

# Is Sour Grapes Learnable? A Computational and Experimental Approach

Brandon Prickett

*University of Massachusetts Amherst*

## 1 Introduction

Vowel harmony is a typologically common phonological pattern that restricts which kinds of vowels can appear together in a surface form, often causing all of a word’s vowels to agree in their value for one or more features (Rose & Walker, 2011). Harmony patterns are also commonly used in artificial language learning studies (see, e.g., Pycha et al., 2003; Finley, 2008; Moreton, 2012; Lai, 2015), with the focus of most studies often being a comparison between the relative learnability of attested and unattested generalizations. One such unattested pattern—and the focus of this paper—is Sour Grapes Harmony (Bakovic, 2000; Wilson, 2006b; Finley, 2008; Lin & Myers, 2010).

Attested harmony patterns often involve blocker segments, which “block” harmony in the sense that different values of the relevant feature are allowed on either side of the blocker (Rose & Walker, 2011:§3.3.3). For example, if a backness harmony pattern was blocked by [a], the form [tipitaku] would be grammatical, despite the fact that [u] and [i] have different values for the feature [back]. Crucially, in a standard pattern like this, all of the vowels to one side of a blocker (e.g. the first two [i]’s in this example) still share the same value for the harmonizing feature.

In Sour Grapes, disharmonic sequences of vowels are generally not allowed in surface forms, just like attested harmony patterns. However, when a blocker segment is present in a word, otherwise ungrammatical sequences can appear either to its left or right, depending on the kind of Sour Grapes harmony being implemented (Bakovic, 2000; Wilson, 2006b). For example, in a left to-right Sour Grapes pattern, segments to the left of a blocker can disagree in their value for the relevant feature. So, if backness was harmonized in such a pattern, a form like \*[tiputu] would be ungrammatical, but a form like [tiputaku] would be allowed, because, in the latter word, the [a] licenses the vowels to its left to be disharmonic.

Sour Grapes is predicted by a number of constraint based (Prince & Smolensky, 1993) theories of long distance assimilation (Bakovic, 2000; Wilson, 2006b), with different explanations for its absence from phonological typology. These explanations usually take one of two approaches: they either revise the set of constraints used to capture harmony so that Sour Grapes cannot be represented in the phonological grammar (e.g., Wilson, 2006b; McCarthy, 2011) or focus on the formal complexity of Sour Grapes and work toward a theory of phonology that categorically forbids any patterns that are at that level of complexity (e.g., Heinz, 2018; Smith & O’Hara, 2019; Lamont, 2019). Crucially, both of these approaches seek to categorically limit phonological grammars so that they lack the expressive power to ever capture Sour Grapes harmony. However, past experimental work has struggled to find evidence that Sour Grapes is categorically unlearnable (Finley, 2008; Lin & Myers, 2010).

In this paper, I present results from simulations using three different maximum entropy phonotactic models (Hayes & Wilson, 2008; Moreton et al., 2017): one that can only represent Sour Grapes, one that can only represent standard, attested harmony, and one that has the expressive power to capture both patterns. I then present results from an experiment designed to test the predictions of these models and find that humans behave most like the model that can capture both generalizations—challenging the idea that Sour Grapes is

---

\* Thanks to Joe Pater, Gaja Jarosz, John Kingston, Mohit Iyyer, Katya Pertsova, Andrew Lamont, and the audiences of WCCFL 2021 and AMP 2022 for helpful discussion on the topics in this paper.

categorically unlearnable.

The paper proceeds as follows: §2 discusses past theoretical and experimental work on Sour Grapes harmony, §3-6 describe an experiment designed to test the predictions of maximum entropy models that vary in their ability to represent harmony patterns and reports the experiment's results, and §7 concludes.

## 2 Background

**2.1 Theoretical Work on Sour Grapes** Several optimality theoretic (Prince & Smolensky, 1993) approaches to harmony, such as AGREE constraints (Bakovic, 2000; Wilson, 2006b), predict Sour Grapes in systems that include a blocker segment. This is because constraints like AGREE penalize words with blocker segments, regardless of whether vowels to one side of the blocker are harmonized. For example, the words [tiputaku] and [tipitaku] would both incur a violation of AGREE(back)—the former because of the [i...u] sequence in the first two syllables and the latter because of the [i...a] sequence in the second and third syllables. This means that both forms would be expected in a language that has a constraint enforcing [a]'s role as a blocker (such as \*[-high, -back]) ranked or weighted more highly than AGREE, and AGREE ranked or weighted more highly than a constraint enforcing faithfulness to underlying values of [back]. However, Sour Grapes is an unattested phonological pattern, inspiring a range of research attempting to explain this absence.

Wilson (2006b) proposed that the reason for Sour Grapes' unattestedness is that attested harmony patterns are myopic and Sour Grapes fails to meet this criterion. That is, to determine whether a word is grammatical according to attested patterns, one only needs to look at each of its vowels two at a time, with no need to check pairs of non-adjacent vowels or to look at more than two vowels at once. Since blockers in Sour Grapes license strings of disharmonic vowels of arbitrary length, one needs to look at a potentially infinite window of vowels to determine whether a word is grammatical, to ensure a blocker isn't at its end. Wilson (2006b) suggested that a variant of optimality theory (Prince & Smolensky, 1993) that uses targeted constraints (Wilson, 2001), would solve this issue by only predicting myopic harmony patterns. Targeted constraints do this by spreading the harmonizing feature in a word one vowel at a time and never looking ahead to non-adjacent vowels.

McCarthy (2011) proposed SHARE constraints, paired with Harmonic Serialism (McCarthy, 2000) as a similar method for avoiding non-myopic harmony patterns, including Sour Grapes. SHARE constraints penalize any adjacent pair of segments that fail to share the same feature value. For example, the form [tiputaku] would incur three violations of the constraint SHARE(back), because no pair of adjacent vowels has the same value for the feature [back] (assuming [a] is a central vowel). But the form [tipitaku] would only incur two violations of the SHARE constraint, since the [i...i] pair in the word's first two syllables does have the same value for the crucial feature. This, combined with Harmonic Serialism's limited GEN and a constraint motivating [a]'s status as a blocker, allows attested blocking patterns while categorically preventing Sour Grapes from being represented in the grammar.

However, these approaches have been critiqued, since a number of non-myopic phonological patterns do seem to exist (Walker, 2010; Jardine, 2016; McCollum & Essegbey, 2018; Stanton, 2018). Since targeted constraints and SHARE constraints both categorically restrict a grammar so that it can only represent myopic harmony, these patterns are evidence that myopia is not the crucial factor causing Sour Grapes to be unattested.

Another approach for explaining Sour Grapes' absence is Formal language Theory (FLT; Chomsky, 1956). FLT is a way of describing how complex a pattern is in terms of the computational machinery needed to represent it. The framework was originally designed to demonstrate that natural language syntax was more complex than the set of regular patterns (i.e. those that could be represented using finite state machines). However, Johnson (1972) showed that all known phonological mappings could be considered regular. Recent work has supported this finding, showing that attested phonological patterns can be characterized as subregular (Heinz, 2018), and suggesting that this is due to a categorial restriction on phonological learning (Heinz, 2010; Heinz & Idsardi, 2013; Jardine & Heinz, 2016).

Specifically, it has been argued that all phonological patterns can be characterized as either Strictly Local or Tier-based Strictly Local (TSL; Heinz et al., 2011). The former level of complexity includes any pattern that bans a finite set of substrings occurring in a word, while the latter includes any pattern that does so over a tier of segments (i.e., certain classes of segments can be ignored by the pattern). An example of a Strictly

Local pattern that commonly occurs in natural language is the typologically common restriction banning voiceless sounds after nasals (henceforth  $*NC_{\sigma}$ ; Pater, 1999). This pattern is Strictly Local since it bans any word containing the finite set of strings that results from combining all nasals with all voiceless sounds (e.g. [nt], [np], [mt], [mp], etc.). TSL patterns are also common in phonology and most harmony patterns belong to this region of the subregular hierarchy (Heinz et al., 2011; Heinz, 2018). Figure 1 shows the full Subregular Hierarchy and where each of these two types of patterns are located on it.

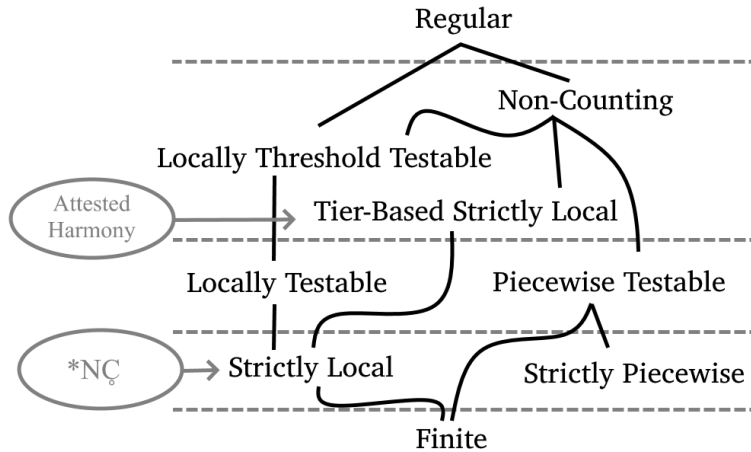


Figure 1. The Subregular Hierarchy, with examples of Strictly Local and Tier-based Strictly Local patterns given. Dashed grey lines indicate different orders of logic. From lowest to highest, these are: Conjunction of Negative Literals, Propositional, First Order, and Monadic Second Order. Solid black lines indicate subset relationships.

While Sour Grapes is still a subregular pattern, it does not reside in the same region of the subregular hierarchy as standard, attested harmony patterns (Heinz, 2018; Smith & O’Hara, 2019; Lamont, 2019). This is because all TSL patterns can be defined by a set of banned substrings, but Sour Grapes allows almost any sequence of segments on one side of a blocker. For example, a standard backness harmony pattern in a language with three vowels in its inventory ([i], [u], and [a]) could be represented as a ban on the sequences [iu] and [ui] on the vowel tier. This would make words like \*[tipu] and \*[tipuka] ungrammatical, but words like [tupukati] grammatical, since the former two have [ui] and [iu] vowel bigrams in their vowel tier, while the latter only has [uu], [ua], and [ai]—all of which are allowed.

However, an analogous left-to-right Sour Grapes pattern cannot be represented so easily. The sequence [iu] would be allowed in a word like [tipuka], but not allowed in the word [tipu], because only the former has an [a] to license the disharmonic pair [iu]. Lamont (2019) showed that Sour Grapes is located in the non-counting region of the subregular hierarchy, which is more complex than TSL. This is illustrated in Figure 2.

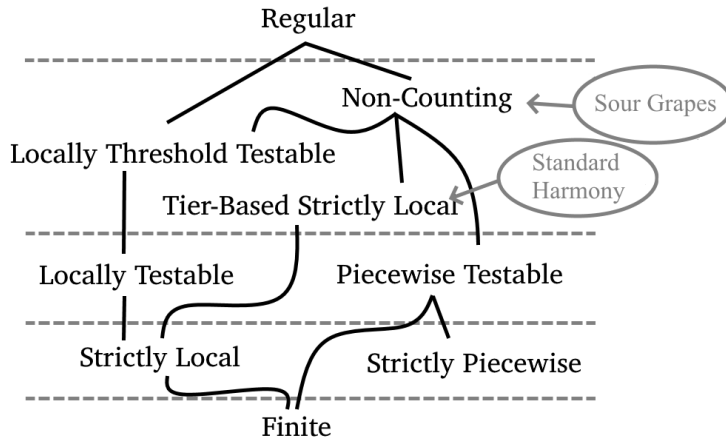


Figure 2. The Subregular Hierarchy, with the complexity of Sour Grapes and attested harmony shown.

Lamont (2019) also shows that several non-myopic harmony patterns that have been compared to Sour Grapes in the past (e.g., tone in Copperbelt Bemba; Jardine, 2016) belong to less complex regions of the subregular hierarchy. This means that FLT provides a way to categorically differentiate between Sour Grapes and attested harmony patterns. Learning algorithms have been developed that are limited to learning TSL languages (see, e.g., Jardine & Heinz, 2016; Jardine & McMullin, 2017) and any theory that used these would correctly predict a lack of Sour Grapes in phonological typology.

**2.2 Past Experimental Approaches to Sour Grapes** The learnability of Sour Grapes relative to attested harmony patterns has also been explored in experimental work, however neither of the past studies that explored the topic were able to explain the pattern’s typological absence. Artificial language learning studies usually seek to show that unattested patterns are more difficult to learn or generalize in the lab than their attested counterparts (see Moreton & Pater, 2012a, 2012b for a review of such experiments) and use this difficulty as a way to explain their typological absence. The idea behind such studies is that if a difficult pattern arises diachronically, it might be quickly replaced by an easier one or never properly phonologized at all by the community using the language (Moreton, 2008).

Finley (2008) attempted to train participants on phonological alternations that were ambiguous between Sour Grapes and an attested harmony pattern in an artificial language learning design to see whether participants generalized as if they had learned the latter. This poverty of the stimulus method of artificial language learning experiment has been used effectively to demonstrate a learning bias away from an unattested pattern (e.g., Wilson, 2006a; Finley & Badecker, 2009). However, Finley’s (2008) participants were unable to learn either Sour Grapes or the attested harmony—possibly because the combination of blockers and alternations introduced too much complexity into the learning task for participants to handle over the course of a short experiment.

Lin and Myers (2010) also attempted to explain Sour Grapes’ absence using an artificial language learning experiment. Specifically, they used nasal harmony patterns that were blocked by [s] and [k]. Their participants were native speakers of Taiwan Southern Min, a language with phonemic contrasts for nasality on both consonants and vowels and thus were able to properly perceive a novel pattern that spread nasality. Lin and Myers (2010) trained half of these participants on an attested nasal harmony pattern and the other half on Sour Grapes, then compared the groups’ accuracies in a separate testing phase. While they did find a marginal difference between the accuracies in the two conditions, the difference suggested that Sour Grapes was easier to learn than its attested counterpart. Lin and Myers (2010) proposed that Sour Grapes’ absence could be due to factors other than learnability, such as limitations introduced by the phonetic origins of harmonic phonological patterns.

### 3 Design

To test whether humans have a preference against Sour Grapes patterns, I designed an experiment that combined the two previous attempts to find such a bias (Finley, 2008; Lin & Myers, 2010). That is, like Finley’s (2008) study, I trained participants on data that is ambiguous between a Sour Grapes and attested harmony language (a poverty of the stimulus design). However, following Lin and Myers (2010), I will only be presenting participants with the surface forms of the language—meaning that they will not be exposed to any phonological alternations. The goal of this hybridization is to see whether the poverty of the stimulus approach might find evidence of a learning bias against Sour Grapes that more easily explains its typological absence, while hopefully presenting participants with a simpler learning scenario than Finley (2008), so that they succeed in learning the patterns in the first place.

Sour Grapes is more permissive than attested harmony patterns with blockers—that is, the set of surface forms attested harmony allows is always a subset of the ones allowed in a minimally different Sour Grapes pattern. This is because Sour Grapes allows all of the surface forms that are grammatical in attested harmony, as well as words with disharmonic sequences that occur to one side of a blocking segment. This is illustrated in the Venn diagram in Figure 3, for left-to-right backness harmony.

The experiment I describe here consisted of two parts: a training phase and a testing phase. In the training phase, participants were presented with words that were grammatical in both Sour Grapes and attested harmony. The feature that was harmonizing was [back] and the vowel inventory of the language consisted of [i], [u], and [a] (with [a] acting as the blocker for the harmony process). The consonantal inventory was made

up of [p], [t], [tʃ], and [k] and all words used one of the following consonantal templates: [p...t...tʃ], [k...tʃ...t], [tʃ...k...p], [t...p...k], [p...t...tʃ...k], [k...tʃ...t...p], [tʃ...k...p...t], or [t...p...k...tʃ]. All syllables were CV, meaning that all words were either three or four syllables long.

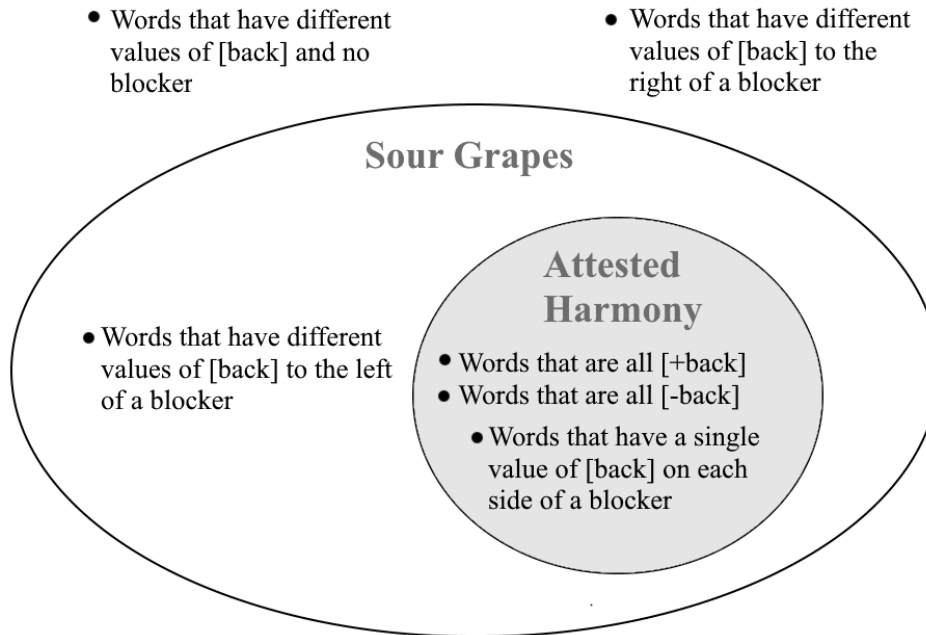


Figure 3. Venn diagram demonstrating the SRs that are grammatical in Sour Grapes and Attested Harmony, when [back] is the feature being spread from left to right. Bulleted items describe the words that are grammatical in each language, with descriptions outside of both circles being banned by both patterns. For examples of each of these types from my experiment, see Table 1 and Table 2.

There were four types of trials in the training phase, each of which was grammatical in both Sour Grapes and attested harmony and conveyed different kinds of information about the pattern being learned. The first trial type was Faithful, which consisted of words that had at most one [i] or [u], with [a] filling in the rest of the word's vowel positions. These words taught participants little about the pattern, and mostly served as fillers in the experiment. The next two trial types were Harmonic-[i] and Harmonic-[u]. These presented words that had greater than one [i] (and no [u]'s) or greater than one [u] (and no [i]'s), respectively. These demonstrated that words with multiple high vowels that agreed in their value for [back] were allowed in the language. The fourth trial type in training was Disharmonic, which presented words that had at least one [i] and at least one [u], with an [a] in between them. These demonstrated the fact that [i] and [u] could cooccur in a word, despite their different values for [back], as long as there was an [a] in the middle of them. Examples from each of the trial types in training are shown in Table 1.

Table 1. Examples of the stimuli presented in the Sour Grapes experiments' training and testing phases.

	Example Stimulus	Example Stimulus (Vowel Tier)
Faithful	[tapukatʃa]	[a...u...a...a]
Harmonic-[i]	[tipikitʃi]	[i...i...i...i]
Harmonic-[u]	[tupukutʃu]	[u...u...u...u]
Disharmonic	[tupukatʃi]	[u...u...a...i]

The testing phase presented participants with novel words that belonged either to the trial types from training or to two additional trial types: Ungrammatical Both and Ungrammatical-AH trials. These novel trial types were designed to test whether participants learned a Sour Grapes pattern or an attested harmony pattern from the ambiguous data in training. Words in the Ungrammatical-Both trials would be banned by both Sour

Grapes and attested harmony, meaning they contained an [iu] or [ui] pair in their vowel tier without an [a] to the right of the disharmonic sequence. Ungrammatical-AH trials consisted of words that were ungrammatical in attested harmony but allowed by Sour Grapes. These words also contained an [iu] or [ui] sequence in their vowel tier, but there was always an [a] somewhere to the right of it, licensing the disharmonic string. Examples of these additional trial types are shown in Table 2.

Table 2. Examples of the stimuli presented in the Sour Grapes experiments' testing phase that were not present in training.

	Example Stimulus	Examples Stimulu (Vowel Tier)
<i>Ungrammatical-Both</i>	[tipukutʃu]	[i...u...u...u]
<i>Ungrammatical-AH</i>	[tipukatʃu]	[i...u...a...u]

## 4 Methods

The consonantal templates described in §3 were combined with all relevant vowel sequences to create a pool of possible stimuli. Sets of stimuli were randomly sampled for each participant such that no individual saw the same item more than once in training, but across individuals, words could be reused. A broad transcription for each stimulus was fed through the text-to-speech software, Tacotron 2 (Shen et al., 2018) to convert the transcriptions into .wav files. These audio files were the only version of the stimuli that participants ever encountered, with no orthographic stimuli being presented at any point.

The experiment took an average of 29 minutes to complete and participants (N=40) were paid \$7.50 for their time (UMass Amherst IRB protocol number 2017-4040). All participants were recruited using Prolific (www.prolific.co) with the following text advertising the experiment on the website:

You will be learning the words of a new, made-up language. At first, you'll participate in a training phase in which you'll hear words, and then be asked to repeat them. Later there will be a testing phase that will ask you to judge whether new words sound like they belong to the language of interest.

After agreeing to participate on Prolific, participants were forwarded to the experiment, which was hosted on Ibex Farm (<http://spellout.net/ibexfarm/>). The first page of the experiment presented them with the following instructions (which were mostly a repetition of the instructions on Prolific), as well as an informed consent form:

In this experiment you will be learning the words of a new, made-up language. At first, you'll participate in a training phase in which you'll hear words, and then be asked to repeat them. Later there will be a testing phase that will ask you to judge whether new words sound like they belong to the language of interest. Be sure to be using headphones so that you can carefully listen to each recording. Also, please use a computer (as opposed to a smart phone) so that all of the audio plays clearly.

They were then taken to a page that played a non-linguistic audio file, which gave them an opportunity to ensure their computer's audio was working properly. Next, the training phase of the experiment began, which presented 30 stimuli from each of the four training trial types discussed in §3, in a randomized order. Each trial began with a page that said "Please listen to the following word" and played an audio file of the relevant stimulus. A new page would then appear that said "Now repeat", giving participants a chance to repeat the stimulus out loud. After the training phase, a new set of instructions was presented:

Nice work! You've completed the training phase of the experiment. Next you will participate in the testing phase. In this phase, you will be asked to judge whether new words sound like they belong to the language you've been learning. Be sure to let the full word play before selecting an option. Your response times are being recorded.

Each trial in testing presented participants with a stimulus and then asked them "Does this word sound like it could belong to the language?". They were given the options "Yes" and "No" and were not given any

feedback about their answers. The stimuli in the testing phase included 5 words for each trial type repeated from training (20 total), 10 novel stimuli from each of the four training trial types (40 total), and ten novel items from each of the two testing trial types described in §3.

## 5 Predictions

To better understand what predictions models might make for this experiment, depending on whether they have the expressive power to represent Sour Grapes, I implemented three maximum entropy phonotactic learners (Hayes & Wilson, 2008; Moreton et al., 2017): one with a constraint that represented Sour Grapes harmony (violated by any word that would appear in the Ungrammatical-Both trials), one with a constraint that represented the attested harmony pattern (violated by any word that would appear in either the Ungrammatical-Both or Ungrammatical-AH trials), and a model that had both of these constraints. Each learner also had a base set of constraints that allowed it to penalize words containing specific strings of vowels that were four or three segments long (e.g., \*uiu, \*iiuu, \*iuii, etc.). This was meant to represent the possibility that participants were simply memorizing the words in the experiment’s training phase and favoring words in testing that had similar segments in their vowel tier.

Each of the models was trained for 10 separate runs on a set of data that was randomly produced in the same manner as the stimuli in the human participants’ training phase (described in §4). That is, 30 unique words were randomly produced for each of the 4 training trial types and then given to each of the models. All simulations used batch gradient descent with a learning rate of .01 and initial weights of 1.0. The predictions discussed below are the model’s probability estimates after being trained for a single epoch (i.e., one full pass through the data and a single weight update), since this corresponds to the same amount of exposure to stimuli that the participants were given in the experiment. However, these results are qualitatively identical to results after 50 epochs—a more standard amount of exposure to give to a maximum entropy learner.

The test items were also randomly produced in the same manner as the human participants’ stimuli (described in §4) and were divided into three categories: words that were ungrammatical in both Sour Grapes and attested harmony (\*SG,\*AH below, analogous to Ungrammatical-Both trials), words that were grammatical in Sour Grapes but ungrammatical in its attested counterpart (SG,\*AH below, analogous to Ungrammatical-AH trials), and words that were grammatical according to both patterns (SG,AH below, analogous to the trial types found in training). The models’ average probabilities for these three categories of stimulus after a single epoch of training are shown in Figure 4.

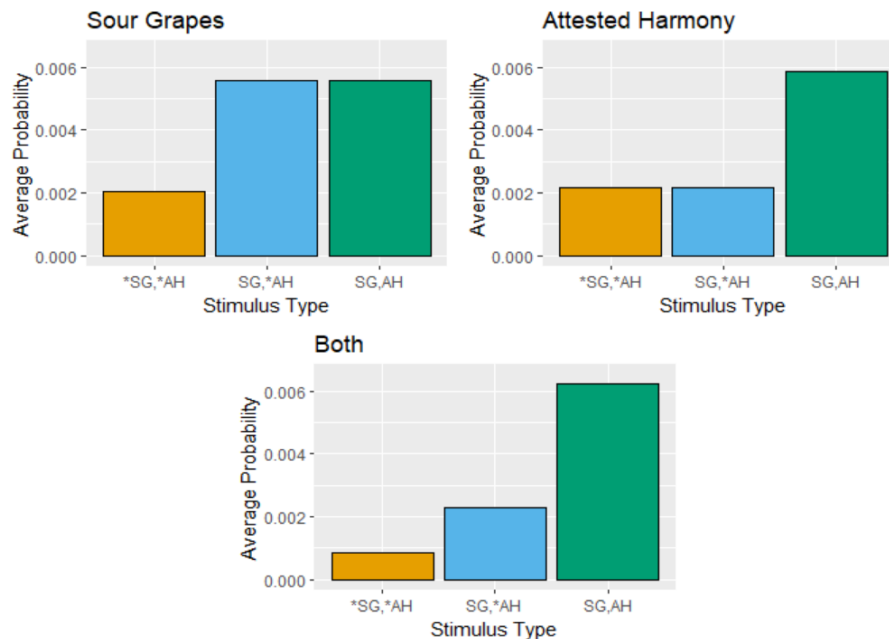


Figure 4. Probability estimates from the three maximum entropy phonotactic models for test data after a single epoch of training, averaged over 10 unique runs for each model.

The results from the maximum entropy learners show that the three different kinds of models each make unique predictions for the experiment. The model that can only represent Sour Grapes predicted no difference in the acceptability between the SG,\*AH and SG,AH words, but predicts that \*SG,\*AH will be significantly less acceptable. This is because that model’s constraint enforcing Sour Grapes is blind to the difference between SG,\*AH and SG,AH forms, assigning no violations to either type. The model that can only represent attested harmony predicts that there will be no difference between the \*SG,\*AH words and the SG,\*AH words, but that SG,AH words will be significantly more acceptable than both of the other categories. Similarly, this is because the constraint enforcing attested harmony is blind to difference between \*SG,\*AH and SG,\*AH words, assigning violations to both equally. And finally, the model that can represent both Sour Grapes and attested harmony predicts a three way contrast with \*SG,\*AH words at the bottom, SG,\*AH words in the middle, and SG,AH words at the top, in terms of acceptability.

## 6 Results

The proportion of “Yes” responses for the human participants, by stimulus type, is shown in Figure 5. On average, participants were least likely to judge \*SG,\*AH stimuli as belonging to the language they had been trained on. They were more likely to judge SG,\*AH as belonging to the language—however, SG,AH stimuli had the highest likelihood of receiving a “Yes” response from participants.




Figure 5. Proportion of ‘Yes’ responses in the testing phase of the Sour Grapes experiment, by stimulus type.

To ensure these differences were unlikely to arise from chance, I ran a logistic regression on the data using the lme4 (Bates et al., 2015) package in R (R Core Team, 2016). The stimulus type of each participant’s response was given to the model using the coding scheme in Table 4. This allowed for a comparison of participants’ likelihood of saying “Yes” to SG,\*AH items with their likelihood of saying “Yes” to the other two trial types. The predictor SG,\*AH vs. \*SG,\*AH compared SG,\*AH items with \*SG,\*AH ones, while the predictor SG,\*AH vs. SG,AH compared SG,\*AH and SG,AH items. The model was also given random effects of participant and item on the intercept and the two slopes, and its results are shown in Table 5.

Table 4. Coding for the Sour Grapes experiment.

	<i>*SG,*AH vs. SG,*AH</i>	<i>SG,*AH vs. SG,AH</i>
<i>*SG,*AH</i>	-.333	-.333
<i>SG,*AH</i>	.666	-.333
<i>SG,AH</i>	-.333	.666



Table 5. Results from the logistic regression in the Sour Grapes experiment.

Predictor	Estimate	Std. Error	z-value	p-value
Intercept	1.26	0.25	5.08	3.84e-07***
SG, *AH vs. *SG, *AH	-0.51	0.20	-2.61	0.00914**
SG, *AH vs. SG, AH	0.41	0.17	2.35	0.01865*

Both predictors in this model were found to be significant ( $p < .01$  for SG,\*AH vs. \*SG,\*AH and  $p < .05$  for SG,\*AH vs. SG,AH), meaning that the observations about Figure 5 discussed above were unlikely to have arisen from chance. Specifically, SG,\*AH items were significantly more likely to get “Yes” responses from participants than \*SG,\*AH stimuli, but significantly less likely to get “Yes” responses than SG,AH items.

These results are consistent with a model that can represent both Sour Grapes and attested harmony patterns, according to the simulation results presented in §5 and are consistent with the evidence that Lin and Myers’ (2010) found using a different methodology that suggested both patterns were learnable.

## 7 Conclusions

Despite being representable in constraint-based theories of harmony, Sour Grapes patterns are absent from phonological typology. Past attempts to explain this have sought to design phonological theories that categorically exclude Sour Grapes from being representable in the phonological grammar, either by restricting constraint sets (Wilson, 2006b; McCarthy, 2011) or by focusing on the computational complexity of the pattern (Heinz, 2018; Smith & O’Hara, 2019; Lamont, 2019). However, past experimental work (Lin & Myers, 2010) has failed to find a bias against learning Sour Grapes in artificial language learning.

Here I showed that a maximum entropy model that had the expressive power to represent both Sour Grapes and an attested harmony pattern could correctly predict human generalization in an artificial language learning experiment with a poverty of the stimulus design. Future work should test such a model to see if it can help explain Sour Grapes’ typological absence despite the model’s ability to represent the pattern.

## References

- Bakovic, Eric (2000). *Harmony, dominance and control* [PhD Thesis].
- Bates, Douglas, Maechler, Martin, Bolker, Ben, Walker, Steven, Christensen, Rune H. B., Singmann, Henrik, Dai, Bin, & Grothendieck, Gabor (2015). Package ‘lme4.’ *Convergence*, 12(1).
- Chomsky, Noam (1956). Three models for the description of language. *IRE Transactions on Information Theory*, 2(3), 113–124.
- Finley, Sara (2008). *Formal and cognitive restrictions on vowel harmony* [PhD Thesis].
- Finley, Sara, & Badecker, William (2009). Artificial language learning and feature-based generalization. *Journal of Memory and Language*, 61(3), 423–437.
- Gafos, Adamantios, & Dye, Amanda (2011). Vowel harmony: Transparent and opaque vowels. *The Blackwell Companion to Phonology*, 4, 2164–2189.
- Gorman, Kyle (2013). *Generative phonotactics*. University of Pennsylvania.
- Hayes, Bruce, & Wilson, Colin (2008). A maximum entropy model of phonotactics and phonotactic learning. *Linguistic Inquiry*, 39(3), 379–440.
- Heinz, Jeffrey (2010). Learning long-distance phonotactics. *Linguistic Inquiry*, 41(4), 623–661.
- Heinz, Jeffrey (2018). The computational nature of phonological generalizations. *Phonological Typology, Phonetics and Phonology*, 126–195.
- Heinz, Jeffrey, & Idsardi, William. (2013). What complexity differences reveal about domains in language. *Topics in Cognitive Science*, 5(1), 111–131.
- Heinz, Jeffrey, Rawal, Chetan, & Tanner, Herbert G. (2011). Tier-based strictly local constraints for phonology. *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies: Short Papers-Volume 2*, 58–64.
- Jardine, Adam (2016). Computationally, tone is different. *Phonology*, 33(2), 247–283.
- Jardine, Adam, & Heinz, Jeffrey (2016). Learning tier-based strictly 2-local languages. *Transactions of the Association for Computational Linguistics*, 4, 87–98.
- Jardine, Adam, & McMullin, Kevin (2017). Efficient learning of tier-based strictly k-local languages. *International Conference on Language and Automata Theory and Applications*, 64–76.
- Johnson, C. Douglas (1972). *Formal aspects of phonological description*. Mouton & Co. N.V.
- Lai, Regine (2015). Learnable vs. Unlearnable harmony patterns. *Linguistic Inquiry*, 46(3), 425–451.

- Lamont, Andrew (2018). Precedence is Pathological: The Problem of Alphabetical Sorting. *Proceedings of the 36th West Coast Conference on Formal Linguistics*. <http://blogs.umass.edu/alamont/files/2019/03/WCCFL-2018-paper.pdf>
- Lamont, Andrew (2019). *Sour Grapes is phonotactically complex*. Linguistic Society of America. <https://blogs.umass.edu/alamont/files/2019/01/LSA-2019-handout.pdf>
- Lin, Yu-Leng, & Myers, James (2010). *Testing Universal Grammar in phonological artificial grammar learning* [PhD Thesis]. Master's thesis. National Chung Cheng University, Taiwan.
- Mailhot, Frédéric (2013). Modeling the emergence of vowel harmony through iterated learning. *Origins of Sound Change: Approaches to Phonologization*, 247–61.
- McCarthy, John J. (2000). Harmonic serialism and parallelism. *Linguistics Department Faculty Publication Series*, 40.
- McCarthy, John J. (2011). Autosegmental spreading in optimality theory. *Tones and Features: Phonetic and Phonological Perspectives*, 195–222.
- McCollum, Adam G., & Essegbey, James (2018). Unbounded Harmony Is Not Always Myopic: Evidence from Tutrugbu. In W. G. Bennett, L. Hrats, & D. R. Storoshenko (Eds.), *Proceedings of the 35th West Coast Conference on Formal Linguistics* (pp. 251–258).
- McMullin, Kevin J. (2016). *Tier-based locality in long-distance phonotactics: Learnability and typology* [PhD Thesis]. University of British Columbia.
- Moreton, Elliott (2008). Analytic bias and phonological typology. *Phonology*, 25(1), 83–127.
- Moreton, Elliott (2012). Inter- and intra-dimensional dependencies in implicit phonotactic learning. *Journal of Memory and Language*, 67(1), 165–183.
- Moreton, Elliott, & Pater, Joe (2012a). Structure and Substance in Artificial-phonology Learning, Part I: Structure. *Language and Linguistics Compass*, 6(11), 686–701.
- Moreton, Elliott, & Pater, Joe (2012b). Structure and substance in artificial-phonology learning, part II: Substance. *Language and Linguistics Compass*, 6(11), 702–718.
- Moreton, Elliott, Pater, Joe, & Pertsova, Katya (2017). Phonological Concept Learning. *Cognitive Science*, 41(1), 4–69.
- Pater, Joe (1999). Austronesian nasal substitution and other NC effects. *The Prosody-Morphology Interface*, 310–343.
- Prince, Alan, & Smolensky, Paul (1993). *Optimality Theory: Constraint Interaction in Generative Grammar*. Blackwell Publishing.
- Prince, Alan, & Tesar, Bruce (2004). Learning phonotactic distributions. *Constraints in Phonological Acquisition*, 245–291.
- Pycha, Anne, Nowak, Pawel, Shin, Eurie, & Shosted, Ryan (2003). Phonological rule-learning and its implications for a theory of vowel harmony. *WCCFL*, 22, 423–435.
- R Core Team (2016). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing. <http://www.R-project.org/>
- Rose, Sharon, & Walker, Rachel (2011). Harmony Systems. In J. Goldsmith, J. Riggle, & A. C. L. Yu (Eds.), *The Handbook of Phonological Theory* (2nd ed., pp. 240–290). Wiley-Blackwell.
- Shen, Jonathan, Pang, Ruoming, Weiss, Ron J., Schuster, Mike, Jaitly, Navdeep, Yang, Zongheng, Chen, Zhifeng, Zhang, Yu, Wang, Yuxuan, & Skerrv-Ryan, Rj (2018). Natural tts synthesis by conditioning wavenet on mel spectrogram predictions. *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 4779–4783.
- Smith, Caitlin, & O'Hara, Charlie (2019). Formal Characterizations of True and False Sour Grapes. *Proceedings of the Society for Computation in Linguistics*, 2(1), 338–341.
- Stanton, Juliet (2018). Gurindji nasal cluster dissimilation as trigger deletion. *Journal of Linguistics*, 1–39. <https://doi.org/10.1017/S0022226718000506>
- Walker, Rachel (2010). Nonmyopic harmony and the nature of derivations. *Linguistic Inquiry*, 41(1), 169–179.
- Wilson, Colin (2001). Consonant cluster neutralisation and targeted constraints. *Phonology*, 18(1), 147–197.
- Wilson, Colin (2006a). Learning phonology with substantive bias: An experimental and computational study of velar palatalization. *Cognitive Science*, 30(5), 945–982.
- Wilson, Colin (2006b). Unbounded spreading is myopic. *Current Perspectives on Phonology Workshop, Phonology Fest*.