



Received 6 January 2020; revised 24 July 2020; accepted 26 September 2020. © 2020 Gillian Gallagher.
Published by the Linguistic Society of America with permission of the authors under a [CC BY 3.0 license](#).

Strident harmony from the perspective of an inductive learner

Gillian Gallagher*

New York University – gillian@nyu.edu

In languages with strident harmony, stridents within a particular domain are required to have the same minor place of articulation. Harmony is often required only of stridents within a root or stem morpheme, and doesn't trigger alternations. Harmony is also often quite local, applying exclusively or more strongly between stridents in the same or adjacent syllables. Finally, harmony may be morpheme specific, triggering alternations in some affixes but not others. All of these specifics of a given harmony pattern give rise to exceptions to harmony at the level of the word, and may require a morphologically parsed learning corpus in order to be acquired. This paper explores the learnability of strident harmony in text corpora from three languages: Nkore-Kiga (Bantu), Papantla Totonac (Totonacan) and Navajo (Athapaskan). The analyses show that word level exceptions largely obscure the harmony pattern as an overall phonotactic in a language. The three languages also serve as a test of the Projection Induction Learner (Gouskova & Gallagher 2020), which is found to be successful when the generalizations in the data are strong but may fail in the face of patterned exceptions.

Keywords: strident harmony; phonotactic learning; maximum entropy; constraint induction

1 Introduction

Minor place harmony among strident consonants is the most common type of nonlocal consonant interaction, and has been the topic of several surveys and lines of analysis (Shaw 1991; Gafos 1999; Hansson 2001). This paper looks at strident harmony from the perspective of an inductive learner, evaluating the statistical evidence for harmony in language corpora. The starting point is two macro-generalizations that emerge when looking through the existing surveys. First, strident harmony is often described as being obligatory only between stridents in adjacent syllables, and second, harmony is often morphologically specific, either holding only of roots/stems or triggering alternations in only a subset of affixes. Both of these factors raise the question of whether strident harmony is observable to a learner as a general phonotactic over words, or if a more elaborated morphological or phonological representation must be acquired before the harmony pattern can be identified.

I investigate the learnability of strident harmony through a quantitative assessment of strident cooccurrence in dictionary and web corpora for three languages from three different regions and families – Nkore-Kiga (Bantu), Papantla Totonac (Totonacan) and Navajo (Na-Dene) – and by testing the ability of

* I am grateful to Maria Gouskova for comments on this paper, as well as three anonymous reviewers and PDA editor Megan Crowhurst. Will Bennett provided assistance in interpreting the Nkore-Kiga literature and data; Alicia Chatten was instrumental in putting together the corpora of Navajo and in summarizing the typology of strident harmony systems; Ildikó Emese Szabó assisted in analyzing cooccurrence statistics in the corpora. This work was supported by NSF BCS-1724753.

the Inductive Projection Learner (Gouskova & Gallagher 2020; Gallagher et al. 2019) to learn the nonlocal harmony pattern from the distribution of stridents in these corpora. These three languages were chosen because they have sufficient materials available for study and because strident harmony manifests differently in each language.

The Inductive Projection Learner is based on the UCLA Phonotactic Learner (Hayes & Wilson 2008), which constructs a grammar of n -gram constraints that fits the distribution of natural classes in a set of learning data. The Inductive Projection Learner takes the output of this learning procedure and looks for evidence of nonlocal phonotactics in the form of **placeholder trigram** constraints. If the grammar contains a constraint of the form X-any_segment-Y (where X and Y are natural classes, either the same class or two distinct classes), this is interpreted as evidence that X and Y may interact nonlocally in the language. To test for nonlocal interactions, the model constructs a nonlocal projection that includes both X and Y and then builds a final grammar that includes constraints on this nonlocal projection. The learning simulations included here serve to assess the evidence for strident harmony in the learning data as well as the viability of this learning procedure itself.

The case studies address four main points. First, existing descriptions of harmony patterns are compared to the distribution of stridents in the corpora, uncovering both exceptions to harmony and asymmetries based on feature value (anterior...posterior sequences have different attestation than posterior...anterior sequences).

Second, the strength of a harmony preference at the word level is examined in languages with morphologically sensitive harmony (Totonac and Navajo). In both languages, harmony is observable at the trigram level, that is, between fairly local stridents, but is not substantiated as a general phonotactic over all stridents in a word. The number of exceptions at the word level are too numerous.

Third, the generality of the Inductive Projection Learner is assessed by testing its ability to induce a strident projection based on the distribution of stridents within a trigram. The finding is that generally the model is successful, though the settings of the learner need to be calibrated in a specific way to achieve the desired result. The scenarios where the model fails highlight the difficulties in distinguishing between systematic and accidental gaps on statistical grounds alone.

Fourth, and finally, harmony patterns in word corpora from two different sources are compared: the web corpora available on An Crúbadán (<http://crubadan.org/>) and word lists extracted from dictionaries. For both languages for which this comparison is available (Totonac and Navajo), the strident harmony pattern is stronger in the dictionary data than the web data.

The paper is organized as follows. The typology of strident harmony is briefly discussed in Section 2 and the Inductive Projection Learner is introduced in Section 3. The case studies of Nkore-Kiga, Papantla Totonac and Navajo are presented in Sections 4, 5 and 6. Section 7 provides a general discussion and concludes.

2 An overview of strident harmony

Hansson's (2001) dissertation surveys dozens of cases of consonant harmony, of which strident harmony is the most frequent. Minor place harmony among stridents is identified in 44 languages in 15 language families. Harmony, however, is rarely if ever a general requirement of words in a language; instead, it is limited by morphological or locality factors. Strident harmony is often quite local, holding only between transvocalic stridents or stridents in adjacent syllables. In many languages, strident harmony at longer distances either doesn't hold at all or is optional. Harmony is often also morphologically sensitive, applying only to roots or stems but not triggering alternations, or applying to specific affixes only. Furthermore, some sources mention that harmony is variable, or exists as a tendency in the language.

Based on the impressionistic descriptions, the question is raised whether strident harmony is observable as a general phonotactic in a language, or whether the distribution of stridents and the limitations on harmony obscure the pattern in the language as a whole. The three case studies in this paper aim to elucidate

the learning problem for strident harmony by quantifying harmony in three languages. Each of these languages shows a slightly different type of pattern.

Nkore-Kiga is described as having strident harmony within stems, with supporting alternations triggered by morphological mutations that change the anteriority of a stem final strident (e.g., [-ʃa:ʃa] ~ [sa:s-ire] ‘be in pain’). Nkore-Kiga is the subject of Bennett & Pulleyblank’s (2018) insightful squib, and is documented in Taylor’s (1959) dictionary and (1985) grammar as well as Morris & Kirwan (1957). Harmony in the language is featurally asymmetrical, in that anterior-posterior combinations are nearly absent while posterior-anterior combinations of stridents are attested. Harmony also holds more strongly among transvocalic stridents than at longer distances.

The Totonacan languages show strident harmony as a root cooccurrence restriction only (e.g., [ʃufʃut] ‘water’, *[tʃufʃut]). Within roots, stridents agree for anteriority, but no alternations are triggered in affixes so harmony does not hold as a general phonotactic of words. There is both a dictionary (Aschmann & Aschmann 1973) and an An Crúbadán web corpus of the Papantla variety of Totonac; the phonology of Mistantla Totonac, including a short discussion of strident harmony, is described in MacKay (1999). Because roots are generally quite short, Totonac can also be described as a case where strident harmony holds more strongly in adjacent syllables than within words as a whole.

Finally, Navajo is one of the most oft-cited cases of strident harmony. Harmony applies within the root, as well as to certain prefixes within the stem (the “conjunction” prefixes), but does not generally hold within words (e.g., [si-sází] ‘my ancestor’ but [ʃó-joo-s-t’e] ‘he has acquired it’). The language has been given many descriptive and phonological treatments (Sapir & Hoijer 1967; McDonough 2003; Hansson 2001; Martin 2011), and is documented in Young and Morgan’s dictionaries and grammars (1972, 1987) as well as an An Crúbadán web corpus. Harmony is described as being required among adjacent syllables, but optional at further distances.

3 The Inductive Projection Learner

3.1 Description of the model

The Inductive Projection Learner (Gouskova & Gallagher 2020; Gallagher et al. 2019) builds on the UCLA Phonotactic Learner (Hayes & Wilson 2008), implementing an inductive procedure for postulating nonlocal projections.¹ The UCLA Phonotactic Learner takes as input a list of attested word forms in a language (the training data) and a feature set that uniquely defines each segment in the list. The model then constructs a phonotactic grammar based on natural class n -gram constraints and assigns these constraints weights using a Maximum Entropy procedure, with the goal of maximizing the likelihood of the observed data. The resulting grammar can be tested and evaluated based on the scores it assigns to nonce forms (Daland et al. 2011; Berent et al. 2012; Hayes & White 2013).

The UCLA Phonotactic Learner allows the analyst to define nonlocal projections to handle nonlocal phonological interactions. For example, Hayes & Wilson (2008) show that vowel harmony in Shona can be accounted for with bigram constraints on a projection that includes only [+syllabic] sounds, mimicking an autosegmental analysis. While the original version of the learner requires the analyst to stipulate what projections the grammar should include, the Inductive Projection Learner zeroes in on properties of the baseline grammar with no projections to diagnose what projections may be useful in a given language. The key insight is that languages that have nonlocal phonological dependencies may show those dependencies within a trigram, and relevant trigrams may be used by the learner as evidence to induce a nonlocal projection. To see how this works, consider the example of a nonlocal laryngeal restriction in Quechua, shown in (1). Ejectives may appear either initially or medially in a root, but crucially may not co-occur, regardless of the number of intervening segments.

¹ The learner is available at https://github.com/gouskova/inductive_projection_learner.

(1) a. p'afa 'clothing'
 t'uλu 'bone'

 b. rit'i 'snow'
 wajk'uj 'to cook'

 c. *p'at'a
 *p'ant'a
 *p'amit'a

A typical analysis of this phenomenon would postulate a constraint against [+constricted glottis][+constricted glottis] ([+cg][+cg]) bigrams, which holds on a nonlocal projection that includes only ejective consonants. The nonlocal projection allows the number of intervening segments to be irrelevant to the evaluation of the constraint, as all of the ungrammatical forms in (1c) contain ejective-ejective bigrams at this level of representation, as shown in (2).

(2) baseline projection: p' a t' a p' a n t' a p' a m i t' a
 p' t' p' t' p' t'

When the UCLA Phonotactic Learner is run on Quechua data without any stipulated nonlocal projections (that is, just the ‘baseline’ or linear string), it includes a trigram constraint of the form * [+cg] [] [+cg], where [] stands for ‘any segment’. This is the crucial aspect of nonlocal phonology that the Inductive Projection Learner is built on. While nonlocal phonological dependencies are characterized by allowing an arbitrary amount of intervening material, the restricted segments are also restricted within a trigram across a single segment, and this trigram-level dependency can be observed without a nonlocal projection. A ‘placeholder trigram’ X-[]-Y is a signal that the natural classes X and Y interact nonlocally, since the identity of the intervening segment is irrelevant.

The Inductive Projection Learner works as follows. First, a baseline grammar is built based on just the linear string with no nonlocal projections and this grammar is evaluated for placeholder trigrams of the form X-[]-Y, where X and Y are natural classes. The model then postulates a nonlocal projection for each placeholder trigram constraint, and a final grammar is built that includes the baseline projection as well as any motivated nonlocal projections. The nonlocal projection is defined as the smallest natural class that includes all of the segments in X and Y. For example, in Quechua, the * [+cg] [] [+cg] constraint motivates a [+cg] projection. In a language with height harmony between vowels a placeholder trigram constraint like * [+high] [] [-high] would motivate a [+syllabic] projection, since [+syllabic] (the class of all vowels) includes all the segments in the [+high] (the class of high vowels) and [-high] (the class of non-high vowels) classes.² Previous work has shown that laryngeal restrictions in Quechua and Aymara and vowel-vowel cooccurrence restrictions in Shona can be successfully learned through this inductive procedure (Gouskova and Gallagher 2020).

3.2 Methodology for building and assessing grammars

One of the goals of the current paper is to test the generality of the placeholder trigram approach to learning nonlocal projections by applying it to more languages, namely, those with strident harmony. The UCLA Phonotactic Learner does not learn every possible constraint on unattested or underattested combinations, so there is no guarantee that every language with a nonlocal phonological dependency will include a corresponding placeholder trigram constraint in the baseline grammar. In some cases, the learner might miss a placeholder trigram entirely either because there are too many exceptions, or because the target constraint holds over small natural classes and thus doesn’t substantially improve the fit of the grammar.

² This assumes that glides are not specified for [+high]; if they are, then the smallest natural class to build from the constraint [+high] [] [-high] would be [-consonantal], the class of all glides and vowels. The projection that is motivated by a particular placeholder trigram constraint is particular to the natural class structure of that language.

The learner could also include a more general constraint that targets larger natural classes, often with more exceptions than the target placeholder trigram (e.g., a constraint on all [+anterior] consonants, as opposed to just the [+strident, +anterior] ones). In other cases, the learner might include a constraint that is more specific. For example, instead of a placeholder trigram with [] ‘any segment’ as the medial gram, the learner might include a trigram constraint where the medial gram is a more specific natural class, like [-low] or [-high], if this results in a more accurate constraint due to the distribution of exceptions to harmony.

What constraints are included in the grammar is affected by several parameters that determine the number of constraints induced, the length of the segmental strings that they scope over, and how tightly the grammar is fit to the data. For all of the models presented in this paper, the number of constraints to be induced is set to a high threshold so that the grammars include as many constraints as meet the other criteria (e.g., a model is allowed to find up to 100 constraints but routinely returns 60-70). Therefore, the size of the constraint set, as determined by the analyst, is not playing a role in the results presented here. The length of the n -gram constraints on the baseline projection was 1-3, and constraints on nonlocal projections could be bigrams or trigrams.

The two parameters that were actively manipulated in the reported simulations are **gain** and **gamma**. The gain parameter (Della Pietra et al. 1997; Wilson & Gallagher 2018) replaces the O/E threshold in the original version of the learner proposed in Hayes & Wilson (2008). A constraint’s gain is proportional to the reduction in the Kullback-Leibler divergence between the data and the current grammar compared to the grammar with the constraint added, with the weights of all other constraints unchanged. The gain of a constraint is higher when the probability distributions in the training data are close to those generated by the grammar if the constraint is added to the grammar. The higher the gain parameter, the harder it is to add new constraints to the grammar, so grammars with higher gain tend to have fewer constraints.

The gamma parameter affects how the objective function of the learner is calculated with each new constraint. It scales the harmony score of the constraint relative to the negative log probability, with the effect that the impact of constraint violations by individual forms is increased. Higher gamma results in a grammar with fewer low-weighted constraints (usually, those with many violations observed in the training data), often favoring more specific and accurate constraints.

The gain and gamma parameters are chosen by the analyst and have a tremendous impact on the resulting grammar. At this point, it is unknown what parameters best approximate how an actual human learner assesses phonotactic constraints, so the appropriate gain and gamma for a given data set cannot be determined in advance. For all of the simulations presented in this paper, many gain and gamma combinations were tried and the results from a ‘best fit’ grammar are presented. Procedurally, the ‘best’ gain and gamma combination was found by beginning with a gamma of 0 and a low gain of 5, and then gradually raising gain until the baseline grammar was fairly small (100 constraints or fewer). Once the gain setting was returning a moderately sized grammar, gamma was slowly increased. Gain and gamma were then tweaked together to assess the range of values at which strident harmony could be captured, and where and how the model failed when it did. In addition to reporting the best-fit grammar, the range of successful gain and gamma combinations are summarized for each data set (or the range of unsuccessful settings that were attempted, in cases where strident harmony was not captured by the model).

The outcome of the learning procedure was assessed through three metrics. First, the baseline grammar was searched for a placeholder trigram constraint that could motivate a strident projection. Second, the strident projection was searched for constraints that favored harmonic strident combinations over disharmonic combinations. Third, the scores that the final grammar assigns to a large set of nonce words were compared, to see if the grammar prefers harmonic to disharmonic stridents in a general way.

A grammar that captures strident harmony will have different average scores between harmonic and disharmonic nonce words, but will also have a different range of scores such that harmonic words are consistently preferred to disharmonic words, regardless of other phonotactic structures. The grammars learned through this model of constraint induction often include many constraints, some of which may penalize accidental gaps, or be constraints that capture the relative frequency of licit phonological

structures. For a grammar to capture strident harmony, the penalty for disharmonic stridents must be greater than the relative range of penalties assigned to otherwise well formed nonce words.

4 Case study 1: Nkore-Kiga

Nkore-Kiga has four strident fricatives [s zʃʒ], which are described as harmonizing for minor place within the stem (Hansson 2001; Bennett & Pulleyblank 2018). Prefixes do not participate, and the affricate stridents [tsʃ] are not reported to participate either. Prior descriptions are based on alternations, as shown in (3) and (4) (tones and morpheme boundaries marked as in the source). In their thorough analysis of alternations, Bennett & Pulleyblank explain that the root alternations are partially governed by morphology. Harmonic alternations are triggered when a suffix requires a stem-final strident to shift from anterior to postalveolar.

(3)	Alternations from Hansson (2001)
	<i>stem</i> <i>perfective</i>
	-ʃa:ʃa ‘be in pain’ sa:s-ire
	-ʃíʃa ‘compensate’ sis-ire
	-ʃinʒa ‘testify against’ sinz-ire

(4)	Alternations from Bennett & Pulleyblank (2018)
	<i>stem</i> <i>derived</i>
	-ʃíʃa ‘be fat’ -sisire
	-ʃè:ʃa ‘pull down (house’ -òmù-sè:si-
	-ʃà:ʃa ‘be in pain’ -òbù-sà:si

The Morris & Kirwan (1957) grammar also provides a few examples of harmonic alternations with the causative suffix /-isa/, shown in (5).

(5)	Alternations from Morris & Kirwan (1957)
	<i>stem</i> <i>derived</i>
	-tagaiʒa ‘walk painfully’ -tagaizisa
	-raʃa ‘shoot’ -rasisa

The three suffixes that predominantly trigger this alternation are the perfective suffix /-ire/, the agentive nominalizer /-i/, and the ‘short’ causative extension suffix /-j/. Hansson further states that the harmony pattern is featurally asymmetrical and distance sensitive: anterior-postalveolar combinations are allowed at longer distances (more than a single V(n) intervening between the two stridents), while postalveolar-anterior combinations are prohibited at any distance.

To substantiate these previous descriptions and explore the distributional and alternation-based evidence available to an inductive learner, a corpus was constructed from Taylor’s (1959) dictionary (available in electronic form at <http://www.cbold.ish-lyon.cnrs.fr/>).

4.1 The corpus

The electronic form of Taylor’s (1959) dictionary contains 12,574 entries, with stems marked. Once entries with odd characters (e.g., ?, @, #, -) were removed, there were 12,147 stems remaining for analysis. Forms with spaces or dashes were included, with the space or dash removed, and segmentally equivalent stems were retained as separate entries (there were 6,849 unique stem types). Taylor (1985) describes the orthographic sequences <ky> and <gy> as palatalized velars in Nkore and as palatoalveolar affricates in Kiga; these sequences were transcribed as clusters of a velar followed by a palatal glide in the corpus. Otherwise, the transcription followed the transparent grapheme-to-phoneme correspondence described by Taylor.

4.2 Alternations

To find alternating forms, the list of stems was searched for items with two stridents. Morphological relatedness was then assessed based on phonological form and meaning. The examples in (6) show all of the plausibly morphologically related forms with alternating anteriority in stridents that were found in the dictionary corpus. Each lettered sub-heading groups together forms with the same putative root, and gives one example of each unique stem form that appears with that root (the same stem may appear with multiple prefixes).

(6)	a. ekj-obusa:si kw-esa:sa kw-esa:sira ku-sa:sa obu-sa:si ku-sa:sibwa oku-sa:sira oku-sa:sirwa	‘unfortunately’ ‘hurt oneself’ ‘pity oneself’ ‘hurt’ ‘trouble, pain’ ‘need’ ‘pity’ ‘pity’	ku-ʃa:ʃa	‘to be in pain’
	b. kw-esi:sibanisa	‘massage oneself’	kw-eʃi:ʃa	‘spoil oneself’
	c. ebi-rikusinza ku-sinza aba-sinzi o-sinzire	‘intoxicants’ ‘intoxicate’ ‘drunkards’ ‘drunk’	eki-ʃinʒo	‘rage’
	d. ku-sangunuza	‘collect’	ku-ʃanguruʒa	‘collect’
	e. ku-se:sa	‘spend the whole night’	obu-ʃe:ʃe	‘early morning’
	f. omu-se:si eki-se:sero	‘emptier, spiller’ ‘urinal’	ku-ʃe:ʃa ku-ʃe:ʃeka	‘pull down (a house)’ ‘be poured out’
	g. obu-si:si	‘sin, evil’	oku-ʃi:ʃa	‘spoil, do wrong’
	h. ku-za:za omu-za:zi	‘cause birth, give rise to’ ‘obstetrician (for cattle)’	aba-ʒa:ʒabi	‘nurses’
	i. aba-sigisi	‘stirrers of porridge’	ku-ʃigisa	‘stir (millet)’
	j. eki-segesi	‘porcupine’	eki-ʃegjeʃu	‘blood from cupping’

While the examples in (6) demonstrate that harmonic alternations in anteriority are attested, it is not always clear what determines the anteriority value in a given word. Moreover, harmony is not exceptionless even in the cases where it is better understood. In particular, Bennett & Pulleyblank report on pairs like [-ʃà:ʃà] ~ [sa:s-ire] ‘be in pain’, where the perfective suffix /-ire/ triggers an alternation in the stem-final strident, which then in turn triggers harmony in the stem-initial strident (c.f. *[ʃa:s-ire] where only the stem final strident alternates, resulting in a disharmonic form). My search of the dictionary revealed, however, that an alternation in the stem final strident does not always result in harmony. Forms were extracted from the dictionary that contained two stridents and either the nominalizing /-i/ or perfective /-ire/ suffix (the causative /-j/ suffix may or may not be overtly present in the output form, making it difficult to assess whether it is underlyingly present, so forms of this sort were not sought out). Of these, 4/5 /-ire/ forms show agreement for minor place among the stridents, and 31/63 /-i/ forms do. Some of the non-harmonic forms are shown in (7).

(7)	àbà-jànzìrè	‘wife’s sisters’ husbands’	(expected àbà-sànzìrè)
	àb-è:jàzì	‘cattle-waterers’	(expected àb-è:sàzì)
	òmù-jàzì	‘madman’	(expected òmù-sàzì)

These disharmonic forms also contradict Hansson’s claim that harmony is obligatory within adjacent syllables, since these examples involve disagreeing stridents separated by a single vowel or Vn sequence. A search through the dictionary did confirm that there are no alternations involving the affricates, as previously described. While strident agreement is thus certainly attested in the lexicon, and seems to be supported by alternations, it is by no means exceptionless. The next section looks at the overall distributional support for harmony in the dictionary corpus.

4.3 Phonotactics

While many forms in the dictionary contain agreeing stridents, there are also many forms with multiple stridents that disagree in anteriority, as shown in the previous section. The pattern appears to be featurally asymmetrical: anterior-postalveolar combinations are nearly unattested, but postalveolar-anterior combinations are frequent. The dispreference for anterior-postalveolar combinations is observable both within a trigram across a single vowel, and at a further distance on a strident projection, as shown by the observed/expected counts in Tables 1 and 2.³ As noted by Hansson (2001), postalveolar-anterior combinations are less frequent trans-vocally than at further distances, though they are by no means absent even across just a single vowel.

Table 1: Observed/expected cooccurrence of anterior and postalveolar stridents in Nkore-Kiga stems, transvocalic

	s, z	ʃ, ʒ
s, z	115/77 = 1.5	1/39 = 0.03
ʃ, ʒ	50/88 = 0.57	84/46 = 1.83

Table 2: Observed/expected cooccurrence of anterior and postalveolar stridents in Nkore-Kiga stems, across any amount of intervening material

	s, z	ʃ, ʒ
s, z	252/170 = 1.48	4/86 = 0.05
ʃ, ʒ	205/287 = 0.71	230/144 = 1.59

In Table 3, the observed counts for each combination of individual strident consonants is shown. Here, the affricates are included as well. While affricates are not described as participating in the harmony restriction, and do not alternate or trigger alternations, these counts allow us to see that the phonotactic distribution of affricates is actually fairly consistent with that of the fricatives. The postalveolar fricatives [ʃ ʒ] are rare following anterior fricatives, and they are also rare following the anterior affricate [ts]; the postalveolar affricate [tʃ] is completely unattested following both anterior fricatives and affricates.

³ Expected values were computed based on position-specific occurrence in the relevant string type (X₁ and X₃ in linear X₁X₂X₃ trigrams for “transvocalic” and projection-based X₁X₂ bigrams for “strident projection”). For example, the expected number of [s z]...[s z] bigrams in Table 1b is computed as the probability of [s z] as X₁ in a X₁X₂ bigram on the strident projection (256/687 = 0.37), multiplied by the probability of [s z] as X₂ in a X₁X₂ bigram on the strident projection (457/687 = 0.67), multiplied by the total number of X₁X₂ bigrams on the strident projection in the whole corpus (687): 0.37 x 0.67 x 687 = 170.3.

Table 3: Observed/expected cooccurrence of strident bigrams on a strident projection, by segment

	s	z	ts	ʃ	ʒ	ʃ̪
s	100/33 = 3.07	85/84 = 1.01	1/7 = 0.14	2/27 = 0.07	1/26 = 0.04	0/12 = 0
z	14/12 = 1.17	53/32 = 1.66	3/3 = 1.00	1/11 = 0.09	0/10 = 0	0/4 = 0
ts	4/7 = 0.57	19/17 = 1.12	14/1 = 14	1/6 = 0.17	1/5 = 0.20	0/2 = 0
ʃ	20/60 = 0.33	159/154 = 1.03	4/13 = 0.31	118/51 = 2.31	48/49 = 0.98	0/22 = 0
ʒ	6/16 = 0.38	20/42 = 0.48	6/3 = 2	1/14 = 0.07	59/13 = 4.5	2/6 = 0.33
ʃ̪	2/18 = 0.11	39/46 = 0.85	3/4 = 0.75	1/15 = 0.07	9/15 = 0.60	51/7 = 7.29

The counts in Tables 1-3 show that the harmony pattern among stridents looks somewhat different from previous descriptions, when assessed by phonotactic distribution across the dictionary corpus. First, stem alternations support harmony between anterior and postalveolar fricatives only, while the phonotactic distribution shows that the affricates are restricted as well (Table 3). Second, there are many disharmonic forms and only anterior-postalveolar combinations are restricted from a phonotactic perspective ([ʃ, ʒ]...[s, z] O/E = 0.57 in a transvocalic context in Table 1 and 0.71 more generally in Table 2).

4.4 Phonotactic learning simulations with projection induction

This section shows that the phonotactic restriction on anterior-postalveolar strident combinations is learnable from a baseline trigram. A relevant subset of the features given to the learner are shown in Table 3. The privative place features [anterior] and [postalveolar] were used to distinguish minor place in coronals, and [continuant], [voice] and [strident] were used to make the further necessary distinctions.⁴

Table 4: Feature specifications given to the learner to distinguish among coronal obstruents in Nkore-Kiga

	[anterior]	[postalveolar]	[voice]	[continuant]	[strident]
t	+	0	-	-	0
d	+	0	+	-	0
ts	+	0	-	-	+
ʃ̪	0	+	-	-	+
s	+	0	-	+	+
z	+	0	+	+	+
ʃ	0	+	-	+	+
ʒ	0	+	+	+	+

The learner finds a placeholder trigram constraint *[+anterior, +strident][][+postalveolar] in the baseline grammar at a range of settings. When gain is between 100-150 and gamma 10-20, this constraint is included in the grammar, which ranges from 40-75 constraints in size. The simulation reported here has a gain of 125 and a gamma of 10. The baseline grammar finds 58 constraints, one of which is *[+anterior, +strident][][+postalveolar]. This constraint penalizes [ts s z]-[-ʃ ʒ ʃ̪] trigrams and motivates the model to build a [+strident] projection, since [+strident] is the smallest natural class that includes all segments in both the [+anterior, +strident] and [+postalveolar] classes.⁵ The final grammar includes three constraint on the [+strident] projection, given in Table 5.

⁴ Through experience, it has been found that privative feature specifications lead to more interpretable results and more sensible looking grammars. Privative place features in particular allow each place class to be picked out with a single feature, as opposed to the two-feature specification in a more traditional analysis (e.g., [+postalveolar] as opposed to [+coronal, -anterior]).

⁵ Other placeholder trigram constraints in the baseline grammar motivate two other nonlocal projections: [+long] (all long vowels) and [-sonorant] (all obstruents). These projections appear to be superfluous, given my understanding of the phonotactics of Nkore-Kiga, though it is possible that they are capturing useful generalizations that have not been noted in the literature.

Table 5: Constraints found on the [+strident] projection for the Nkore-Kiga stem corpus, given in the order they were added to the grammar.

constraint	weight	sequences penalized
a. [+anterior][+postalveolar]	3.87	[s z ts]...[ʃ ʒ ʃʃ]
b. [+continuant][-continuant]	2.66	[s z ʃ ʒ]...[ts ʃʃ]
c. [-continuant][+continuant, -voice]	2.67	[ts ʃʃ]...[s ʃ]

The constraints in Table 5 penalize all combinations of anterior-postalveolar stridents, at any distance, whether the stridents are fricatives or affricates. Overall, the grammar captures the harmony restriction with constraint a, assigning a higher penalty to combinations of anterior-postalveolar stridents than to other combinations of stridents. The grammar also includes two additional constraints on affricates (constraints b and c), which do not distinguish harmonic from disharmonic combinations (though they capture other aspects of the distribution of stridents).

The grammar was evaluated based on the scores assigned to a set of test words. Test words were nonce words, constructed based on real words. The real words used to form the testing set were all of those that contained two harmonic strident fricatives, neither of which were followed by [i u j w]. Bennett & Pulleyblank (2018) show that the anteriority contrast is neutralized before high vowels and glides in the language, so forms that would be under the purview of these CV phonotactics were removed to assess the status of anteriority harmony specifically. From each real word, three nonce test words were made. One form switched the anteriority value of the first strident, one form switched the anteriority value of the second strident, and one form switched the anteriority value of the both stridents. For example, the real word [koze:sa] led to the creation of the three nonce words [koʒe:sa], [koze:ʃa] and [koʒe:ʃa]. The test set consisted of 252 nonce words, with a range of phonotactic shapes and with stridents at a range of distances. There were 48 anterior-anterior forms, 36 postalveolar-postalveolar forms, 84 anterior-postalveolar forms and 84 postalveolar-anterior forms. The number of test forms is not balanced between minor place combinations because the lexicon itself is not balanced.




Figure 1: Average (red diamond) and distribution of scores assigned to nonce test words by the final grammar trained on Nkore-Kiga stems

Figure 1 shows the average and distribution of scores assigned to these test words, by place of articulation of the strident fricatives. Forms with anterior-postalveolar (anterior-postalveolar) combinations are given lower scores on average than all other combinations, indicating that they are penalized more by the grammar. Importantly, no anterior-postalveolar form is given a comparably high score to the best forms

in the other categories: the highest score assigned to an anterior-postalveolar form is -10 while the highest score assigned to forms in each of the other three categories is -6.

The grammar generally does not distinguish between anterior-anterior, postalveolar-postalveolar and postalveolar-anterior combinations, all of which have average scores around -10. Moreover, while the average scores do differ between the restricted anterior-postalveolar combinations and the other three, there is substantial overlap in scores for all nonce words. This indicates that the weight of the constraints in Table 5 is not so great that it eclipses any other phonotactic constraints in the grammar.

4.5 Summary and discussion

The investigation of Taylor's dictionary corpus reveals a somewhat different picture of Nkore-Kiga strident harmony than what is reported in the previous literature. First, harmony triggered by morphological alternations seems to be variable. Some stridents harmonize in minor place to agree with a following strident that alternates under affixation, but by no means all do. Second, Hansson describes harmony holding transvocally for both combinations (anterior-postalveolar and postalveolar-anterior) of stridents, but the statistical evidence only supports a prohibition on anterior-postalveolar combinations. This restriction is observable both transvocally and at longer distances. While there are exceptions to harmony, the dispreference for anterior-postalveolar combinations is strong enough that it is discovered by an inductive constraint procedure both as a local trigram and as a constraint on a non-local strident projection.

The phonotactic distribution uncovered here calls into question some aspects of Bennett and Pulleyblank's analysis. Under their analysis, a form like [-sà:sì] results from an input /sà:sì/ (cf. [-fà:sà] 'be in pain'). The alternation in the final strident from /s/ to [s] is triggered by the /-i/ suffix (analyzed as a morphological mutation, though the sequence [sì] is also generally nearly unattested in the language), and the initial strident then alternates to agree with the final strident. The distribution of stridents in the dictionary corpus, however, suggests that the motivation for this harmony is not phonotactic. A hypothetical form like [sà:sì], with morphological mutation but no harmony, appears to be phonotactically licit. The alternations seen in Nkore-Kiga thus seem to be lexically specific as well as morphologically conditioned, and aren't consistent with the general phonotactics of the language as assessed by the descriptive statistics or learning simulations presented in this section (though there may be alternative assessments of phonotactics that would reveal subtle but useful support for a restriction on postalveolar-anterior strident combinations). A model of strident harmony acquisition in this language may require morphological knowledge and a subsetting of the data into harmonizing and non-harmonizing forms, suggesting that harmony as a productive pattern would emerge later in learning than if it were supported by word-level phonotactic patterns.

5 Case study 2: Papantla Totonac

The Totonacan languages are reported to have minor place harmony between stridents [s ts ʃ ʃ̪] in Hansson (2001). MacKay (1999) describes harmony in Misantla Totonac as morphologically sensitive, applying within the root and triggering alternations in derivational prefixes. Her discussion is brief, and only a single example of an alternation is given. In (8), the prefix normally realized as [ʃ̪a:-] is realized as [ts'a:-] if the root contains a [+anterior] strident, in this case [s].

(8) min-ʃ̪a:-n'i 'your body'
 ts'a:-st'alah 'clean bodied'

Papantla Totonac was chosen for further study because there is both a dictionary and an An Crúbadán web corpus. The goals for this section are to assess harmony in more detail by looking at the prevalence of exceptions to harmony in word forms in both corpora, and to compare the evidence for harmony in words (where exceptions are expected) to those in stems (where harmony is expected to hold more strongly).

5.1 The corpora

5.1.1 *Dictionary stem corpus*

The Aschmann & Aschmann (1973) dictionary was used to create a corpus of stems. This list consists of dictionary head words, with obvious prefixes (mostly, those listed as such in the dictionary) parsed out. When it looked like multiple entries contained the same root with different suffixes, only the shortest form was included. The resulting list of 1814 forms likely still contains some multi-morphemic items, since a complete morphological decomposition isn't possible based on available sources, but it is substantially different from the larger lists of words in the dictionary and web corpora.

For the stems corpus and the word corpora, the transcription followed the description of the orthography in Aschmann & Aschmann. While *<e o>* are used in the orthography, these vowels are found only in the vicinity of uvulars and were changed to [i u] in the transcription. Forms with *<r>* were removed from all lists, as this is a non-native phoneme and indicates a loan word. Stress is contrastive and marked in the dictionary, but was removed from the transcription for the data analyzed here.⁶

5.1.2 *Dictionary word corpus*

The dictionary word corpus was constructed by taking all head words and all words in example sentences from the Aschmann & Aschmann (1973) dictionary. The result was a list of 7480 unique words.

5.1.3 *An Crúbadán web corpus*

The unprocessed An Crúbadán web corpus contains 15,490 forms. After removing English, Spanish and forms with odd characters, there were 10,552 unique word forms for analysis. Glottalization, which is distinctive on vowels and indicated with an apostrophe in the orthography, is very rare in the web corpus (appearing in just 28 forms) and is therefore likely not consistently transcribed in online text. Because of this unreliability and infrequency, forms with glottalization marked were removed from the corpus. Glottalization should be orthogonal to strident harmony, and its role can be assessed based on the other two data sets which include it.

5.2 Descriptive statistics and learning simulations

5.2.1 *Dictionary stem corpus*

The stem corpus extracted from the dictionary shows a strong harmony preference. Combinations of anterior and postalveolar stridents are underattested, both within a trigram and at further distances, as shown in Tables 6 and 7. There are tri-consonantal clusters in Totonac, so stridents within a trigram may be separated by either a vowel or a consonant.

Table 6: Observed/expected cooccurrence of anterior and postalveolar stridents in the Totonac dictionary stems corpus, transvocalic

	s, ts	ʃ, ɺ
s, ts	33/20 = 1.65	0/13 = 0
ʃ, ɺ	1/14 = 0.07	23/10 = 2.30

⁶ The removal of mid vowels and stress were both done to limit the number of natural classes available to the learner, resulting in cleaner and more interpretable grammars for present purposes.

Table 7: Observed/expected cooccurrence of anterior and postalveolar stridents in the Totonac dictionary stems corpus, across any amount of intervening material

	s, z	f, ʃ
s, z	56/32 = 1.75	2/27 = 0.07
f, ʃ	2/26 = 0.08	45/21 = 2.14

The four exceptions to harmony are given in (9). While these may be poly-morphemic, and thus not constitute morpheme-level exceptions, no obvious decomposition was available based on the dictionary or MacKay's grammar.

(9) a. postalveolar...anterior
 qantʃa:stu:n 'where the wall forms a corner with the floor'
 ma:k-ʃkatsi:ni 'prevent it'

b. anterior...postalveolar
 maspu:ʃtu 'the heart of a tree'
 suktʃa:ʃ 'calandria (kind of bird)'

Table 8 shows a relevant subset of the features given to the learner. The privative place features [anterior] and [postalveolar] were used to distinguish minor place in coronals, and [lateral], [continuant] and [strident] were used to make the further necessary distinctions. These same feature specifications were used for all three Totonac training sets (dictionary stems, dictionary words, and the web word corpus).

Table 8: Feature specifications given to the learner to distinguish among coronal obstruents in Totonac (all simulations)

	[anterior]	[postalveolar]	[lateral]	[continuant]	[strident]
t	+	0	-	-	0
tl	+	0	+	-	0
ts	+	0	-	-	+
ʃ	0	+	0	-	+
ɬ	+	0	+	+	0
z	+	0	-	+	+
ʃ	0	+	0	+	+

Given the stem list as training data, the baseline grammar includes the exceptionless constraint on [+anterior, +strident]-[]-[+postalveolar] combinations at a range of settings. Gains of 15-30 and gamma of 0-10 yielded moderately sized grammars of 40-90 constraints. The single exception to the target trigram constraint [+postalveolar]-[]-[+anterior, +strident] results in this constraint not being included in the baseline grammar. Instead, the infrequency of postalveolar-anterior combinations is accounted for with the more specific and accurate constraint [+postalveolar]-[-glottal]-[+anterior, +strident]. This constraint only penalizes postalveolar-anterior combinations across non-glottalized vowels, thus dancing around the one exception [qantʃa:stu:n] which contains a glottalized vowel between the interacting stridents. If gain is raised above 30 (which should encourage more general constraints), no constraints on strident cooccurrence are found at all (include the target placeholder trigram constraint). This result shows the fragility of the placeholder trigram approach to learning nonlocal projections in the face of exceptions in natural language data.

In Totonac, however, the [+strident] projection can still be learned based on the one exceptionless placeholder trigram constraint, [+anterior, +strident]-[]-[+postalveolar]. A final model with this projection does include both bigram constraints on disharmonic strident combinations, accounting for the harmony pattern within stems. The model reported here had a gain of 30 and a gamma of 5, and includes the two constraints in Table 9 on the [+strident] projection.

Table 9: Constraints found on the [+strident] projection for the Totonac dictionary stems corpus, given in the order they were added to the grammar.

constraint	weight	sequences penalized
a. [+postalveolar][+anterior]	4.67	[ʃ tʃ]...[s ts]
b. [+anterior][+postalveolar]	3.95	[s ts]...[ʃ tʃ]

While the presence of exceptions at the level of a baseline trigram does pose difficulties for the learner in finding the target placeholder trigrams, the smaller and more targeted search space of a nonlocal projection renders the small number of exceptions in the data here unproblematic for finding the appropriately general bigram constraints on disharmonic strident pairs.

The grammar was tested on a set of nonce forms with two stridents, to see how it captures strident harmony. The testing words were constructed in the same manner as for Nkore-Kiga, by changing the anteriority value of stridents in real words with the goal of including a range of phonotactic shapes that is characteristic of the language as a whole. The resulting testing set had 43 anterior-anterior combinations, 88 anterior-postalveolar combinations, 88 postalveolar-anterior combinations and 50 postalveolar-postalveolar combinations.




Figure 2: Average (red diamond) and distribution of scores assigned to nonce test words by the final grammar trained on Totonac dictionary stems

As can be seen in Figure 2, both disharmonic combinations of stridents receive lower scores on average than both harmonic combinations. There is some overlap in the distribution of scores, due to a small number of forms that receive very low scores, though in general the model makes a clear distinction between harmonic and disharmonic forms. The forms with very low scores violate constraints that are orthogonal to strident harmony, and may or may not represent systematic restrictions. For example, the grammar includes a constraint on [ts] followed by non-glottalized long vowels. This constraint may be an example of the grammar over-fitting, or it could be a real restriction in speakers' grammars.

To summarize, the list of Totonac stems taken from the dictionary shows minor place harmony among stridents, though there are a few exceptions. The inductive learning model can use the presence of a placeholder trigram constraint in the baseline grammar to build a nonlocal [+strident] projection. A final grammar with this nonlocal projection includes constraints enforcing strident harmony, and distinguishes between harmonic and disharmonic nonce forms. The simulations also show, however, that even a single

exception to a target placeholder trigram can prevent the constraint from being included in the baseline grammar. We return to this issue in the general discussion.

5.2.2 Dictionary word corpus

The dictionary word corpus contains many morphologically complex forms, and there are many counter-examples to strident harmony. Tables 10 and 11 show that a dispreference for anterior-postalveolar combinations is observable within a trigram, but is much weaker across longer distances. There is also only a weak dispreference for postalveolar-anterior combinations at either distance.

Table 10: Observed/expected cooccurrence of anterior and postalveolar stridents in the Totonac dictionary word corpus, transvocalic

	s, ts	ʃ, ʃ̄
s, ts	122/63 = 1.94	3/63 = 0.05
ʃ, ʃ̄	29/88 = 0.33	147/88 = 1.67

Table 11: Observed/expected cooccurrence of anterior and postalveolar stridents in the Totonac dictionary word corpus, across any amount of intervening material

	s, z	ʃ, ʃ̄
s, z	225/130 = 1.73	148/243 = 0.61
ʃ, ʃ̄	35/60 = 0.58	337/112 = 3.01

The learner does find a placeholder trigram constraint on anterior-postalveolar strident combinations, though it only does so when gamma is quite low. With a gamma of 5, the model finds [+anterior, +strident]-[]-[+postalveolar] at gains of 30-60. If gamma is raised, the model includes more specific and accurate constraints, e.g., *[+anterior, +strident]-[+continuant]-[+postalveolar], to which there are only two exceptions, instead of the placeholder trigram constraint.

While the model can induce the [+strident] projection from a baseline trigram, the final grammar only weakly enforces harmony. This is unsurprising, given that harmony is only weakly observed outside of trigrams. The model reported here has a gain of 60 and a gamma of 5. There are two constraints on the [+strident] projection, given in Table 12. Constraint a enforces harmony by penalizing anterior-postalveolar combinations, but it has a low weight. Constraint b penalizes both harmonic and disharmonic combinations equally.

Table 12: Constraints found on the [+strident] projection for the Totonac dictionary word corpus, given in the order they were added to the grammar.

constraint	weight	sequences penalized
a. [+anterior][+postalveolar]	2.56	[s ts]...[ʃ ʃ̄]
b. [+postalveolar, -continuant][][]	11.33	[ʃ]...[s ts ʃ ʃ̄]...[s ts ʃ ʃ̄]

The grammar was tested on a set of nonce forms with two stridents, constructed in the same fashion as for previous simulations. The resulting testing set had 278 anterior-anterior combinations, 446 anterior-postalveolar combinations, 447 postalveolar-anterior combinations and 177 postalveolar-postalveolar combinations. As can be seen in Figure 3, the model only weakly enforces a harmony preference, with very slightly lower average scores for disharmonic combinations of stridents than for harmonic combinations. This is expected, given the constraints in Table 9 and the O/E scores in Table 8. There is only a weak harmony effect in the dictionary word list (O/Es of 0.58 and 0.61 for disharmonic combinations on the strident projection), and the grammar here confirms that this mild under-attestation is not strong enough to have a major effect on the constraints in the grammar. Moreover, the harmony constraint in (a) in Table 9 ensures that no nonce word with an anterior-postalveolar combination receives a score higher than -8.5, but some words from each of the other categories receive scores of -6, the highest score assigned by the

grammar.⁷ The substantial overlap in the distributions of scores shows that even an explicit harmony constraint is not highly weighted enough relative to other phonotactic constraints to achieve a strong preference for harmony. The testing forms were created from real words, and so variation in assigned scores reflects gradient phonotactic effects among ‘licit’ structures.




Figure 3: Average (red diamond) and distribution of scores assigned to nonce test words by the final grammar trained on Totonac dictionary words

The comparison between the stem list and the word list from the dictionary shows that strident harmony is morphologically sensitive, and that the number of exceptions that arise in word forms are sufficient to at least partially obscure the pattern as a whole. While a placeholder trigram is still observable in the baseline grammar trained on words, the large number of exceptions to strident harmony on a strident projection result in only a very weak representation of harmony. Totonac learners, then, must be performing an analysis over roots or stems if they are to arrive at a grammar with a clear harmony pattern.

5.2.3 Web word corpus

Like the dictionary word corpus, the web corpus of words contains many exceptions to strident harmony, both within a trigram and at longer distances. Tables 13 and 14 show that anterior-postalveolar combinations are dispreferred within a trigram, though there are many more such combinations on a strident projection. This is the same pattern seen in the dictionary corpus. In the other order, postalveolar-anterior combinations are also somewhat underattested within a trigram, but are well-attested at longer distances. By contrast, in the dictionary word corpus postalveolar-anterior combinations are still somewhat underattested on a strident projection, so harmony is generally even weaker in the web corpus than in the dictionary corpus.

⁷ The final grammar often includes a constraint ${}^*[]$, which serves to penalize longer words by assigning a violation for each segment. Grammars with this constraint do not assign a perfect score of 0 to any form, hence the highest score in this simulation being -6.

Table 13: Observed/expected cooccurrence of anterior and postalveolar stridents in the Totonac web word corpus, transvocalic

	s, ts	ʃ, ʃ̪
s, ts	45/17 = 2.65	3/31 = 0.10
ʃ, ʃ̪	20/48 = 0.42	116/88 = 1.32

Table 14: Observed/expected cooccurrence of anterior and postalveolar stridents in the Totonac web word corpus, across any amount of intervening material

	s, z	ʃ, ʃ̪
s, z	105/64 = 1.64	33/74 = 0.45
ʃ, ʃ̪	553/594 = 0.93	720/679 = 1.06

A round of learning simulations did not find any placeholder trigram constraints on strident combinations, regardless of gain and gamma settings (gain was varied from 10-150 and gamma from 0-50, yielding grammars of 60-200 constraints). Further inspection of the baseline grammars found that the model was finding constraints that accounted for the exceptions to harmony within a trigram. Specifically, the three trans-segmental anterior-postalveolar combinations all contain the same sequence, [...tsaʃ...], with a low vowel intervening between the two stridents, so the model includes a constraint [+strident, +anterior][+high][+postalveolar] as opposed to the more general [+strident, +anterior][[]][+postalveolar]. Instead of learning a placeholder trigram, and thus postulating a non-local projection, the baseline grammar includes a more specific and accurate trigram constraint that specifies a certain type of intervening segment. In this case, the model is not missing much, because there is no strong harmony pattern on the strident projection, but this result again shows that even a few exceptions may be problematic for the placeholder trigram approach to learning nonlocal projections, if those exceptions happen to contain an accidental generalization that the learner can work around.

5.3 Summary

The exploration of three different Totonac corpora show that strident harmony is present in stems, but only weakly observable in words. In this case, the number of word-level exceptions does obscure the pattern seen in stems when looking at the strident projection, though some trigram constraints are found that reflect more local harmony. When looking at stems, our model is able to find placeholder trigram constraints on the baseline and build a nonlocal projection and final grammar that enforces harmony in general terms. In the dictionary words corpus, a placeholder trigram was found on the baseline but the final grammar included only a single low-weighted constraint on one of the two disharmonic combinations. While the qualitative pattern in the training data is similar in the web words corpus and the dictionary words corpus, the patterning of exceptions in the web words corpus is such that the baseline model does not include a placeholder trigram constraint at all.

All three simulations show the fragility of the placeholder trigram approach. If there are exceptions to a nonlocal restriction within a trigram, and there is some regularity to the intervening segments in these exceptions, the model may include a more specific and accurate trigram constraint in the baseline grammar as opposed to the target placeholder trigram.

6 Case study 3: Navajo

Navajo is one of the most frequently cited cases of strident harmony. Minor place harmony is reported to hold between all stridents [dʒ ts ts' s z dʒ ʃ ʃ̪ ʃ̪̪ ʃ̪̪̪ ʃ̪̪̪̪ ʃ̪̪̪̪̪], both within a root and via alternations in certain prefixes. Sapir & Hoijer (1967) report that alternations are optional in slow speech, and that alternations occur more often across a single vowel than at longer distance. The alternations are also morphologically sensitive: some prefixes alternate and others do not. The alternating prefixes are analyzed as occupying positions

closer to the root (referred to as the “conjunction stem” in McDonough (2003)), and the non-alternating prefixes are more peripheral (comprising the “disjunction stem”). Sapir & Hoijer give examples of alternations in four conjunction prefixes and one disjunction prefix, as shown here in (10) through (14). For each prefix, alternations are shown by allomorph (which may vary in presence/absence of a vowel, voicing and continuancy). Nasalization is indicated with a cedilla (e.g., [í] indicates a nasal vowel with a high tone), following the widely used Young & Morgan (1972) orthography. The alternating morpheme is in bold.

(10) Perfective prefix [si] ~ [ʃi], [s] ~ [ʃ], [z] ~ [ʒ]

a.	si-? _á	‘a round object lies’	ʃi-? _{íʃ}	‘it is bent, curved’
	si- _í	‘he is lying’	ʃi- _í éé _ʒ	‘they two are lying’
			ʃi- _í zaa?	‘a mass lies’
			ʃi- _í ʃid	‘he is stooping over’
b.	dee-s-gai	‘whiteness extends’	ji-ʃ-taʃ	‘he has given it a tap’
		outward’	ʃ-ii- _í éé _ʒ	‘blackness extends outward’
			dah-da-dʒii-ʃ-ʃa?	‘he tied them here and there’
			b-aa-da-? _a -ʃ-ʃ’os	‘he has become putrid with maggots’
c.	?a-z-? _á	‘some round object is lying’	ʒi-ʒ-? _{íʃ}	‘he (4 th p) is stooped over’
	ʒi- _í z- _í	‘he (4 th p) is lying’	ʒi-ʒ- _í éé _ʒ	‘they (4 th p) two are lying’
			ha-haa-ʒ-dʒéé?	‘several sticklike objects lie across’
			ʒi-ʒ-? _í ʃid	‘he (4 th p) is stooped over’
			?a-ʒ-ʃ’ah	‘he has fallen into fire’

(11) 1st sg. poss [ʃi] ~ [si]

ʃi-taa?	‘my father’	si-sázi	‘my ancestor’
ʃi-lí?	‘my horse’	si-zid	‘my scar’
		si-dzástis	‘my shin bone’
		si-tsíli	‘my younger brother’
		si-ts’aa?	‘my basket’

(12) 1st sg. subject [ʃ] ~ [s]

di-ʃ-bááh	‘I start off to war’	tá-? _á -dí-s-gis	‘I wash myself’
ha-ʃ-taal	‘I sing a ceremony’	n-s-neež	‘I am tall’
		da-s-caah	‘I am dying’
		dii-s-ts’aa?	‘I hear about it’

(13) 4th p subject ʒi] ~ [dʒi], [ʃ] ~ [s], [ʒ] ~ [z]

a.	ʒi-di-bááh	‘he (4 th) starts off to war’	ʒi-siil	‘he (4 th) steams it’
	ha-ʒi-taal	‘he (4 th) sings a ceremony’	n-ʒi-z-? _á	‘they (4 th) lay around’
			da-ʒi-dziih	‘they (4 th) survive’
			da-ʒi-tsaah	‘they (4 th) are dying’
			?á-da-ʒi-l-tsíísi	‘they (4 th) are small’
b.	da-ʃ-dii-tíjh	‘he (4 th) lifts up a sticklike object’	da-s-doo-l-? _í is	‘he (4 th) has his foot raised’
	da-ʃ-dii-l-teeh	‘he (4 th) lifts him up’	da-s-di-dii-l-? _í eez	‘he (4 th) has raised his foot’
			da-s-dii-l-tsooz	‘he (4 th) has started off with a fabric-like object’

c. ná-**z**-dii-l-tí ‘he (4th) picked him up’ tá-?á-**z**-di-gis ‘he (4th) washes himself’
 ná-**z**-dii-káah ‘they (4th) started back’ da-**z**-dée-**z**-?íí? ‘they (4th) looked up’
 home’

ná-**z**-dii-dzá ‘he (4th) has started back home’
 da-**z**-doo-**tsaal** ‘he (4th) will die’
 ho-**z**-dii-**ts'a?** ‘he (4th) heard about things’

(14) Disjunct prefix [dzi] ~ [dʒi]

dzí -bá ‘it is gray-striped’	dʒi -tl’iʒ	‘it is blue-striped’
dzí -gai ‘it is white-striped’	dʒi -dʒin	‘it is black-striped’
	dʒi -ʃíi?	‘it is red-striped’

Examples of words with prefixes that don’t alternate are given in (15), and the examples in (16) show that enclitics do not trigger alternations in a preceding stem.

(15) Disharmonic forms with non-alternating prefixes

bi- za -dʒi-l-téeh	‘he bridles (a horse)’
bi- zé -ná-ʃ-nih	‘I embrace him’
tsí -i-t’aj	‘we two go along seeking safety’
tsí -dʒi-káah	‘they (4 th) go along seeking safety’
bi- ts’á -dʒi-l-γod	‘he ran away from him’
ʃó-joo-s-t’e	‘he has acquired it’
dʒi -s-í-ba?	‘I have done a kind act’
tʃáah -di-s-maas	‘I stumble and roll over’
ha- tʃíi? -nee-z-dá	‘he sat down before him’

(16) Disharmonic forms with enclitics

ná-hoo-kos- dʒí?	‘to the north’
kó-dzaa-go- ʃíi	‘when apparently this happened’
ji-da-?-nii-l-tseed- ʃa?	‘it appears that they were killing them’
?al-ts’áah- dʒí	‘on opposite sides’

The data above show that strident harmony is morpheme specific in Navajo. This likely presents a challenge to the learner, since the phonotactic distribution of stridents within words as a whole may not represent strident harmony. Instead, morphologically defined subsets of the lexicon must be examined to discover the harmony pattern. Moreover, harmony is optional in some cases. Martin (2011) looked at 211 compounds with multiple stridents, and found that harmony held 70% of the time for stridents in adjacent syllables and 44% of the time for stridents in non-adjacent syllables. Berkson (2013) found that evidence for harmony in the 1st singular possessive prefix [ʃí]~[sí] ([ʃí] in forms with no stridents) was weak, across several different measures: the [ʃí-] form was preferred for all roots in a judgment study of orthographic forms, a small production study found consistent production of the prefix as [ʃí-], and very low rates of harmony were found in an analysis of online written Navajo assessed via the number of Google hits. Both of these findings suggest that harmonic alternations are both variable and perhaps undergoing change.

In the rest of this section, evidence for harmony is evaluated in three different corpora. The first corpus is a list of monomorphemic stems taken from the Young, Morgan & Midgette (1992) dictionary. Then, a corpus of words from Young & Morgan (1972) is examined and compared to the An Crúbadán web corpus, as for Totonac.

6.1 The corpora

6.1.1 *Dictionary stem corpus*

The stem corpus was a list of verb and noun stems from the Young, Morgan & Midgette (1992) dictionary. The verb stems were compiled for Eddington & Lachler (2006), and the noun stems were extracted from the dictionary directly.⁸ The stem list contains 917 unique items. The items classified as “stems” in this dictionary are monomorphemic, so they could be considered roots. Tone was not included in the transcription used here (as in Nkore-Kiga & Totonac, where suprasegmental information was excluded as well). Otherwise, the corpus was transcribed into IPA from the orthography based on the correspondences described in Young & Morgan (1972).

6.1.2 *Dictionary word corpus*

The dictionary word corpus was constructed by extracting word forms from the Young & Morgan (1972) dictionary. Only inflected, stand-alone words were included. Head words were often abstract stems, which were not included in the list of words. The result was 7830 unique words.

6.1.3 *An Crúbadán web corpus*

The web word corpus was based on the An Crúbadán web corpus of 30,526 forms. After cleaning the corpus to remove English and Spanish, as well as odd characters, hashtags and web addresses, there were 19,007 word forms. Nasality is not transcribed in the corpus, so while vowels are contrastively nasalized in Navajo, this distinction is not represented in the web word corpus.

6.2 Descriptive statistics and learning simulations

6.2.1 *Dictionary stem corpus*

Harmony in the stem corpus is nearly exceptionless. Stems are primarily monosyllabic (particularly verb stems), and while the majority of stridents are separated by just a single vowel, there are some consonant clusters and polysyllabic stems where stridents are separated by more segmental material. The counts in Tables 15 and 16 show all pairs of stridents, those separated by a single vowel as well as those at longer distances, and show that there are just two exceptions to minor place harmony within stems.

Table 15: Observed/expected cooccurrence of anterior and postalveolar stridents in the Navajo stems corpus, transvocalic

	[+ant]	[-ant]
[+ant]	40/20 = 2	1/21 = 0.05
[-ant]	1/21 = 0.05	41/21 = 1.95

Table 16: Observed/expected cooccurrence of anterior and postalveolar stridents in the Navajo stems corpus, across any amount of intervening material

	[+ant]	[-ant]
[+ant]	52/23 = 2.26	1/30 = 0.03
[-ant]	1/30 = 0.03	67/38 = 1.76

The two exceptions to minor place harmony both come from compound nouns (the elements of the compound are listed separately in the stem list), where the disharmony could be a result of local assimilation to an immediately adjacent strident. These nouns are not listed separately outside of these compounds,

⁸ Available at <https://linguistics.byu.edu/faculty/deddingt/navajo>.

however, so it is unknown whether the tautomorphemic stridents are underlying harmonic or not (e.g., if [ts'af] is [ts'as] when not followed by a postalveolar strident). These compounds are given in (17).

(17) ts'af + dʒiʃ ‘diaphragm’
dʒas + tso:? ‘clitoris’

The Navajo models were given a feature set that distinguishes stridents based on the features shown in Table 17. As for previous languages, two privative place features are used to make the minor place distinctions.

Table 17: Feature specifications given to the learner to distinguish among coronal obstruents in Navajo (all simulations)

	[anterior]	[postalveolar]	[lateral]	[cont.]	[const. gl.]	[voice]	[strident]
t	+	0	-	-	0	-	0
t'	+	0	-	-	+	-	0
d	+	0	-	-	0	+	0
tł	+	0	+	-	0	-	0
tł'	+	0	+	-	+	-	0
dl	+	0	+	-	0	+	0
ts	+	0	-	-	0	-	+
ts'	+	0	-	-	+	-	+
dʒ	+	0	-	-	0	+	+
ʃ	0	+	0	-	0	-	+
ʃ'	0	+	0	-	+	-	+
dʒ'	0	+	0	-	0	+	+
ł	+	0	+	+	0	-	0
s	+	0	-	+	0	-	+
z	+	0	-	+	0	+	+
ʃ	0	+	0	+	0	-	+
ʒ	0	+	0	+	0	+	+

When the learner is given the root corpus as training data, the baseline model includes general placeholder trigram constraints that enforce harmony. The two constraints, [+anterior, +strident][][][+postalveolar] and [+postalveolar][][][+anterior, +strident], are found in grammars with a gain of 5-15 and a gamma of 0-3 (grammars of 50-100 constraints). When gamma is above 3 neither constraint is found. Gain must be relatively low because of the small number of training items, and gamma must be low because of the two exceptions to harmony. The model reported has a gain of 15 and a gamma of zero. The baseline grammar includes both placeholder trigram constraints, which motivate a [+strident] projection. On the strident projection, the constraints in Table 18 are found.

Table 18: Constraints found on the [+strident] projection for the Navajo stem corpus, given in the order they were added to the grammar.

constraint	weight	sequences penalized
a. [][-continuant]	0.084	[s z dʒ ts ts' ʃ ʒ dʒ ʃ ʃ']...[dʒ ts ts' dʒ ʃ ʃ']
b. [+wb][-continuant, +voice]	1.42	#...[dʒ dʒ]
c. [+postalveolar][+anterior]	3.74	[dʒ ʃ ʃ' ʃ ʒ]...[dʒ ts ts' s z]
d. [anterior][+postalveolar]	3.71	[dʒ ts ts' s z]...[dʒ ʃ ʃ' ʃ ʒ]
e. [+wb][+continuant, +voice]	0.43	#...[z ʒ]

Constraints c and d straightforwardly enforce harmony, while the other constraints reflect the distribution of stridents in ways that are orthogonal to harmony. The harmony enforcing constraints have higher weights than the other constraints. The class [+wb] is a word boundary (either beginning or end).

A testing set was constructed based on the root training data in the same manner as for previous simulations. There were 39 forms with anterior-anterior combination, 69 with a postalveolar-postalveolar combination, 69 with a anterior-postalveolar combination and 40 with a postalveolar-anterior combination. The grammar assigns lower average scores to disharmonic strident combinations than to harmonic strident combinations, as can be seen in Figure 4.




Figure 4: Average (red diamond) and distribution of scores assigned to nonce test words by the final grammar trained on Totonac dictionary stems

While there is substantial overlap in the scores assigned to harmonic and disharmonic forms (due to constraints in the grammar that are orthogonal to strident harmony), no disharmonic form receives a score as high as the highest harmonic forms. The highest score assigned to a harmonic nonce form is -7, and the highest score assigned to a disharmonic nonce form is -10. As in previous simulations, the lowest scores are assigned to forms that violate constraints that are orthogonal to harmony and may reflect overfitting by the learner.

6.2.2 Dictionary word corpus

In the dictionary word corpus, harmony is still strongly observed, despite some exceptions. Tables 19 and 20 show that disharmonic stridents are underattested both transvocalically and on a strident projection. Harmony appears to be symmetric, with both anterior-postalveolar and postalveolar-anterior combinations being comparably underattested, and harmony also appears to hold with equal strength at both distances.

Table 19: Observed/expected cooccurrence of anterior and postalveolar stridents in the Navajo dictionary word corpus, transvocalic

	[+ant]	[-ant]
[+ant]	213/109 = 1.95	14/104 = 0.13
[-ant]	14/104 = 0.13	203/99 = 2.05

Table 20: Observed/expected cooccurrence of anterior and postalveolar stridents in the Navajo dictionary word corpus, across any amount of intervening material

	[+ant]	[-ant]
[+ant]	1035/558 = 1.85	83/560 = 0.15
[-ant]	61/538 = 0.11	1016/539 = 1.88

Interestingly, transvocalic stridents show about the same distribution as all stridents, despite the observation in previous descriptions that harmony is optional at further distances. Looking more closely, it turns out that part of the strength of harmony on the strident projection is due to the large number of strictly adjacent stridents, which show an almost exceptionless harmony pattern. Tables 21 and 22 compare the distribution of minor place in stridents in strident-strident clusters and strident...strident bigrams in non-adjacent syllables. Indeed, when looked at this way, the harmony preference is stronger in adjacent stridents, and weakened in non-adjacent syllables. Overall, 99% of string adjacent stridents are harmonic, 95% of syllable adjacent stridents are harmonic, and just 71% of more distant stridents are harmonic.

Table 21: Observed/expected combinations of strictly string adjacent stridents in the Navajo dictionary words corpus, by place of articulation

	[+ant]	[-ant]
[+ant]	490/263 = 1.86	3/230 = 0.01
[-ant]	2/229 = 0.01	427/200 = 2.14

Table 22: Observed/expected combinations of stridents in non-adjacent syllables in the Navajo dictionary words corpus, by place of articulation

	[+ant]	[-ant]
[+ant]	84/53 = 1.59	61/92 = 0.66
[-ant]	35/66 = 0.53	147/116 = 1.27

While the strength of harmony does decay with distance, this difference won't be noticeable to the learner. The model has access only to the linear string and to a strident projection (provided one is induced based on the linear string), and thus can only assess the usefulness of harmony preferring constraints based on the distributions shown in Table 20. This lack of distance sensitivity is a property of projections as currently defined.

The baseline models include placeholder trigram constraints on combinations of disharmonic stridents when gamma is low (0 or 5) and gain is 125 or lower. These settings yield grammars of 50-100 constraints. The model reported here has a gain of 125 and a gamma of 5. The baseline grammar includes the two placeholder trigram constraints on both disharmonic combinations of stridents: [+anterior, +strident][][][+postalveolar] and [+postalveolar][][][+anterior, +strident]. Both of these constraints motivate the learner to search through a [+strident] projection when building the final grammar. The constraints on the [+strident] projection that are included in the final grammar are given in Table 23. None of these constraints straightforwardly enforces harmony, though the constraints in c and d do specifically target some disharmonic combinations.

Table 23: Constraints found on the [+strident] projection for the Navajo dictionary words corpus, given in the order they were added to the grammar.

constraint	weight	sequences penalized
a. [-continuant][-continuant]	2.34	[ts ts'dzʃʃdʒ]...[ts ts' dzʃʃdʒ]
b. [+continuant, +voice][-continuant]	2.40	[z ʒ]...[ts ts'dzʃʃdʒ]
c. [+postalveolar][+voice, +dental]	3.90	[ʃʃdʒʃʃdʒʃʃdʒ]...[dʒ z]
d. [-cont, -voice, -lateral][+postalv.]	3.49	[s]...[ʃʃdʒʃʃdʒʃʃdʒ]

The grammar was tested on a set of nonce words. The testing set had 478 anterior-anterior forms, 836 anterior-postalveolar forms, 833 postalveolar-anterior forms and 369 postalveolar-postalveolar forms. As can be seen in Figure 5, the final grammar does not clearly capture harmony, as expected based on the constraints in Table 16. The average scores for disharmonic combinations are lower than for harmonic combinations, but all four minor place combinations have large concentrations of forms that receive a perfect score of 0. Disharmonic forms do receive lower scores on average, since certain disharmonic segmental combinations are penalized by the grammar. The lower scores for anterior-postalveolar combinations are due to the relatively large number of [s]-postalveolar testing items (segmental frequency is not balanced in the testing set, since the testing words are based off of real words), which are penalized by constraint d.




Figure 5: Average (red diamond) and distribution of scores assigned to nonce test words by the final grammar trained on Totonac dictionary words

The lack of a clearer harmony preference in the final grammar may be surprising, since the O/E of disharmonic forms is not terribly different on a strident projection than it is in a baseline trigram (see Table 16 above). However, the model does not use O/E to assess the usefulness of constraints, and the raw number of observed exceptions to harmony on the strident projection is apparently sufficient to prevent the final grammar from including general harmony constraints like $*[+anterior][+postalveolar]$ and $*[+postalveolar][+anterior]$. Instead, the model zeros in on a few specific disharmonic combinations and learns constraints against these segmental combinations only, as well as other constraints on stridents that don't distinguish harmonic from disharmonic combinations.

6.2.3 Web word corpus

The web word corpus of Navajo contains even more exceptions to strident harmony than the dictionary corpus, as shown in Tables 24 and 25. As in the dictionary corpus, however, the harmony preference is comparable for both anterior-postalveolar and postalveolar-anterior combinations.

Table 24: Observed/expected cooccurrence of anterior and postalveolar stridents in the Navajo web word corpus, transvocalic

	[+ant]	[-ant]
[+ant]	541/432 = 1.25	49/158 = 0.31
[-ant]	72/181 = 0.40	176/67 = 2.63

Table 25: Observed/expected cooccurrence of anterior and postalveolar stridents in the Navajo web word corpus, across any amount of intervening material

	[+ant]	[-ant]
[+ant]	1204/1038 = 1.16	189/355 = 0.53
[-ant]	203/369 = 0.55	292/126 = 2.32

The relatively weak harmony preference is consistent with Berkson's (2013) findings working with contemporary speakers and web data. In contrast, the dictionary word forms are carefully elicited and reflect data collected at an earlier time, unlike the more current and uncurated web data. It may be that harmony is falling out among contemporary speakers, and/or that it is represented inconsistently in written form (written web materials may also be produced by second language or heritage speakers).

The baseline grammar trained on the web word corpus doesn't include general placeholder trigram constraints that enforce harmony. This null result was determined through multiple runs of the model at a variety of parameter combinations. When gain is between 75 and 125 and gamma is 5 or 10, some trigram constraints on disharmonic combinations are found, and these are sometimes placeholder trigrams (other models include constraints with a medial gram that designates a subset of the vowels, e.g., [+long]). At a gain of 125 and a gamma of 5, for example, the baseline grammar includes *[+anterior, +strident][][+continuant, -voice, +postalveolar], which penalizes [ts ts' dz s z]-[]-[ʃ] trigrams. This constraint motivates the learner to construct a [+strident] projection, on which the constraints in Table 26 are found. These constraints penalize both harmonic and disharmonic strident combinations, and thus don't enforce harmony in any general sense.

Table 26: Constraints found on the [+strident] projection for the Navajo web words corpus, given in the order they were added to the grammar.

constraint	weight	sequences penalized
a. [-continuant][-continuant, -voice]	1.65	[ts ts' dz ʃ ʃ' dz]...[ts ts' ʃ ʃ']
b. [+anterior][+continuant, +postalv.]	3.16	[ts ts' dz s z]...[ʃ ʒ]
c. [+cont., -lat.][-cont., -voi., +postalv.]	4.39	[s z]...[ʃ ʃ']
d. [+cont., +voice, +postalv.][[]]	2.17	[ʒ]...[ts ts' dz s z ʃ ʃ' dz ʒ]
e. [-voice, +postalv.][+anterior]	2.56	[ʃ ʃ' ʒ]...[ts ts' dz s z]
f. [-cont., -voice][-cont., +voice]	1.66	[ts ts' ʃ ʃ']...[dʒ dʒ]
g. [-cont., +postalv.][-cont., +voi., +ant.]	2.88	[ʃ ʃ' dʒ]...[dʒ]

A set of testing forms was constructed following the established procedure. There were a total of 3970 testing forms, 278 with two anterior stridents, 1052 with two postalveolar stridents, 1320 with a postalveolar-anterior pair and 1320 with a anterior-postalveolar pair. The distribution of scores assigned to testing forms is shown in Figure 6, and demonstrates that the grammar does penalize disharmonic forms somewhat more than harmonic forms. There is a wide distribution of scores for all place combinations, however, as the grammar does not capture harmony in a general way. There are disharmonic combinations of stridents that violate no constraints (e.g., [dʒ...s] and [z...dʒ]), and forms with these combinations are given perfect scores just like some harmonic combinations.

The web word corpus, like the dictionary word corpus, does show a weak harmony preference among stridents that is captured by our learner. The grammars don't capture harmony in a general way, however, with constraints on minor-place combinations (as is done in the stems simulation). Instead, there are multiple constraints on a variety of segmental combinations. The overall result is slightly lower scores for disharmonic forms than harmonic forms on average, but the distribution of scores does not represent a clear harmonic preference.




Figure 6: Average (red diamond) and distribution of scores assigned to nonce test words by the final grammar trained on Navajo web words

6.3 Summary and discussion

The comparison of three Navajo corpora show that harmony is morphologically sensitive, holding much more strongly of stems than of words as a whole. As has been remarked on in previous work, harmony at the word level is characterized by numerous exceptions. In learning simulations, the frequency and patterning of exceptions at the word level is such that the model does not capture a general harmony preference, in contrast with the clear harmony pattern learned from the stem data. Even when the word grammars include a strident projection, the constraints on this projection do not capture harmony in a general way. Instead, the grammar includes a range of constraints on subsets of strident combinations, some of which distinguish harmonic from disharmonic combinations but most of which do not.

Only certain prefixes undergo harmony. In an unparsed word corpus, it may look instead like certain segments are more likely to undergo harmony, because of the frequency of these segments in alternating prefixes, and because of the frequency of those prefixes. For example, many of the alternating prefixes in Navajo contain [s]~[ʃ], and the grammar indeed includes a constraint on the strident projection against [s] followed by postalveolar strident. The Navajo simulations show that there are at least two potential consequences of morphological sensitivity for word-level phonotactics: the first is that exceptions may obscure the overall pattern, and the second is that harmony may appear to be a segment-specific quirk.

7 General discussion and conclusion

This paper set out to quantify the learnability of strident harmony patterns for an inductive phonotactic learner. The main conclusion is that word-level exceptions in languages with a morphologically sensitive system may make harmony unlearnable or harder to learn as a general phonotactic pattern over words. This is true in both Totonac, where harmony holds of stems but not words, and in Navajo, where only certain affixes participate in harmony. In both of these languages, however, harmony is observable and learnable as a trigram constraint, applying between very local stridents and reflecting the very strong harmony preference within stems.

The results show the limitations of looking at phonological patterns through a purely phonotactic lens. Based on the phonotactic structures in words, it may appear that some languages have very weak or even no strident harmony. Speakers, however, may have a strong harmony preference (or not), which is presumably learned from a morphologically informed analysis of the data. In languages where harmony

holds only of roots or stems, a phonotactic learner may be more successful if morpheme boundaries are represented in the learning data (e.g., the simulation of a root-bound laryngeal restriction in Aymara in Gallagher et al. 2019). Morpheme-specific alternations require a model of morphophonological learning (e.g., Albright & Hayes (2003) Gouskova & Becker (2013), Allen & Becker (2015)), that may represent knowledge of alternations as distinct from phonotactics.

The evaluation of several corpora for each of the three languages largely confirmed the descriptions of these languages available in the literature. The biggest discrepancy was found for Nkore-Kiga, where harmony was found to be featureally asymmetrical in the corpus. In both Totonac and Navajo, a robust harmony pattern was found in stems, with more exceptions in word forms. Comparison of web corpora and dictionary corpora found that web corpora have more exceptions to harmony; this is unsurprising considering that web corpora are uncurated as compared to dictionaries. It is unclear whether the conclusion should be that dictionary data are more reliable for phonological study or that dictionary data are less reflective of real-life language use than web data.

Another goal of the paper was to evaluate the generality of the Induction Projection Learner, which induces non-local projections based on local trigrams. The results show that this procedure is workable and may lead to reasonable final grammars, but is quite sensitive to exceptions. Given non-categorical patterns, the model often learns constraints that are overly specific and dance around the exceptions, and then fails to induce the appropriate non-local projection. Future work must improve on how the learner induces constraints, so that it is more robust to exceptions and avoids overfitting or pursue other methods for identifying non-local interactions from the baseline. On this point, the simulations above show that even when the placeholder trigram may be more difficult to find in the baseline grammar, the model is more robust to exceptions when searching for constraints on a projection.

The simulations reported here are also a set of case studies for the impact of parameter settings in the UCLA Phonotactic Learner. The success of the learner in capturing strident harmony (or any other property of a language's phonology) is dependent on the parameter settings given by the analyst. If gain and gamma are too low, the learned grammar includes hundreds of constraints and likely overfits the data. If gain and gamma are too high, the learned grammar will include very few constraints, and only those with zero or few exceptions, and will thus fail to capture many phonotactic patterns. Ideally, the appropriate gain and gamma parameters, as well as the success of a model, would be assessed based on fit to behavioral data that approximates speakers' full phonotactic grammar. The assumption made here is that the right grammar is the one that achieves separation in the range of scores assigned to harmonic and disharmonic forms, reflecting the hypothesis that a phonological pattern like strident harmony should have a bigger impact on wellformedness than other gradient phonotactic patterns that may reflect accidental gaps or underattestations, as supported by some behavioral studies (Hayes et al. 2009; Hayes & White 2013).

The gain and gamma combinations that were effective in the data sets analyzed in this paper varied considerably, as summarized here in Table 27. In general, it appears that smaller values are better for smaller data sets. Other factors that are likely relevant are the number of natural classes in a language and the simplicity or complexity of grammatical structures in the data set. For example, if the training data primarily includes CVC strings (e.g., because it is a training set of roots), there model does not need to evaluate and select among the large number of hypothetical constraints on consonant clusters that may be relevant in a training set with more varied syllable structure.

Table 27: Summary of parameters in the learning simulations that accounted for strident harmony.

training set	# of training items	gain	gamma	# of constraints
Nkore-Kiga stems	12,147	125	10	61
Totonac stems	1,814	30	5	47
Totonac words	7,480	60	5	73
Navajo stems	917	15	0	60
Navajo words	7830	125	5	77

In sum, this paper has shown that non-local phonology can be induced from the baseline via placeholder trigram constraints, so long as the pattern is robustly attested. Exceptions at the word level may make a general harmony constraint hard or impossible for the model to find as a baseline trigram, in which case the model will not include the strident projection and be unable to represent harmony in a general way. Future work must be concerned with elaborating on the integration of phonotactic and morphological learning, and on evaluating grammatical models with comparison to behavioral data.

References

Albright, Adam & Bruce Hayes. 2003. Rules vs. analogy in English past tenses: A computational/experimental study. *Cognition* 90. 119–161. [https://doi.org/10.1016/S0010-0277\(03\)00146-X](https://doi.org/10.1016/S0010-0277(03)00146-X).

Allen, Blake & Michael Becker. 2015. Learning alternations from surface forms with sublexical phonology. Manuscript. <https://ling.auf.net/lingbuzz/002503>.

Aschmann, Herman & Elisa Dawson de Aschmann. 1973. *Diccionario totonaco de Papantla, Veracruz: Totonaco-español, español-totonaco*. Dallas, TX: Summer Institute of Linguistics. <https://www.sil.org/resources/archives/29566>.

Bennett, William G. & Douglas Pulleyblank. 2018. Directionality in Nkore-Kiga sibilant harmony: Arbitrary or emergent? *Linguistic Inquiry* 49(1). 1–22. https://doi.org/10.1162/ling_a_00264

Berent, Iris, Colin Wilson, Gary Marcus & David Bemis. 2012. On the role of variables in phonology: Remarks on Hayes and Wilson 2008. *Linguistic Inquiry* 43. 97–119. https://doi.org/10.1162/LING_a_00075.

Berkson, Kelly Harper. 2013. Optionality and locality: Evidence from Navajo sibilant harmony. *Laboratory Phonology* 4(2). 287–337. <https://doi.org/10.1515/lp-2013-0011>.

Daland, Robert, Bruce Hayes, James White & Marc Garellek. 2011. Explaining sonority projection effects. *Phonology* 28.197–234. <https://doi.org/10.1017/S0952675711000145>.

Della Pietra, Stephen, Vincent Della Pietra & John Lafferty. 1997. Inducing features of random fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 19. 380–393.

Eddington, David & Jordan Lachler. 2006. A computational analysis of Navajo verb stems. In Sally Rice & John Newman (eds.), *Empirical and experimental methods in cognitive/functional research*. Stanford, CA: CSLI Publications.

Gafos, Adamantios. 1999. *The articulatory basis of locality in phonology*. New York: Garland.

Gallagher, Gillian, Maria Gouskova & Gladys Camacho Rios. 2019. Phonotactic restrictions and morphology in Aymara. *Glossa: A Journal of General Linguistics* 4(1). 29. <http://doi.org/10.5334/gjgl.826>.

Gouskova, Maria & Gillian Gallagher. 2020. Inducing nonlocal constraints from baseline phonotactics. *Natural Language & Linguistic Theory* 38(1). 77–116. <https://doi.org/10.1007/s11049-019-09446-x>.

Hansson, Gunnar Olafur. 2001. *Theoretical and typological issues in consonant harmony*. Berkeley, CA: University of California dissertation.

Hayes, Bruce & Colin Wilson. 2008. A maximum entropy model of phonotactics and phonotactic learning. *Linguistic Inquiry* 39. 379–440. <https://doi.org/10.1162/ling.2008.39.3.379>.

Hayes, Bruce & James White. 2013. Phonological naturalness and phonotactic learning. *Linguistic Inquiry* 44. 45–75. https://doi.org/10.1162/LING_a_00119.

MacKay, Carolyn Joyce. 1999. *A grammar of Misantla Totonac*. Salt Lake City: University of Utah Press.

Martin, Andrew. 2011. Grammars leak: Modeling how phonotactic generalizations interact within the grammar. *Language* 87(4). 751–770. <https://doi.org/10.1353/lan.2011.0096>.

McDonough, Joyce. 2003. *The Navajo sound system*. Dordrecht: Kluwer Academic Publishers.

Morris, Henry Francis & Brian Edmond Renshaw Kirwan. 1957. *A Runyankore grammar*. Nairobi: East African Literature Bureau.

Sapir, Edward & Harry Hoijer. 1967. *The phonology and morphology of the Navajo language*. Berkeley, CA: University of California Press.

Shaw, Patricia. 1991. Consonant harmony systems: the special status of coronal harmony. In Carole Paradis & Jean-François Prunet (eds.), *The special status of coronals: Internal and external evidence*, 125–127. San Diego: Academic Press.

Taylor, Charles. 1959. *A simplified Runyankore-Rukiga-English and English-Runyankore-Rukiga dictionary*. Kampala: Eagle Press.

Wilson, Colin & Gillian Gallagher. 2018. Accidental gaps and surface-based phonotactic learning: A case study of South Bolivian Quechua. *Linguistic Inquiry* 49(3). 610-623.
https://doi.org/10.1162/ling_a_00285.

Young, Robert & William Morgan. 1972. *The Navaho language*. Salt Lake City: Deseret Book Company.
<https://archive.org/details/TheNavahoLanguage/>.

Young, Robert & William Morgan. 1987. *The Navajo Language: A grammar and colloquial dictionary*. Albuquerque: University of New Mexico Press.

Young, Robert W., William Morgan, & Sally Midgette. 1992. *Analytical lexicon of Navajo*. Albuquerque: University of New Mexico Press.

Gillian Gallagher
Department of Linguistics
New York University
10 Washington Place
New York, NY 10003 USA
ggillian@nyu.edu