Learning French liaison with Gradient Symbolic Representations: Errors, Predictions, Consequences

Anne-Michelle Tessier, Karen Jesney University of British Columbia, Carleton University

1 Introduction

Smolensky & Goldrick (2016) first made the case for Gradient Symbolic Representations (GSRs) as the inputs to phonological grammar using the phenomena of French liaison. Liaison involves a rich set of morpho-phonological alternations in which some consonants come and go at word boundaries for reasons that are largely, but not exclusively, phonological (Côté, 2011 provides one comprehensive overview). One of the most controversial aspects of liaison has been the proper analysis of the alternating segments' underlying forms (discussed extensively by e.g. Tranel 1995 and Smith, 2015). Smolensky & Goldrick argue that positing inputs with gradiently-represented segments can resolve much of this controversy and provide a more complete analysis of the phenomenon.

This paper is concerned with the implications of GSRs for phonological learning, and particularly the extent to which a learner can effectively use GSRs to acquire the complex patterns of liaison (see also Hsu 2018; Rosen 2016, 2019; Smolensky, Rosen & Goldrick, 2020). Section 2 lays out the basics of liaison in the most obligatory of contexts – i.e., within noun phrases – and explains how gradiently-activated segments are used in Smolensky & Goldrick's account. Section 3 then describes the types of child production errors known to emerge during the acquisition of liaison, and compares these errors to those made by the GSR learner in our simulations. Section 4 turns to predictions about the end state grammar of this GSR learner, focusing particularly on the treatment of nonce words. We report the results of a pilot online study designed to probe to what extent these predictions are problematic. Finally, in section 5, we touch on the ways in which generalizations about liaison are captured in this approach, both with respect to individual lexical items and across the entire grammar.

2 How does GSR capture patterns like liaison?

2.1 The basic phenomenon: French pre-nominal liaison Liaison occurs in many different morphosyntactic environments with varying degrees of optionality. This study focuses on contexts where liaison is generally obligatory: within the noun phrase, at the juncture between a preceding element (Word1 or W1) and the noun (Word2 or W2). The set of W1s includes determiners, numerals, and a small closed class of frequent pre-nominal adjectives.

Examples are given in (1) for the W1s *les* (plural definite determiner) and the pre-nominal adjective *petit*. When the noun begins with a consonant -[n] in the case of *nuage* or [b] in the case of *bébé* - no liaison consonant appears between Word1 and Word2. When the noun is vowel-initial - as in *ami* or *ours* - however, the liaison consonant appears, surfacing as [z] in the case of *les* and [t] in the case of *petit*.

^{*} For discussion of this work we thank especially Eric Rosen, Matt Goldrick and audiences at UMass Amherst, NYU, UBC, Carleton and AMP2020, particularly Canaan Breiss, Chris Barker, Sam Bowman, Eleanor Glewwe, Gunnar Hansson, Bruce Hayes, Carla Hudson Kam, Doug Pulleyblank and Joe Pater. We thank Anne Bertrand and Roger Lo for crucial help with the study discussed in section 4, and Michael Becker and Starr Sandoval for continuous help with French corpus data. We thank Amanda Rysling, Ryan Bennett and all the AMP organizers at UCSC, and finally we thank the reader for tactfully ignoring all of our remaining errors.

(1) a.	Word1: determiner <i>les</i>				Word1: masculine adjective <i>petit</i>			
	les nuages	[le.ny.ˈaʒ]	'the clouds'		petit nuage	[pœ.ti.ny. 'aʒ]	'little cloud'	
	les bébés	[le.be'be]	'the profs'		petit bébé	[pœ.ti.be. be]	'little baby'	
	les ami	[le.za.ˈmi]	'the friends'		petit ami	[pœ.ti.ta. mi]	'little friend'	
	les ours	[le.ˈzuks]	'the bears'		petit ours	[pœ.ti. 'turs]	'little bear'	

Broadly speaking, liaison occurs when two conditions obtain: W1 comes from a defined set of functional items, and W2 begins with a vowel. In (1a), we see plural nouns combining with a liaison-triggering W1; in (1b) we see these nouns in the masculine singular, with some other morphologically-appropriate W1s. Each W1 is associated with a single liaison consonant – e.g., [z] for *les*, but [t] for *petit* and [n] for the masculine determiner *un*. Vowel-initial W2s, on the other hand, can appear with any liaison consonant; the quality of the consonant is determined by W1.

But not all W1 + W2 combinations result in liaison. In (2a), the [n] of the prenominal adjective *jeune* surfaces regardless of the following segment. In (2b), the prenominal adjective *joli* consistently appears without any liaison consonant, resulting in hiatus when W2 is vowel initial.

```
(2) a. Word1: adjective jeune

jeune bebe [ʒœn.be'be] 'young baby'

jeune ami [ʒœ.na'mi] 'young friend'

b. Word1: adjective joli

joli bébé [ʒɔ.li.be.'be] 'pretty baby'

joli ami [ʒɔ.li.a.'mi] 'pretty friend'
```

These phrases also some crucial aspects of French syllable phonotactics -- V.V hiatus is tolerated like in *joli ami*, as are C.C coda-onset sequences such as in *jeune bébé*. Thus, the liaison observed in (1) cannot be driven simply by syllable structure constraints like NoCoda and Onset.

A final observation is that some nouns systematically resist liaison, despite meeting all the criteria for W2s. Some such words are shown in (3); these words are almost all spelled with an initial 'h', and they are known in the French literature as 'h-aspiré'. Zuraw & Hayes (2017) refer to the broader class of words that resist liaison as 'alignant'.

```
(3) a. Word1: plural determiner les
                                                         b. Word1: masculine adjective petit
         les héros
                       [le'e, Ro]
                                      'the heroes'
                                                             petit héro
                                                                             [bœ.ti.e'ko]
                                                                                               'little hero'
         les hibous
                       [le.i'bu]
                                      'the owls'
                                                             petit hibou
                                                                             [pœ.ti.i. bu]
                                                                                               'little owl'
         les haricots
                      [le.a.ĸiˈko]
                                     'the beans'
                                                             petit haricot
                                                                            [bœ.ti.a.ĸi,ko]
                                                                                              'little bean'
```

2.2 The core analysis: GSR's partially active and blended segments As suggested above, liaison has been extensively described and analyzed in a variety of ways. Much of the variation in analytical approaches stems from the fact that the full liaison pattern involves complications far beyond those directly considered in this paper. In some morphosyntactic environments liaison is optional as opposed to obligatory; between forms of the verb être 'to be' and a following verb, for instance, both pronunciations with a liaison consonant and pronunciations without a liaison consonant are admitted. Thus with the third person plural form sont, we can find both e.g. elles sont allées [ɛl.sɔ̃.a.le] ~ [ɛl.sɔ̃.ta.le] (see especially de Jong, 1990). To follow this example further, different forms of the W1 verb être trigger liaison to different degrees, with the singular third person form est triggering liaison at higher rates than forms like sont; in general more frequent forms trigger liaison at higher rates, but frequency cannot fully explain the lexical effects (Bybee 2001). Zuraw & Hayes (2017) find similar lexical idiosyncracies among the set of Word2's spelled with initial <h>.

From the current paper's perspective, the key point of variation is whether an alternating liaison consonant is underlyingly assigned to Word1 or Word2. To use a concrete example, the phrase *petit ami* – realized as [pœ.ti.ta'mi] – might have any of the following inputs (and perhaps others; see esp. Côté, 2011):

```
    (4) a. Liaison consonant associated with W1
    b. Liaison consonant associated with W2
    c. Liaison consonant absent from the input
    /pœti/+ /tami/
    e.g. Ternes (1977); also §3.1
    e.g. Tranel (1981); Côté (2008)
```

Smolensky & Goldrick (2016) argue that the range of such disagreements over where liaison should be housed, coupled with the fact that all analyses recognize that liaison is in some ways lexically specified,

support the remaining logical possibility: liaison consonants are stored in the inputs of **both** W1 **and** W2. Their analysis, which we adopt here, is couched in the framework of Gradient Symbolic Representations (GSR). In the present case of liaison, it assigns a partially-activated input segment at the end of W1s that trigger liaison, and a blend of partially-activated input segments at the beginning of otherwise V-initial W2s.

As shown in (5a), the W1 *petit* has 0.5 activation of a word-final [t], the W1 *les* has 0.5 activation of a word-final [z], and the W1 *un* has 0.5 activation of word-final [z]. All of the other segments in these words, which do not alternate, have full 1.0 underlying activation. W1s with non-alternating final consonants like *jeune* have a fully activated final segment, and W1s that never trigger liaison like *joli* have no final consonant. Partial activation makes segments more liable to alternate; deleting one is less costly that deleting a fully-activated segment, while realizing one is more costly than realizing a fully-activated segment.

(5) a. Underlying representation of liaison W1s /pœtit0.5/ /lez0.5/ /c̃en0.5/
b. Underlying representation of stable C-final W1s /3œn1.0/
c. Underlying representation of stable V-final W1s /3oli/

Meanwhile, underlying forms of Word2's are represented as having blended word-initial segments. A noun like *ami* has 0.3 activation of **each** of the possible liaison consonants /t/, /n/, and /z/ in the initial position. Forms like *bébé* that do not participate in liaison have full underlying activation of a single initial consonant. Finally, h-aspiré forms that resist liaison simply have no consonant in word-initial position.

(6) a. Underlying representation of W2s that undergo liaison
 b. Underlying representation of stable C-initial W2s
 c. Underlying representation of stable V-initial W2s

Within the GSR framework, these gradient representations form part of a lexical item's input, but their gradience must be resolved when choosing an optimal output – in the winning output candidate, either the final /t/ of *petit* is fully realized, or it is not at ll. The usual constraints of MAX and DEP mediate this segment's behavior, with the degree of violation scaled to the underlying activation of the segment. The underlying form /pœtito.s/, for example, will incur 0.5 violations of MAX if the final [t] is not realized, because a "half-present" segment has been deleted. Conversely, if the final [t] of /pœtito.s/ is realized, 0.5 violations of DEP will be incurred because an addition "half-a-segment" has been added to allow a full [t] to surface.

These inputs are designed to result in liaison only when W1 and W2 combined have a sufficient degree of underlying activation for the **same** consonant. The liaison [t] in the phrase *petit ami* [pœ.ti.ta. 'mi] is a <u>coalescence</u> of the 0.5 activated /t/ from *petit* and the 0.3 activated /t/ from *ami* – producing 0.8 of a faithful /t/ in total, and thus incurring only 0.2 violations of DEP. The basic optimization is sketched in (7).

(7)	/ pœtit _{0.5} / + /{t _{0.3} , n _{0.3} , z _{0.3} }ami/	DEP $ w = -10$	MAX $w = -2$	ONSET $w = -0.9$	Н
	r pœ.ti. t a. 'mi	1.0 - 0.5 - 0.3	0.3 + 0.3		-3.2
		= 0.2	= 0.6		
	pœ.ti.a. 'mi		0.5 + 0.3 + 0.3 + 0.3	1	-3.7
			= 1.4		
	pœ.ti. z a. ˈmi	1.0 - 0.3	0.5 + 0.3 + 0.3		-9.2
		= 0.7	= 1.1		

The first candidate in (7) is the winning output. It incurs the 0.2 violations of DEP needed to fully realize a [t], plus 0.6 violations of MAX because it does not realize the partially-activated $/n_{0.3}/$ and $/z_{0.3}/$ segments of *ami*. Given the weights of the constraints, this is preferable to the second candidate, which fails to realize any of the partially-activated underlying segments, and also violates ONSET. The third candidate, which employs the incorrect liaison segment [z] loses because it incurs worse violations of DEP and MAX.

The case of *petit ami* contrasts with the case of *joli ami*, where liaison never applies. As shown in (8), without underlying activation of a W1-final consonant, the cost of epenthesizing a full segment is too great.

(8)	/ʒɔli/ + /{t₀.3, n₀.3, z₀.3}ami/	DEP $ w = -10$	MAX $w = -2$	ONSET $w = -0.9$	Н
	ʒɔ.li. t a.ˈmi	1.0 - 0.3	0.3 + 0.3		-8.2
	-	= 0.7	= 0.6		
	ு ʒɔ.li.a.ˈmi		0.3 + 0.3 + 0.3	1	-2.7
			= 0.9		

Smolensky & Goldrick (2016) include additional constraints to account for a wide range of data, and formalize a threshold of underlying activation that is needed for a segment to be spelled out under different phonotactic conditions. Given their constraint weights, the activation values in (5) and (6) produce the correct results. The quality of the liaison consonant is determined by the partially-activated segment at the end of W1, but liaison only surfaces if it is also partially activated in the underlying representation of W2.

3 Comparing a GSR learner's liaison patterns with child production data

3.1 Children's acquisition of liaison To evaluate a GSR learner and its ability to capture the development of liaison, we must know something of how French-acquiring children in fact produce liaison and what types of errors they make. Several studies by Chevrot, Dugua and colleagues have reported data from both cross-sectional and longitudinal studies of multiple children learning European French; their data includes results from several elicitation experiments (see also Narby 2003). With respect to liaison accuracy in obligatory contexts, which includes most of the prenominal contexts discussed in section 2, Chevrot, Dabanal & Dugua's (2007) longitudinal study of five children found that 3-4 year olds were 71% accurate on average; Dugua (2006) founds that 4-5 year olds were at 80% accuracy, and at older ages they approached ceiling. During these early years, however, performance can remain variable: Dugua, Chevrot & Côté (2003) found that children omitted liaison 20% of the time at age 2-3, and 14% of the time at ages 5-6, with no significant difference in mean correct liaison across the entire age range. Chevrot, Dabanal and Dugua (2007: 116) also report certain liaison errors which persist in optional contexts as late at age 10-11.

In the W1+W2 context of this paper, there are three main child error types, which we name in (9):

(9)	Error Type	Intended	Child	Comments
	a. Liaison substitution	les arbres [le.zarbr]	[le: u arpr]	V-initial Word2,
		un arbre [œ̃. n aʁbʁ]	[œ.zarpr]	with wrong liaison C
	b. Liaison omission	les arbres [le.zaubu]	[le:arpr]	V-initial Word2,
				failure to realize liaison C
	c. Root substitution	les nuages [le.ny.aʒ]	[le. z y.aʒ]	C-initial Word2,
		un zèbre [æ. z ɛbʁ]	[æ.nɛbʁ]	wrong initial invariant C

At least with respect to the noun phrase contexts, it appears that children's initial bias is to attribute liaison exclusively to W2 – that is, that they misanalyze a liaison C as the first segment of a vowel-initial W2 noun, ignoring its relationship to W1. Thus, children may produce *ami* 'friend' variably as [nami] even in contexts where liaison is not expected, such as utterance-initially. Chevrot & Fayol's (2001) diary study of one child (age 2;1-3;7) reported 41 spontaneous errors in which the child produced utterance-initial Word2's with an unlicensed liaison consonant. In Chevrot, Dugua & Favrol (2009)'s experiment 2 (see also Dugua 2002), children ages 2;4-6;0 were asked to produce nouns in the vocative and thus utterance-initial – e.g. 'Bear, come here!' where *ours* 'bear' is a Word2 liaison target in other contexts. Up at age 4, children produced 'extra' liaison consonants quite commonly (on average 35% of trials ages 2;4-3;1 and 19% of trials ages 3;2-4;1), and they continued to do so occasionally up until age 6. In this task, the added liaison C was almost always [n] (93% of the total errors), but [z], [t], and [l]-initial errors were also found. In diary studies, frequent Word2 nouns are observed with spontaneous productions of liaison substitutions: Chevrot et al. (2007)'s data again show [n] to be the most common liaison substitution, with [z] and [t] also frequent with certain nouns.

4

_

² These 41 errors constitute 6.2% of the total 665 reported liaison errors in this study suggest that misattributing a C to a V-initial W2 is not a uniform or, even exceptionally-frequent, mistake.

As we will see, in our GSR learning simulations, root substitution (9c) has a special status in that it is an error that our learner does **not** produce. One notable aspect of root substitution is that, unlike liaison omissions and liaison substitutions seen more often at ages 2-3, root substitution appears to be a somewhat later error. Chevrot et al.'s (2009) experiment 4 (see also Dugua 2002) revealed that children ages 4-5 were the most likely to produce errors like *les zuages* (instead of *nuages*) and *un nèbre* (rather than *zèbre*): 11% of tokens (40 total) showed root substitution at this age, and 25 of the 50 children in this age range produced at least one such error.

3.2 The GSR input learner The point of departure for our GSR learner is the sketch found in Smolensky & Goldrick's (2016) section 5, which summarizes how a learner could adjust consonant activations at word boundaries in the face of observed data, and ultimately reach a stable, error-free system. Synthesizing both their reasoning and the elaborated child results detailed in the previous section, we assume that W1s are initially stored with no liaison C, and that at least *some* W2s are initially stored with one fully-activated initial C. (This is already probably a considerable idealization, insofar as vowel-initial W2s may be produced with multiple different liaison Cs in the earliest stages of child language acquisition, but we leave this to future work.)

(10) Representations in the initial state of GSR learning

	Types of inputs	Examples
Word1s	liaison triggers, stored with NO final C	les [le], un [@], petit [pæti]
	non liaison triggers, true V-final forms (sanity checks)	joli [ʒəli]
Word2s	liaison triggers, stored with a fully-activated initial C	ami [z _{1.0} ami], ours [n _{1.0} uks]
	liaison triggers, stored with NO initial C	ane [an]
	non-liaison triggers, true C-initial forms (sanity checks)	prof [psof], nuage [nyaʒ]

Again following Smolensky & Goldrick (2016), our phonotactic grammar is formalized as a weighted Harmonic Grammar, so that the optimum is the candidate with the greatest Harmony (Legendre et al. 1990; Smolensky & Legendre 2006; Potts et al. 2010). The particular constraints and weights we use are borrowed directly from Smolensky & Goldrick (2016); these include DEP, MAX, and ONSET, which were used in (7) and (8) above, as well as ALIGN-L and UNIFORMITY.³

(11) A HG grammar of French Liaison – based on Smolensky & Goldrick (2016)

Constraint	Weight	Comments
DEP	-10	prohibits epenthesis; evaluated gradiently (see section 2.2)
Max	-2	prohibits deletion; evaluated gradiently (see section 2.2)
ALIGN-L	-1	prohibits a segment at the left edge of an input morpheme not being
		aligned with the left edge of an output syllable; evaluated categorically
ONSET	-0.9	prohibits onsetless syllables; evaluated categorically
Uniformity	-0.7	prohibits segmental fusion – e.g., mapping input W1 and W2 liaison
		consonants onto a single output segment; evaluated categorically

In the current simulations, we assume that by the time productive liaison is being acquired, the child has already established a phonotactic grammar something like that in (11), presumably through the examination of morphologically-unanalyzed forms (see esp. Hayes 2004). We assume therefore that when the child makes liaison errors, these errors are taken as evidence not of a need to change constraint weights, but instead as motivation to adjust the degree activation on word-edge segments (*cf.* the parallel learning of phonotactic

5

_

³ Smolensky & Goldrick (2016) formalize MAX and ALIGN-L as positive constraints that assign rewards. In our implementation all constraints are negative; the two approaches are equivalent here.

grammar and gradient inputs for liaison in Rosen et al. 2020.)

The learning algorithm we adopt, exemplified in (12), is a simple error-driven update rule, as in the Perceptron learning algorithm (Rosenblatt, 1958) and many other applications to weighted constraint learning (Boersma and Pater, 2016 and references therein). When an error is made, the learner increases the activation of any liaison C found in the target form (the winner) but missing from the learner's current output (the loser). Increases in activation are made to the liaison C both at the end of W1 and at the beginning of W2; the learner is agnostic as to the "proper" affiliation of the liaison consonant, and the increase in activation is split between the two. The learner simultaneously decreases the activation of any liaison C found in the learner's current output (the loser) but not the target output form (the winner), on whichever word(s) it is currently associated. (For discussion of how the learner knows they have in fact made an error, see section 5.)

(12)	Error type	Target output	Learner error	Update
		(winner)	(loser)	
	a. Liaison	deux ours	*[dø. n uʁs]	• Increase /z/ on <i>both</i> W1 & W2 by 0.025
	substitution	[dø. z uʁs]		• Decrease /n/ on W1 by 0.05
	b. Liaison	petit âne	*[pœ.ti.an]	• Increase /t/ on both W1 and W2 by 0.025
	omission	[pœ.ti.tan]		
	c. Liaison	joli ami	*[ʒɔ.li. t ami]	• Decrease /t/ on W2 by 0.05
	insertion	[ʒɔ.li.ami]		-

To study the workings of this learner, we used a series of computational simulations coded in R (R Core Team 2019) to study the change in activation of liaison consonants over time. We are particularly interested in the learner's success in reaching a stable and error-free end state, and the types of errors that are produced along the way. Our key point of comparison is how similar the learner's errors are those observed in the course of child language acquisition as in (9).

In each simulation run, the learner is initialized with a lexicon akin to (10) above. The learner randomly selects a W1 and W2 to combine and submits the combination to their current grammar, which assesses whether or not the current optimal output matches the target. If the output does not match target – i.e., if there is an error – the activation values of the word-initial and word final consonants are updated as described above. Throughout our initial simulations, the maximal update is 0.05, split between W1 and W2 segments where applicable. The maximum activation in any input string position is capped at 1.0. A small amount of noise drawn from a normal distribution (SD = 0.01) is added to the constraint weights on each iteration of Eval (but recall that the underlying constraint weights are not undergoing change or any further learning.)

To map the course of learning, we sample to find the optimal output for each possible W1 + W2 combination after every five pieces of learning data. This sampling provides the developmental trajectories we graph below. After 1000 cycles of learning the process stops – at which point, given the parameter settings of these simulations, the system has either stopped making errors and is thus stable, or (rarely) is in a state of perpetual variability.

3.3 Results 1: Proof of Concept Our first result is, encouragingly, that learning works – that is, our simulations all end up with a final set of lexical representations from which the phonotactic grammar will produce all and only the correct liaison combinations. A representative example of initial and final state input representations is in (13) below. Precise final activation values vary slightly across trials, but the overall pattern is consistent.

(13)	Word1	Initial	Final
	les (-z liaison)	/le/	/le z _{0.65} /
	un (-n liaison)	/ œ /	$/\tilde{\alpha}_{\mathbf{n}_{0.63}}/$
	petit (-t liaison)	/pœti/	/pœti t _{0.75} /
	joli (no liaison)	/ʒɔli/	/ʒəli/

Word2	Initial	Final
ami	/ z 1.0ami/	$/\{\mathbf{z}_{0.65}, \mathbf{n}_{0.22}, \mathbf{t}_{0.13}\}$ ami/
ours	/ n 1.0uks/	$/\{\mathbf{z}_{0.20},\mathbf{n}_{0.71},\mathbf{t}_{0.09}\}$ uks/
âne	/an/	$/\{\mathbf{z}_{0.33}, \mathbf{n}_{0.25}, \mathbf{t}_{0.20}\}$ an/
prof	/brot/	/brot/
nuage	/nyaʒ/	/nyaʒ/
héro	\eRO\	\eRO\

We now demonstrate how these representations result in the correct outputs, for four representative phrasal combinations. First, when both W1 and W2 have some underlying activation of liaison consonants, the correct liaison consonant appears. Looking at the 'initial' column of (13), this is true both when neither word initially had a liaison C in their input (*petit âne*, *un âne*, etc.) and when the W2 initially had the wrong initial liaison C for the context (*les ours*, *un ami*, etc.).⁴

(14) Tableau for *petit âne* – target [t] liaison

$/\text{pætit}_{0.75}^{1} + \{z_{0.33}, n_{0.25}, t_{0.2}\}^{2}$ an	DEP -10	Max -2	ALIGN-L -1	Onset -0.9	Unif -0.7	Н
pœ.ti.an		1.53	1	1		-4.96
pœ.ti. t ¹an	0.25	0.78				-4.06
pœ.ti. t ¹,²an	0.05	0.58			1	-2.36
pœ.ti t¹.z² an	0.92	0.45				-10.10
pœ.ti. z ² an	0.67	1.20				-9.10

(15) Tableau for les ours – target [z] liaison

$/lez_{0.65}^{1} + \{z_{0.20}, n_{0.71}, t_{0.09}\}^{2}uss/$	DEP -10	Max -2	ALIGN-L -1	Onset -0.9	Unif -0.7	Н
le.urs		1.65	1	1		-5.20
le. n ² uʁs	0.29	0.94				-4.78
le.z ^{1,2} uks	0.15	0.80			1	-3.80

Second, no liaison occurs when only one of the words has some underlying activation of a liaison consonant – i.e., no liaison occurs with W1 *joli* nor W2 *nuage*. The latter is shown in (16). In addition, no liaison occurs when W2 is an h-aspiré word like *héro*, illustrated in (17). In both cases, there is insufficient activation of the liaison Cs in the inputs to overcome the high-weighted pressures of DEP.

(16) Tableau for *petit nuage* – target no liaison

/pœtit _{0.75} ¹ + nyaʒ/	Dер -10	Max -2	ALIGN-L -1	Onset -0.9	Unif -0.7	Н
pœ.ti.ny.aʒ		0.75				-1.50
pœ.ti t ¹.ny.aʒ	0.25					-2.50

(17) Tableau for *petit héro* – target no liaison

/pœtit _{0.75} ¹ + еко/	Dер -10	Max -2	ALIGN-L -1	Onset -0.9	Unif -0.7	Н
pœ.ti.e.ko		0.75		1		-2.40
pœ.ti.t¹e.so	0.25		1			-3.50

We note that the final activation values learned here are similar, but not identical, to the hand-calculated values in Smolensky & Goldrick (2016). There, the input activation for W1 liaison consonants was posited to be 0.5, and the input activation for W2 liaison consonants was posited to be 0.3 for each for the possible liaison segments. As they discuss: given a particular set of constraint weights, liaison consonants will surface just when their underlying activation meets some critical *threshold* such that it is more costly to not spell them out than it is to violate other constraints. An error-based learning algorithm like the one we employ here finds a set of activation values that are consistent with this threshold effect, but there is some variability in the precise values learned on each run of the learner.

As for development, we illustrate the trajectories of our two liaison phrase types (from 14 and 15) using

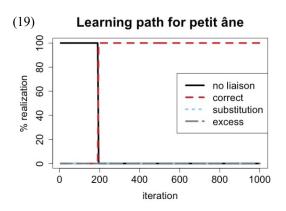
-

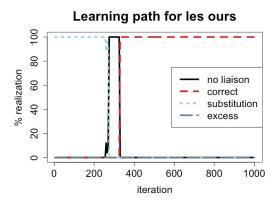
⁴ About notation: As before, the subscript numbers indicate strengths of activation of a segment in the input. In both inputs and outputs, we must also track indices for correspondence relations, as segmental fusion is a crucial output contender in the GSR analysis. Superscripts are used to indicate correspondence.

the proportion of each output type produced when sampled, every five cycles, across the 1000 learning trials. To make the graphs legible we have collapsed all the possible output types into four categories:

(18)		Target: petit âne	Target: les ours	
	No liaison	[pœ.ti.an]	[le.uks]	
	Correct liaison	[pœ.ti.tan]	[le.zurs]	
	Liaison substitution	[pœ.ti.zan], [pœ.ti.nan]	[le.turs], [le.nurs]	
	"Excess" liaison	[pœ.tit.tan], [pœ.tit.nan], etc.	[lez.zurs], [lez.nurs], etc.	

For phrases like *petit âne* – whose W2 was initialized with no word-initial C – the learning curve has just two stages. First we see liaison omission (solid black line). This stage persists until enough activation of the correct liaison C accumulates across both words to emerge faithfully in target-like fashion (dashed red line). No further errors or learning occur.





For phrases like *les ours* – initialized with an initial W2 consonant that is wrong for combining with this W1 – the initial stage is liaison substitution (dotted blue line). This initial error occurs every time that the W2 follows any non-compatible W1, and so this W2's initial activation of /n/ decreases relatively rapidly. Activation of the alternative (correct) consonants increases relatively more slowly, because they are favoured just in a limited set of liaison environments. The learner therefore passes through a stage where the optimal output is liaison omission (solid black line). Eventually, however, the joint activation of the correct liaison C across W1 and W2 is sufficient for it to faithfully emerge in target-like fashion (dashed red line). At that point, errors and learning cease. (Excess liaison does not appear in these simulations; see next section.)

With respect to the no-liaison contexts in this simulation, the h-aspiré phrases are stable throughout, being consistently realized in a target-like fashion with no liaison. This is of course understandable because all such phrases begin with no liaison C activation on either morpheme, and no evidence suggests that any such C should be added to the h-aspiré W2. As seen in tableaux (16) and (17), the grammar is such that if there is no activation of the liaison consonant on either W1 or W2, liaison is not realized (but see section 4.)

3.4 Results 2: Relative Frequency and Overshoot Errors In the simulations above, each two-word phrase is equally likely to be attempted by the learner, and since there are roughly the same number of W1 and W2 words, errors caused the activation of liaison consonants to increase and decrease at roughly the same rate in each individual word. In a real-world learning context, however, these assumptions must be wrong. In particular, while the set of W2s includes all vowel-initial nouns – a large and open lexical class – W1s are a much smaller closed class of mostly-functional items, including determiners, numerals and a special sub-class of adjectives which can appear prenominally. This means that, in practice, a few W1s appear particularly frequently in liaison contexts, while each individual W2 is relatively less frequent. Since the extent of learner's changes to segment activations are tied directly to the number and nature of individual errors, we were interested in examining how more normal lexical frequencies influence our results.

In order to mimic the frequency asymmetry mentioned above, we tried our same simulations again but with an asymmetric change to the update rule from (12). In these new simulations, the maximum update for

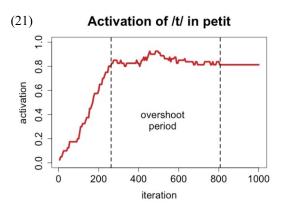
W1 segments is 0.05 per learning trial and the maximal update for W2 segments is 0.0125 per learning trial (i.e., one quarter the update of W1 segments).

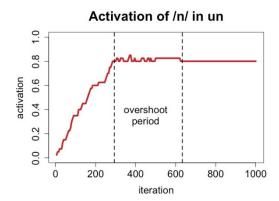
Overall, this change has fairly small and unsurprising quantitative effects on the end state, exemplified in (20). Consonant activations rise much faster on W1s than W2s, based on the same errors. More activation ends up being attributed to the W1s, while the final activation values of W2s remain closer to the initial hypothesis – e.g., more [z] on *ami* if it was initially stored as $/z_{1.0}$ ami/ In addition, however, W1 forms often acquire a small amount of activation of the "incorrect" liaison consonant – e.g., the 0.05 activation of /n on les (cf. 13 in section 3.3). These small amounts of activation emerge as learner tries to account for the presence of [n] in phrases like *petit nuage* where the [n] is stably associated with W2.

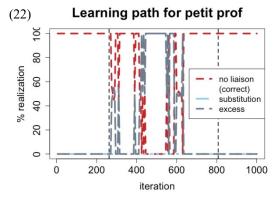
(20)	Word1	Initial	Final
	les (-z liaison)	/le/	/le{z _{0.80} , n _{0.05} }/
	un (-n liaison)	/œ/	$/\tilde{lpha}_{ m n_{0.80}}/$
	petit (-t liaison)	/pœti/	/pœti{t _{0.82} , n _{0.14} }/
	<i>joli</i> (no liaison)	/ʒɔli/	/30li/

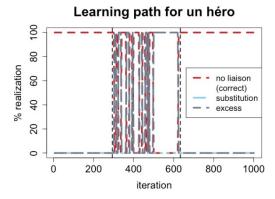
Word2	Initial	Final
ami	/ z 1.0ami/	$/\{z_{0.76}, n_{0.09}, t_{0.07}\}$ ami/
ours	/ n 1.0uks/	$\{z_{0.09}, n_{0.75}, t_{0.11}\}$ uks/
âne	/an/	$/\{z_{0.09}, n_{0.08}, t_{0.08}\}$ an/
prof	/brot/	\brot\
nuage	/nyaʒ/	/nyaʒ/
héro	\ero\	\eRO\

Although still ultimately successful, the learning trajectories in this set of simulations are somewhat bumpier. They also include a new type of error; here the learner commonly passes through a stage where they attribute *too much* liaison consonant activation to one of the W1s. This period of what we call liaison 'overshoot' is observable in (21) for the word-final /t in *petit* and the word final /n of un. In this particular simulation the period lasts for more iterations in the case of *petit* than in the case of un, but with both forms there is a clear timespan where the level of underlying activation is variably greater than the eventual final activation level. (A similar effect is observed with the /z of les.)









During the overshoot period, we see errors in which liaison consonants from the W1 appear in non-target contexts – i.e., excess liaison. This includes liaison clusters, either with a C-initial W2 – e.g., petit prof realized as [pœ.tit.pʁof] – or a V-initial W2 where the initial-state liaison C has not yet been reduced enough – e.g., petit ami realized as [pœ.tit.zami]. It also includes overapplication of liaison with h-aspiré words – e.g., un héro as realized as [æ.ne.ʁo]. Two examples are shown in (22), with the relevant overshoot periods delimited. Note that the error rate during this period is variable, reflecting the ongoing changes in the precise activation values of the W1 consonants.

3.5 Summarizing child vs. GSR learning results The table in (23) compares the child production errors discussed in section 3.1 to the results of the simulations reported in the subsequent sections. Note that patterns (23d-f) are all instances of excess liaison, which to our knowledge is not attested in descriptions of liaison in the child production literature (though see section 5).

(23)	Eri	or Type	Intended	Child	Comments	Simulation?
	a.	Liaison	les arbres	[le:uarpr]	V-initial W2, realized	Yes – given W2
		substitution	un arbre	[@.zarpr]	with wrong liaison C	initiated with a
						liaison C
	b.	Liaison	les arbres	[le:arpr]	V-initial W2, failure to	Yes – in all
		omission			realize liaison C	simulations
	c.	Root	les nuages	[le. z y.aʒ]	C-initial W2, realized	No
		substitution	un zèbre	[œ. n ɛbʁ]	with wrong initial C	
	d.	Regularized	un héro	$[\tilde{\mathbf{w}}.\mathbf{n}e.$ ro]	h-aspiré V-initial W2,	Yes – given
		h-aspiré		attested?	realized with liaison C	asymmetric updates
	e.	Liaison	petit prof	[bœ.ti t.p ʁɔf]	C-initial W2, realized	Yes – given
		clusters v.1		attested?	with liaison C anyway	asymmetric updates
	f.	Liaison	petit ours	[pœtit.nuss]	V-initial W2, realized	Yes – given
		clusters v.2		attested?	with two liaison Cs	asymmetric updates

The only child error that is not found in the GSR simulations is root substitution. As previously mentioned, this tends to be a somewhat later error in child productions so it is possible that additional factors not incorporated into the current simulations could help to explain this. More promisingly, however, liaison substitution and omission errors, consistent with those produced by children, do naturally emerge given error-driven learning of gradient underlying activations.

The GSR simulations can also easily produce liaison clusters (21e,f) and h-aspiré regularization (21d), whose status in the child production data remains unclear to us. To better understand this latter error, we now turn to the broader predictions of the GSR account for French h-aspiré W2s, and to a test of these predictions in adults.

4 Comparing a GSR learner's end-state with adult nonce word data

4.1 A simple but bad prediction about nonce words The learner described in section 3 reaches a stable end state when it stops making errors on combinations of observed W1 and W2s. As it stands, however, this result is restricted to observed lexical items; the learner does not reliably extend liaison to nonce forms.

The general question of how an alternation triggered by a particular subset of lexical items ought to be applied to novel forms is not entirely clear (see, for example, the arguments and experimental results of Pierrehumbert 2006 regarding English velar softening). In the present scenario, the liaison status of each W1 must be learned on a word-by-word basis; *joli* ends with no final liaison consonant while *petit* has a liaison /t/, and this can only be determined by seeing these W1s concatenated with V-initial W2s. In the case of W2s, there are two types of V-initial forms that must again be learned through observation, seeing them together with liaison-triggering W1s. Although there may be tendencies regarding which types of 'V-initial' W2s belong to the h-aspiré class (see esp. Zuraw & Hayes 2017, and also below), it is certainly the case that the majority of eligible W2 words do participate in liaison. Furthermore, anecdotal intuitions from native speakers of all French dialects are that, in the absence of evidence to the contrary, novel V-initial words should participate in liaison – i.e., if they are not known to be spelled with an initial <h>.

There is thus a disconnect between intuitions and the learner's predictions regarding French nonce words. If a learner has heard the phrase "un joli abrola', where 'abrola' is a novel word, the intuition is that it should undergo liaison in other familiar contexts like un abrola, petit abrola and so on. However, simply feeding /abrola/ as an input to the GSR grammar will not result in liaison, even in contexts where the preceding W1 is an established liaison trigger. Without partially-activated liaison consonants, any new word will be treated as h-aspiré – compare the identical violation profiles from tableau (17) for petit héro.

(24) Tableau for nonce petit abrola – failure of liaison

/pœtit _{0.75} ¹ + abʁola/	Dер -10	Max -2	ALIGN-L -1	Onset -0.9	Unif -0.7	Н
рœ.ti.a.bко.la		0.75		1		-2.40
рœ.ti. t ¹a.bкo.la	0.25		1			-3.50

This prediction did not go unnoticed by Smolensky & Goldrick (2016). In their section 7: page 31, they suggest that the French learner must at some point generalize across W2s and posit an abstract blended initial segment that is attributed to novel V-initial W2s. Choosing the activation levels for this blended segment is not actually trivial, however. For our learners in section 3, different initial assumptions about a W2's liaison consonant – i.e., which one it begins with, if any at all – lead to different end-state blends of liaison Cs, all of which are consistent with the target data. In any event, the question of how novel nouns do, in fact, behave in liaison contexts is not one that has been experimentally addressed to our knowledge. We therefore report here on the early results of a pilot study that attempts to gather some quantitative data on this topic.

4.2 A liaison wug test Our online pilot study was primarily designed to elicit nonce word productions in liaison contexts, with a small set of real words in liaison contexts included at the beginning to establish a baseline. Each trial had two parts, exemplified in (25) and (26) below. First, a sentence appeared on the screen with one word missing, indicated with a blank line. At the same time, the participant heard the sentence read aloud, with the blank filled by the trial's target word. All of these heard sentences had the target W2 in a non-liaison context, after *joli*, *jeune* or six – which invariantly end in a vowel, a [n] and a [s], respectively.

(25) Trial Part 1

a.	real word:	Seen:	Voici un jeune	Here is a young
		Heard:	"Voici un jeune [elefa]"	Here is a young elephant
b.	nonce word:	Seen:	Ceci est un joli	This is a pretty
		Heard:	"Ceci est un joli [akidu]"	This is a pretty "akidu"

Then a new sentence appeared on the screen, again with one missing word indicated with a blank line. When ready, the participant pressed a button to begin recording, and then read the sentence aloud, inserting the target word that had filled the blank on the previous screen. In this second set of sentences, the target word was always preceded by one of three W1 liaison triggers: the singular indefinite determiner un, the plural indefinite determiner des or the adjective petit.

(26) Trial Part 2

a.	real word:	Seen:	C'est un	énorme	It's a	enormous
		Produced:	"C'est [œnelefa] é	enorme"		
b.	nonce word:	Seen:	Ils sont des	_ imaginaires	They are some	imaginary
		Produced:	"Ils sont [???] im	aginaires"		

The crucial experimental question was whether participants would produce liaison or not in (26b): [de.a.ki.du] or [de.za.ki.du].

In the initial real word block there were 9 trials, including vowel-initial, consonant-initial, and h-aspiré nouns (see table 27 below for details). Participants who produced liaison correctly for all the three V-initial nonce words but none of the C-initial words had their data included in the study.

In the nonce word block there were 42 trials, including 27 critical V-initial items and 15 C-initial distractors; a sample is provided below in table (27). The critical items were split between monosyllabic

forms (14 words) and trisyllabic forms (13 words), while fillers were a mix of 1, 2 and 3 syllables. We chose to manipulate syllable count in the critical items following work by Zuraw & Hayes (2017) that found short words are more likely to resist liaison, though in a different set of contexts. Preliminary results from our own ongoing corpus work using the Micro Robert dictionary (Rey 1989) are consistent with this; among V-initial nouns we find monosyllables to be most resistant to liaison, disyllables less so, and longer words considerably less so.

(27)		Critical Ite	Filler Items (n=15)		
	1 syl	lable	3 syll	lables	1, 2 and 3 syllables
	[im]	[ibl]	[itikuʃ]	[øtinam]	[bikn]e]
	[Apr]	[yd]	[hpremi]	[aʃeli]	[uipr]
	[øʒ]	[eqr]	[ebraln]]	[aspusin]	[tymɛt]
	[atr]	[uʃ]	[œĸani]	[olakiz]	[lykaʁi]

For the pilot study, participants were adults age 18+ who identified as native speakers of Canadian French who use and speak the language daily. They were recruited via word of mouth and participated on a voluntary basis. For a participant's data to be included in the analysis reported below, they must have completed all trials, and passed the audio/headphones test before the experiment began. These criteria, along with technical difficulties with some participants' recordings, narrowed the dataset to 8 usable participants.

4.3 *Pilot results* To code the data, two anglophone undergraduate students who had some familiarity with the task and French phonology listened to each recording, and coded (i) whether the nonce word had been produced accurately, (ii) whether they heard any amount of a liaison consonant, and (iii) whether they heard any glottal stop in lieu of or preceding a liaison consonant. In cases of uncertainty, coders examined spectrograms for any evidence of consonantal closure; when in doubt they were instructed to err on the side of reporting a liaison consonant.⁵

The overall picture observed from the pilot data is reported in (28) and (29) below. Table (28) gives the liaison behavior for each participant on the nine real words. All 8 usable participants met the criteria of producing liaison on all V-initial words and no C-initial words; they also all produced liaison with $h\hat{o}tel$. On the two h-aspiré words, however, behavior was as mixed as possible, such that across participants some produced liaison on neither word, some produced it on one or the other one, and some produced it on both⁶. As these two words were chosen to be common words invariably reported as h-aspiré in texts and dictionaries, it may be that liaison exceptionality is currently undergoing language change among this population of Canadian French speakers. To better understand these patterns, we are now collecting data from more speakers with a larger range of allegedly h-aspiré targets.

(28) Rates of liaison for real W2 words

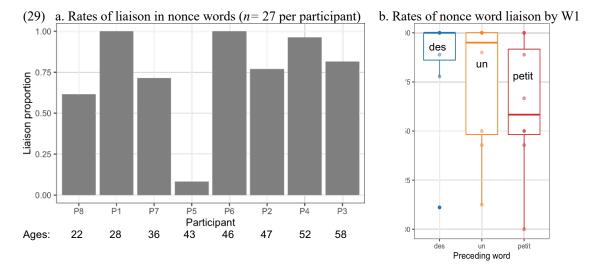
Word Type	Liaison rate
V-initial: ami, éléphant, ours	24/24
C-initial: cactus, telephone, zèbre	0/24
V-initial, though spelled h-initial: <i>hôtel</i>	0/8
h-aspiré 1: <i>héro</i>	4/7 (see footnote 6)
h-aspiré 2: hibou	4/7

With respect to nonce words, recall that our main question was how many V-initial W2s would be realized with liaison following a W1 trigger, when participants had only heard the word in a non-liaison context (and not seen it spelled). The figure in (29a) shows the proportion of each participant's 27 V-initial W2s that were produced with liaison. The overall result is that most participants applied liaison, most of the time. Three (P1, P6 and P4) used liaison nearly always; three others (P8, P7, P2 and P3) used it frequently (between 60% and 80% of trials), and one participant P5 avoided it almost entirely. (We note too that P5's productions were frequently coded with a glottal stop between W1 and W2, which might have been a strategy

⁵ This bias was chosen to work against the GSR 'hypothesis' that nonce words ought **not** to trigger liaison.

⁶ Due to technical difficulties one participant's h-aspiré productions were not recorded.

adopted for this particular task or else might reflect something more consequential about their grammar. We leave this point for future investigation in a larger study.)



The figure in (29b) breaks down the rate of liaison for the three W1s, each of which trigger a different liaison consonant. One notable result is that [z]-liaison, triggered by the indefinite plural determiner des, is the most frequently applied; apart from P5, the other seven participants produced des + nonce noun with liaison more than 75% of the time. Liaison of [n] triggered by W1 un was more variably produced, in that two additional participants (beyond P5) produced liaison on 50% or fewer of un + nonce noun trials. Most variable was [t]-liaison with W1 petit, which runs the gamut from 0-100% liaison depending on the participant. This result accords with typical descriptions of liaison; with respect to real words it is usually reported that determiners like des and un show obligatory liaison before W2 vowel-initial nouns, while liaison with adjectives like petit is highly frequent yet still optional.

Before leaving the pilot data we note a few additional things about our results. First: we found no clear connection between a speaker's age and their rates of liaison (see the participants ages noted in 29a). Almost all of our speakers were born and raised in Montreal, so it seems that dialect differences cannot be explanatory either. We also did not find any effect of word length – if anything, monosyllables were *less* resistant to liaison than the three-syllable wors, although this trend is not significant and much more data would be necessary to make a claim in either direction.

In sum, the two major takeaways from our preliminary experiment are that (i) as expected, liaison can be freely applied to V-initial novel nouns, but also that (ii) liaison is *not* fully automatic in this context, with the additional caveat that the status of h-aspiré words in current Canadian French might also be in flux. Future work is therefore necessary both to establish generalizations about how and when speakers apply liaison productively, and to formalize the mechanism by which a GSR learner would generalize their W2 liaison representations onto novel V-initial words.

5 Discussion and Conclusions

This paper has demonstrated that, given an established phonotactic grammar, an error-driven GSR learner can acquire underlying segmental activations that capture the basics of French liaison. The types of errors that the learner makes in the course of learning are largely similar to those that are made by French-acquiring children. Importantly, however, the learner's end state is one where liaison is *not* applied to novel words. With respect to novel W2 nouns, this seems at odds both with intuition and the overall results of our pilot production study, although participants' performance was somewhat variable.

Looking toward future research, a number of interesting predictions arise from the the GSR approach to modeling liaison when combined with an error-driven learner to acquire input activation values. In the current simulations, combinations of W1s and W2s were sampled randomly, ignoring any asymmetries in collocation frequencies. But surely not all W2s are equally likely to appear with every W1. As an illustration, preliminary

searches of French child-directed speech in the Carbajal, Bouchon, Dupoux & Peperkamp derived corpus on CHILDES (MacWhinney 2000) find that the noun éléphant appears predominantly in singular contexts (un éléphant: 33 times vs. les éléphants 9 times), while the noun oreille (ear) appears overwhelmingly in the plural (un oreille zero times, vs. les oreilles 65 times), In setting up our learner's initial state, we already assume that such especially frequent collocations might result in an initial misparse, e.g. the initial state UR of /nl.ouß/ for ours. However, other more subtle effects are also expected. A W2 element that primarily appears in the singular – following un, and several other [n]-triggering W1 like the singular possessive determiners mon, son and ton – will, in the early stages of learning, have more errors that increase the activation of the initial /n/ than the initial /t/ or /z/. Thus, correct liaison with e.g. éléphant should be mastered more quickly with un, mon, ton and son than with petit, les, and so on. All else being equal, a W2 that primarily appears in the plural will display the opposite pattern, with oreille's liaison produced reliably first with les, des and other [z]-triggering W1s. In our ongoing work we are examining these types of predictions in more detail (Jesney, Tessier, Sandoval and Lo, in prep.)

The importance of collocations also raises important questions about how, exactly, the learner accesses the target form of a phrase. Some models of liaison development (e.g., stages 1 and 2 in Chevrot et al. 2009) store full phrases, but the GSR model discussed here assumes that W1s and W2s are represented separately. In order to assess errors and adjust activations values, however, the learner must have some way of knowing what liaison consonant (if any) particular W1 + W2 combinations are supposed to trigger, regardless of what their current grammar produces. One possibility is that learning occurs just when a specific collocation is encountered in child-directed speech, but the results of our pilot study suggest that liaison is fairly generalized and abstracted (at least by adulthood) and that direct experience with a given collocation is not necessary for speakers to apply liaison.

More generally, though, the variability between our pilot study speakers could potentially be understood within the GSR learning framework. It is possible for learners, depending upon their initial hypotheses and the relative frequency of different forms they encounter, to land on different end-state representations of W1 and W2 items – compare (13) and (20). To the extent that individuals' learned grammars of liaison differ, we might reasonably expect their treatment of exceptional (h-aspiré) and nonce words to differ as well. At the same time, whatever aspects of liaison are determined by constraints within the grammar itself should be reflected in the knowledge and behavior of *all* speakers. One such possible aspect is the relatively greater resistance to liaison in shorter W2s discussed in the literature (see especially Zuraw and Hayes, 2017's appendix). This arguably reflects an increased grammatical pressure of ALIGN-L or CRISPEDGE within the head foot. A larger nonce word production study that we are currently developing will examine these effects and the extent of variability in purported h-aspiré words in greater detail.

Finally, our ongoing corpus work is also collecting additional data about the types of liaison input that children actually receive and the types of errors that they actually make in the course of acquisition. As alluded to above, we have already found (unsurprisingly) that different W2s appears with different degrees of frequency in different liaison contexts. In order to accurately assess the success of the GSR + learning model, it is of course necessary to understand what types of data to feed into the learner. We also suspect that other developmental error types reported in the child production literature may relate to the liaison learning process. As noted in section 3, the current model does not generate root substitution errors, and it is as yet unclear what adjustment to the learner or the input data could create this error pattern. At the same time, our learner does generate instances of excess liaison, which are not directly reported in the literature. One possibility is that, in the case of forms like [pətit.pʁɔf], some child errors have been "mis-described" as errors in gender marking⁸ – but it is important to establish to what extent such errors actually arise in child speech, and whether they extend to h-aspiré target words, as they do in our simulations. In general, a clearer picture of child behaviour can inform both the initial state representations of this type of learner as well as the assessment of its learning success.

References

Boersma, Paul and Joe Pater. 2016. Convergence properties of a Gradual Learning Algorithm for Harmonic Grammar.

⁷ Thanks to Chris Barker for raising this issue.

⁸ Thanks to Gunnar Hansson for suggesting this avenue of research.

- In John J. McCarthy and Joe Pater (eds.) *Harmonic Grammar and Harmonic Serialism*, Sheffield: Equinox. 389-434.
- Bybee, Joan. 2001. Frequency effects on French Liaison. In Joan Bybee and Paul Hopper (eds.), Frequency and the Emergence of Linguistic Structure, Amsterdam: John Benjamins. 337–359.
- Chevrot, Jean-Pierre, Damien Dabanal and Céline Dugua. 2007. Pour un modèle de l'acquisition des liaisons basé sur l'usage: trois études de cas. *French Language Studies* 17: 103-128.
- Chevrot, Jean-Pierre, Céline Dugua and Michel Fayal. 2009. Liaison acquisition, word segmentation and construction in French: a usage-based account. *Journal of Child Language* 36: 557-596.
- Chevrot, Jean-Pierre. and Michel Fayol. 2001. Acquisition of French liaison and related child errors. In M. Amgren, A. Barreña, M. J. Ezeizabarrena, I. Idiazabal, and B. MacWhinney (eds.) *Research on Child Language Acquisition* 2. Somerville, MA: Cascadilla Press. 760–774.
- Côté, Marie-Hélène. 2008. Empty elements in schwa, liaison and h-aspiré: The French Holy Trinity revisited. In Jutta M. Hartmann, Veronika Hegedüs & Henk van Riemsdijk (eds.), Sounds of silence: Empty elements in syntax and phonology, Amsterdam: Elsevier. 61-103.
- Côté, Marie-Hélène. 2011. French liaison. In: Marc van Oostendorp, Colin Ewen, Elizabeth Hume & Keren Rice (eds.), *The Blackwell Companion to Phonology*. Malden, MA: Wiley-Blackwell. 2685–2710.
- de Jong, Daan. 1990. The syntax-phonology interface and French liaison. *Linguistics* 28, 57–88.
- Dugua, Céline. 2002. Liaison et segmentation du lexique en français: vers un scénario développemental. Diplôme d'Etudes Approfondies (MA thesis equivalent), Université Grenoble 3.
- Dugua, Céline. 2006. Liaison, segmentation lexicale et schémas syntaxiques entre 2 et 6 ans. Un modèle développemental basé sur l'usage. Doctoral thesis, University of Stendhal Grenoble 3.
- Dugua, Céline, Jean-Pierre Chevrot and Marie-Hélène Côté. 2003. Liaison et formation des mots: scenario développemental et consequence pour le traitement phonologique. Ms., Montpellier: Université Paul-Valéry.
- Legendre, Géraldine, Yoshiro Miyata and Paul Smolensky. 1990. Harmonic Grammar—a formal multi-level connectionist theory of linguistic well-formedness: An application. *Proceedings of the Cognitive Science Society* 12
- Hayes, Bruce. 2004. Phonological acquisition: the early stages. In René Kager, Joe Pater and Wim Zonneveld (eds.) *Fixing Priorities: Constraints in Phonological Acquisition*. Cambridge: CUP. 158-203.
- Jesney. Karen, Roger Lo, Starr Sandoval and Anne-Michelle Tessier. (in prep.) Error-driven GSR learning of French liaison: insights from lexical frequency and child-directed speech. Msc, UBC and University of Ottawa.
- Narby, A 2003. Production et jugement d'acceptabilité entre 2 et 6 ans: aspects psycholinguistiques et sociolinguistiques de l'acquisition des liaisons. Diplôme d'Etudes Approfondies (MA thesis equivalent), Université Grenoble 3.
- Pierrehumbert, Janet. 2006. The statistical basis of an unnatural alternation. Laboratory Phonology 8: 81-106.
- Potts, Christopher, Joe Pater, Karen Jesney, Rajesh Bhatt and Michael Becker. 2010. Harmonic Grammar with linear programming: from linear systems to linguistic typology. *Phonology* 27: 77-117.
- Rey, Alain (ed.). 1989. Le Micro-Robert Poche. Paris: Dictionnaires le Robert.
- Rosen, Eric. 2016. Predicting the unpredictable: Capturing the apparent semi-regularity of rendaku voicing in Japanese through harmonic grammar. In Emily Clem, et al (eds.) *Proceedings of the Berkeley Linguistics Society* 42: 235-249
- Rosen, Eric. 2019. Learning complex inflectional paradigms through blended gradient inputs. In Gaja Jarosz, Max Nelson, Brendan O'Connor and Joe Pater (eds.) *Proceedings of the Society for Computation in Linguistics*, vol. 2. https://doi.org/10.7275/ccwf-j606
- Selkirk, Elisabeth O. 1974. French liaison and the X-bar notation. *Linguistic Inquiry* 5: 573-590.
- Smith, Brian W. 2015. Phonologically-conditioned allomorphy and UR constraints. Doctoral thesis, University of Massachusetts Amherst.
- Smolensky, Paul & Geraldine Legendre. 2006. *The Harmonic Mind: From Neural Computation to Optimality-Theoretic Grammar*. Cambridge, MA: MIT Press.
- Smolensky, Paul & Matthew Goldrick. 2016. *Gradient Symbolic Representations in Grammar: The case of French Liaison*. Rutgers Optimality Archive 1552, Rutgers University.
- Smolensky, Paul, Eric Rosen & Matthew Goldrick. 2020. Learning a gradient grammar of French liaison. In Hyunah Baek, Chikako Takahashi, Alex Hong-Lun Yeung (eds.) *Proceedings of the Annual Meeting of Phonology* 2019. https://doi.org/10.3765/amp.v8i0.4680
- Tranel, Bernard, 1981. Concreteness in Generative Phonology: Evidence from French. Berkeley: University of California Press.
- Tranel, Bernard. 1995. Current issues in French phonology: liaison and position theories. In *Handbook of Phonological Theory*, Cambridge, MA: Blackwell. 798–816.
- Zuraw, Kie and Bruce Hayes. 2017. Intersecting constraint families: An argument for Harmonic Grammar. Language 93(3): 497-548.