfulness to the schwa (Max 0.25-0.34 Dep) now block forms with schwa, (16b–c), from surfacing.

(16)

une souris	ONSET	MAX	DEP	*C.V	Ĵ-R-C	NoCoda	Unif	Ĵ-R-V	Н
$[^{M}$ упэ·(0.07)] $[\phi_{fem} \cdot 0.35] [^{M}$ suві]	-0.72	0.6	-0.58	-0.47	0.37	-0.34	-0.30	0.02	
(a) 🎏 yn.su.ві					0.34	-0.34			0.00
(b) y.nə.su.ʁi		0.25	-0.34				-0.04	0.001	-0.12
(c) yn.ə.su.ʁi	-0.72	0.25	-0.34	-0.47	0.34		-0.04	0.001	-0.96

<sup>&</sup>lt;sup>13</sup> An anonymous reviewer suggests that "the model crucially relies on morphological boundaries to be phonological elements." The reviewer assumes that such boundaries necessarily need to be phonological objects in order to be gradient. Our response is that there is no reason that morphology cannot have gradient properties independent of the phonology. See, for example Hay & Baayen (2005) for an argument in favour of gradient structure in morphology.

<sup>&</sup>lt;sup>14</sup> In the GSC framework, the probability of a candidate is proportional to the exponential of its Harmony. If two candidates differ only slightly in Harmony, their probabilities will be similar and we predict optionality.

**8.3** Gradient morpheme boundaries: Weak h-aspiré and gradient 'alignancy' In online text, when h-aspiré words occur as  $W_2$ , Zuraw & Hayes (2017) observe interacting gradient tendencies towards alignant behaviour (the tendency for morphemes to align with syllable edges). They propose 5 levels on alignant tendency, shown with some examples in (17). The frequencies were taken from comparing the relative numbers of hits on Google for each pair such as *le hiatus* vs. *l'hiatus*<sup>15</sup>.

	Frequency of aligned alternant in hits on Google										
	Word 2 group $\rightarrow$ 1 1 2 3 3 4 5										
(17)		hameau	héros	haricot	hiéroglyphe	hiatus	hameçon	hiver			
		'hamlet'	'hero'	'bean'	'hieroglyph'	'hiatus'	'fishhook'	'winter'			
	le/l'	99.2%	99.7%	94.8%	47%	52%	1.4%	0.1%			

In their group 3 are words such as *hiéroglyphe* and *hiatus* that are intermediate in their behaviour between alignant h-aspiré words and non-alignant words that are orthographically h-initial. Our account can explain this behaviour if these words have some activation of liaison consonants at the left edge but less than the activation of regular vowel-initial words. The following shows how the variable behaviour with *le* could occur for *hiatus*. First, consider tableaux for non-alignant *l'ami* and alignant *le héros*.

	l'ami	Ĵ-L-V	ONSET	Max	DEP	Ĵ-L-C	*C.V	Ĵ-R-C	Ĵ-R-V	Н
	$\{t, z, n\}_2 = \{.09, .09, .09\}$	1.07	-0.72	0.6	-0.58	0.53	-0.47	0.37	0.02	
(18)	$/19.0.25/ + /\{t, z, n\}_2$ ami/									
(10)	🎏 la.mi									0
	lə.a.mi	0.78	-0.72	0.15	-0.44				0.01	-0.22
	lə.na.mi			0.20	-0.96	0.05			0.01	-0.71

	le héros	Ĵ-L-V 1.07	Onset $-0.72$	MAX 0.6	DEP -0.58	Û-L-С 0.53	*C.V -0.47	Û-R-С 0.37	Û-R-V 0.02	Н
(19)	/lə·0.25/ + /eʁo/									
	læ le.ero	1.07	-0.72	0.15	-0.44				0.01	0.07
	је во									0

Because of the underlying liaison consonants on *amis* that we take to be absent in  $h\acute{e}ros$ , allowing the leftmost vowel on  $W_2$  to surface at a left syllable edge does better with respect to (gradiently assessed) left V-anchoring on  $h\acute{e}ros$  than on amis: the morpheme boundary is gradiently present at the left edge of the vowel in amis and fully present in  $h\acute{e}ros$ . Intermediate degrees of alignancy can be readily produced by manipulating the degree of activation of liaison consonants. For example, suppose that the blended liaison consonants in hiatus have lower activations than those in amis. As shown in (20), under these conditions, the Harmony of the candidate in which schwa surfaces (a) is nearly equal to the Harmony of a candidate in which it does not (b), predicting greater variation than observed in the preceding cases.  $^{16}$ 

(20)	le/l' hiatus	Ĵ-L-V	ONSET	MAX	DEP	Ĵ-L-C	*C.V	Ĵ-R-C	Ĵ-R-V	Н
	$\{t, z, n\}_2' = 0.026, 0.026, 0.026$	1.07	-0.72	0.6	-0.58	0.53	-0.47	0.37	0.02	
	$/1 = / \{t, z, n\}_2'$ iatus/									
	(a) ( [章 ) lə.i.a.tus	0.98	-0.72	0.15	-0.43				0.01	0.013
	(b) 🎏 li.a.tus									0
	(c) lə.ti.a.tus			0.17	-1.00	0.02			0.01	-0.82

In summary, the gradient effects of word 2 can be explained by the degree of input activation on the blend of liaison consonants at the left edge. Words whose behaviour is intermediate between canonical h-aspiré and other vowel-initial words will have a higher degree of activation on that blend of consonants than do regular vowel initial words.

<sup>&</sup>lt;sup>15</sup> Because these data depend on written corpora, where word 1 varies orthographically depending on whether word 2 that follow is alignant, we have no way of knowing from their data, or from other web searches, how these words behave with adjectives such as *petit*.

<sup>&</sup>lt;sup>16</sup> Following note 14. Note that the learned activation of schwa on le was different from the activation of schwa on une.

**8.4** Gradient activation and processing-based variability in liaison Graded underlying activation is not only lexically specified: it can also vary across processing contexts. Kilbourn-Ceron (2017b,a) documents variability in liaison in spontaneous and read speech from French, finding that liaison is more likely to occur when  $W_1$  or  $W_2$  is frequent, and when the conditional probability of  $W_2$  given  $W_1$  is high (see also Côté (2011)). She attributes the impact of  $W_2$  frequency and conditional probability to the Locality of Production Planning Hypothesis (see Kilbourn-Ceron for further discussion of  $W_1$  frequency). This hypothesis claims that variation in liaison (and other external sandhi processes) is due in part to the ease of retrieval and encoding of phonological material participating in liaison. For example, if  $W_2$  is hard to retrieve (due to its low frequency or low predictability), speakers may begin planning and speaking  $W_1$  before they have retrieved  $W_2$ 's phonological structure. In this case,  $W_1$  is planned as if  $W_2$  were not present — and therefore unable to trigger liaison. This situation becomes less and less likely if  $W_2$  is easy to retrieve.

In GSC, we can accommodate the effects of  $W_2$  frequency and conditional probability by assuming that the variation in the activation of phonological structure is impacted by retrieval processes. Difficult-to-retrieve structure will have activation that is lower than the maximum value specified in its underlying lexical representation. This processing-based variation will therefore have the same effect found in other cases where there is lexical variation in activation, as discussed previously in this section: less active material will be less likely to trigger liaison.

### 9 Conclusions

In this work, we have introduced EDGAR, a learning algorithm that can successfully find weights and underlying activations specifying a GSC grammar. This learner grammar successfully accounts for complex interactions between liaison and other processes in French morphophonology. Liaison consonants interact with other gradiently activated morphemes at the  $W_1 \_ W_2$  juncture. The graded activation of underlying material, coupled with ANCHOR constraints sensitive to this graded activation, accounts for variation in the alignancy of different morphemes — including variation in enchaînement and schwa deletion.

We see liaison as providing a paradigmatic example in support of a hypothesis concerning a much broader phenomenon in linguistic theory: long-standing impasses in adequately explaining a set of related linguistic phenomena can be resolved through a gradient blend of two seemingly conflicting approaches. While French liaison is a complex phenomenon that we have by no means fully accounted for, the success of this analysis arguably adds weight to the proposal that gradient symbol structures play an important role in linguistic cognition.

## 10 Appendix: range of the examples fed to EDGAR

- 1. petit 'small, (masc.)' SR: pəti<sup>17</sup>
- 2. petite 'small, (fem.)' (pure activation for fem. morpheme in input) SR: pə.tit
- 3. douze 'twelve' (fixed final C) SR: duz
- 4. deux 'two' (liason final /z/) SR: dø
- 5. bon 'good, (masc.)' SR: bõ
- 6. bonne 'good (fem.)' (pure activation for fem. morpheme in input) SR: ban
- 7. ami 'friend (masc.)' SR: a.mi
- 8. petit ami 'boyfriend' SR: pə.ti.ta.mi
- 9. petite hache 'small axe' (/t/ in petite surfaces as coda) SR: po.tit.as
- 10. bon ami 'good friend (masc.)' (liaison /n/ surfaces) SR: bo.na.mi
- 11. deux amis 'two friends' (liaison /z/ surfaces on deux but not plural /z/ in amis SR: dø.za.mi
- 12. joli ami 'pretty friend' ( $\mathcal{L}_2$  doesn't surface on ami) SR: 30.li.a.mi

<sup>&</sup>lt;sup>17</sup> The algorithm also included some examples that were superfluous inasmuch as they repeated conditions in some examples below. These items were: *petit hasard* 'small chance', *deux hasards* 'two chances', *bon hasard* 'good chance', *le héros* 'the hero', *petit chat* 'little cat', *deux chats* 'two cats' and *bon chat* 'good cat'

- 13. *jolie amie* 'pretty friend (fem.)' (Feminine activation on *jolie* unable to make  $\mathcal{L}_2$  surface on *amie*) SR: 30.li.a.mi
- 14. deux copains 'two friends' ( $\mathcal{L}_1$  on deux does not surface before fixed C) SR: dø.ko.p $\tilde{\epsilon}$
- 15. douze copains 'twelve friends' (fully activated C on douze surfaces before fixed C) SR: duz.ko.pẽ
- 16. bon copain 'good friend' ( $\mathcal{L}_1$  on bon does not surface before fixed C) SR: bɔ̃.ko.pɛ̃
- 17. bonne copine 'good friend (fem.)' (Activation of fem. morpheme causes  $\mathcal{L}_1$  to surface) SR: ban.ko.pin
- 18. petit copain 'boyfriend' ( $\mathcal{L}_1$  does not surface before fixed C) SR: pə.ti.ko.p $\tilde{\epsilon}$
- 19. petite copine 'girlfriend' (fem. activation causes  $\mathcal{L}_1$  to surface) SR: pə.tit.ko.pin
- 20. deux héros 'two heroes' ( $\mathcal{L}_1$  does not surface before h-aspiré)а SR: dø.e.во
- 21. petit héros 'little hero' ( $\mathcal{L}_1$  does not surface before h-aspiré) SR: pə.ti.e.во
- 22. bon héros 'good hero' ( $\mathcal{L}_1$  does not surface before h-aspiré) SR: b $\tilde{s}$ .e.во
- 23. joli héros 'pretty hero' (no  $\mathcal{L}_1$  or  $\mathcal{L}_2$  to surface) SR: 30.li.e.во
- 24. le hasard 'the chance' (schwa surfaces before h-apsiré) SR: la.a.zaʁ
- 25. la hache 'the axe' (/a/ on la surfaces before h-aspiré) SR: la.aʃ
- 26. le chat 'the cat' (schwa surfaces before fixed C) SR: lə.[a
- 27. *l'ami* 'the friend' (schwa does not surface before  $\mathcal{L}_2$ ) SR: la.mi
- 28. *l'amie* 'the friend (fem.)' (/a/ on la does not surface before  $\mathcal{L}_2$ ) SR: la.mi
- 29. *une hache* 'an axe' (schwa optionally surfaces before h-aspiré with minimal harmony difference between schwa and non-schwa candidates) SR: y.nə.aʃ or yn.aʃ
- 30. une souris 'a mouse' (schwa does not surface before fixed C) SR; yn.su.ві
- 31. une abeille 'a bee' (schwa does not surface and /n/ syllabifies as onset before  $\mathcal{L}_2$ ) SR: y.na.bej
- 32. quel ami 'what (a) friend' (fixed /l/ syllabifies as onset of W2) SR: ke.la.mi
- 33. *cafés extras* 'extraordinary coffees' (plural /z/ on *cafeś* surfaces as onset; plural /z/ on *extras* does not surface) SR: ka.fe.zɛk.stʁa
- 34. petits chats 'little cats' (plural /z/ on petits does not surface before fixed C) SR: pa.ti.[a
- 35. serait âgé 'would be old' (phrase with low frequency resists liaison through higher weight of UNIFORMITY) SR: sə.ʁɛ.a.ʒe
- 36. momies énormes 'huge mummies' (liason surfaces in phrase with higher frequency) SR: mo.mi.ze.nэвт
- 37. chers amis 'dear friends' (plural /z/ surfaces before V-initial word) SR: [ɛʁ.za.mi
- 38. petits amis 'little friends' (plural /z/ but not liaison /t/ surfaces before V-initial word) SR: pa.ti.za.mi
- 39. chers chats 'dear cats' (plural /z/ does not surface before C-initial word) SR: [EB.[a
- 40. *tamis énorme* 'huge sieve' (liaison C does not surface before V in low-frequency phrase) SR: ta.mi.e.nэв

#### References

- Cho, Pyeong Whan, Matthew Goldrick & Paul Smolensky (2017). Incremental parsing in a continuous dynamical system: sentence processing in Gradient Symbolic Computation. *Linguistics Vanguard* 3, p. 20160105.
- Côté, Marie-Hélène (2001). Consonant cluster phonotactics: a perceptual approach. Tech. Rep. 548.
- Côté, Marie-Hélène (2008). *In Sounds of silence: Empty elements in syntax and phonology*, Amsterdam: Elsevier, chap. Empty elements in schwa, liaison and h-aspiré: The French Holy Trinity revisited, 61–103.
- Côté, Marie-Hélène (2011). The Blackwell Companion to Phonology, Malden: Wiley-Blackwell, chap. French liaison, 2685–2710
- Faust, Noam (2017). How much for that vowel? Talk, Strength in Grammar Workshop, University of Leipzig.
- Faust, Noam & Paul Smolensky (2017). Activity as an alternative to autosegmental association. *Manchester Phonology Meeting*, Manchester, United Kingdom.

Goldrick, Matthew, Michael Putnam & Lara Schwarz (2016). Coactivation in bilingual grammars: A computational account of code mixing. *Bilingualism: Language and Cognition* 19:5, 857–876. ROA 1441.

Hankamer, Jorge (1977). Mechanisms of Syntactic Change, Texas Press, chap. Multiple Analyses, 583-607.

Hay, Jennifer B. & R. Harald Baayen (2005). Shifting paradigms: gradient structure in morphology. Trends in Cognitive Science 9:7, 342–348.

van Hell, Janet G., Clara Cohen & Sarah Grey (2016). Testing tolerance for lexically-specific factors in gradient symbolic computation. *Bilingualism: Language and Cognition* 19:5, 897–899.

Hsu, Brian (2018). Scalar constraints and gradient symbolic representations generate exceptional prosodification effects without exceptional prosody. Handout, West Coast Conference on Formal Linguistics 36.

Kilbourn-Ceron, Oriana (2017a). Speech production planning affects phonological variability: a case study in french liaison. *Proceedings of the Annual Meetings on Phonology*, vol. 4.

Kilbourn-Ceron, Oriana (2017b). Speech production planning affects variation in external sandhi. Ph.D. thesis, McGill University Libraries.

Morin, Yves Charles (2005). La liaison relève-t-elle d'une tendance à eviter les hiatus? réflexions sur son évolution historique. *Langages* 158, 8–23.

Müller, Gereon (2017). Gradient symbolic representations in syntax. Handout, IGRA 02: Syntax II.

Pater, Joe & Paul Boersma (2013). Convergence properties of a gradual learning algorithm for harmonic grammar.

Rosen, Eric (2016). Predicting the unpredictable: Capturing the apparent semi-regularity of rendaku voicing in japanese through gradient symbolic computation. *Proceedings of the Berkeley Linguistics Society*, vol. 42. ROA 1299.

Rosen, Eric (2018a). Evidence for gradient input features from sino-japanese compound accent. Supplemental proceedings of the Annual Meeting on Phonology, vol. 6.

Rosen, Eric (2018b). Predicting semi-regular patterns in morphologically complex words. *Linguistics Vanguard* 4:1. ROA 1339.

Rosen, Eric (2019). Learning complex inflectional paradigms through blended gradient inputs. *Proceedings of the Society for Computation in Linguistics*, vol. 2.

Smolensky, Paul & Matthew Goldrick (2016). Gradient Symbolic Representations in Grammar: The case of French Liaison. Rutgers Optimality Archive 1552, Rutgers University.

Smolensky, Paul, Matthew Goldrick & Donald Mathis (2014). Optimization and Quantization in Gradient Symbol Systems: A Framework for Integrating the Continuous and the Discrete in Cognition. Cognitive Science 38, 1102– 1138.

Tranel, Bernard (1994). French liaison and elision revisited: A unified account within optimality theory. Tech. Rep. 15-0594.

Tranel, Bernard (1995). *The Handbook of Phonological Theory*, Blackwell: Cambridge MA, chap. Current issues in French phonology: liaison and position theories, 798–816.

Zimmermann, Eva (2017a). Being (slightly) stronger: Lexical stress in Moses Columbian Salish. Handout, Strength in Grammar Workshop, University of Leipzig.

Zimmermann, Eva (2017b). Gradient symbols and gradient markedness: A case study from Mixtec tones. Handout: Manchester Phonology Meeting.

Zimmermann, Eva (2018). Gradient Symbolic Representations in the output: A case study from Moses Columbian Salishan stress. *Proceedings of NELS*, vol. 48.

Zuraw, Kie & Bruce Hayes (2017). Intersecting constraint families: An argument for harmonic grammar. Language 93, 497–548.