Phonology (2025), **Vol. 42**:e24 1–31
doi:10.1017/S0952675725100213

PHONOLOGY

ARTICLE

A large-scale corpus study of phonological opacity in Uyghur

Connor Mayer

Department of Language Science, University of California, Irvine, USA
Email: cjmayer@uci.edu

Received: 16 August 2022; **Revised:** 25 August 2023; **Accepted:** 16 September 2024

Keywords: corpus phonology; opacity; maximum entropy; Uyghur; paradigm uniformity

Abstract

This article examines a case of phonological opacity in Uyghur resulting from an interaction between backness harmony and a vowel reduction process that converts harmonic vowels into transparent vowels. A large-scale corpus study shows that although opaque harmony with the underlying form of a reduced vowel is the dominant pattern, cases of surface-apparent harmony also occur. The rate of surface-apparent harmony varies across roots and is correlated with a number of factors, including root frequency. These data pose problems for standard accounts of opacity, which do not predict such variation. I propose an analysis where variation emerges from conflict between a paradigm uniformity constraint mandating that the harmonising behaviour of a root remains consistent, and surface phonotactic constraints. This is implemented in a parallel model by scaling constraint violations according to certainty in a root's harmonic class. This aligns with past work suggesting some opacity is driven by paradigm uniformity.

Contents

1. Phonological opacity **2**
2. Opacity in Uyghur backness harmony **3**
2.1. Segments involved in backness harmony . . . 4
2.2. A description of Uyghur backness harmony . . . 4
2.3. Vowel reduction . . . 5
2.4. Opaque interactions between backness harmony and vowel reduction . . . 6
2.5. Modelling opacity in serial and parallel models . . . 6
2.6. Past work on opacity in Uyghur . . . 8
3. A corpus study of opacity in Uyghur backness harmony **9**
3.1. Parsing the corpora . . . 10
3.2. Quantitative results . . . 10
3.3. Predicting opacity . . . 12
4. An analysis of opacity in the corpus **15**
4.1. Phonological and lexical effects in backness harmony . . . 15
4.2. Modelling gradience in opacity using Harmonic Uniformity . . . 16
4.3. Calculating $P(\text{HC}|x)$ . . . 17
4.4. Validating the model . . . 18
5. Discussion **20**
5.1. Generalisation . . . 21
5.2. Pathological roots . . . 21
5.3. Alternative analyses using serial models . . . 22

© The Author(s), 2025. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

https://doi.org/10.1017/S0952675725100213 Published online by Cambridge University Press

5.4. Limitations of the corpus 23
5.5. A closing remark on corpus methodologies 23
A. The Uyghur morphological transducer **23**
A.1. Modifying the transducer 24
A.2. Interpreting and sanitising the transducer output 24
B. Additional details on statistical models **25**
C. Maximum Entropy Optimality Theory **27**
D. *P*(HC = *B*|*x*) model coefficients **27**
References **28**

This article presents an exploratory written corpus study of a phonological pattern in Uyghur (Turkic; China) whereby a vowel reduction process converts harmonic vowels into transparent vowels, rendering the vowel harmony pattern opaque. Opaque patterns are of interest to phonological theory because of the challenges they pose for learning and for certain classes of phonological models. In particular, opacity has been a persistent difficulty for strictly parallel phonological models such as classical Optimality Theory (Prince & Smolensky [1993] 2004), which do not straightforwardly predict its existence.

The article has two primary goals. The first is to put claims about opacity in Uyghur on a stronger empirical footing by presenting new data. Uyghur provides a valuable opportunity to study an opaque pattern at scale using text data, because both vowel reduction and harmony are reflected orthographically. The data show that although opaque harmony is the majority pattern, there is variability in rates of opacity both within and between roots: roots sometimes trigger surface-apparent harmony (McCarthy 1999).[1] Rates of opacity are correlated with factors like token frequency and the presence of certain derivational suffixes.

The second goal is to highlight challenges these data pose for standard theories of opacity, which do not predict such variability. I will outline an analysis that is compatible with the data, where observed variation emerges as the result of conflict between lexical knowledge of the harmonising class of a root and sensitivity to surface phonotactic constraints. In addition to accounting for the various correlates of opacity observed in the corpus, this analysis can be implemented in a strictly parallel model using a constraint that mandates paradigm uniformity in the harmony system (Rebrus & Törkenczy 2017, 2021; Rebrus *et al.* 2023). This aligns with previous proposals that lexical factors play a fundamental role in opaque phenomena (Łubowicz 2003; Mielke *et al.* 2003; Sanders 2003; Green 2004; Pater 2010; Nazarov 2019).

§1 provides background on phonological opacity. §2 describes the processes of backness harmony and vowel reduction in Uyghur, and how they interact to produce opacity. §3 presents the results of a large written corpus study looking at rates of opacity. §4 presents an analysis of the corpus data. Finally, §5 discusses implications for theories of opacity in general, some limitations of the work presented here, alternative analyses, and how future research might proceed.

1. Phonological opacity

Opacity is a type of structured phonological exceptionality. Kiparsky (1971, 1973) defines it as follows:

(1) Assume a phonological rule ℝ: A → B / C_D. ℝ is *opaque* iff there are surface forms with either:
   a. *A* in the environment *C*_*D* (underapplication opacity);
   b. *B* derived from *A* in environments other than *C*_*D* (overapplication opacity).

[1] I use the term 'root' in this article to refer to monomorphemic roots as well as polymorphemic stems with only derivational suffixes. This is in part because the morphological transducer that will be described below typically treats these derived stems as single roots, and because derivational suffixes will often impose their own harmonic properties, while inflectional suffixes inherit the harmonic properties of the stem. I hope the reader forgives this slight abuse of terminology in the interests of avoiding awkward prose.

https://doi.org/10.1017/S0952675725100213 Published online by Cambridge University Press

Opacity arises when either a conditioned alternation appears not to occur despite its conditions being met, or appears to occur when its conditions have not been met.

Kiparsky (1973) associates opacity of types (1a) and (1b) with *counterfeeding* and *counterbleeding* rule orders, respectively. In counterfeeding opacity, the structural conditions for rule $\mathbb{R}$ to apply are created by a different rule $\mathbb{P}$ that applies after $\mathbb{R}$; hence, the necessary conditions are not met when $\mathbb{R}$ applies. Changing the rule ordering such that $\mathbb{P}$ applies before $\mathbb{R}$ would produce a *feeding* order where $\mathbb{R}$ applies transparently to the conditioning environment produced by $\mathbb{P}$.

In counterbleeding opacity, the conditions for $\mathbb{R}$ are met when it applies, but are subsequently altered by a different rule $\mathbb{Q}$ that applies after $\mathbb{R}$. Changing the rule ordering such that $\mathbb{Q}$ applies before $\mathbb{R}$ would produce a *bleeding* order where $\mathbb{R}$ transparently fails to apply because $\mathbb{Q}$ removes its conditioning environment.

More recently, interest in opacity has stemmed from debates on the merits of serial models such as SPE-style rules (Chomsky & Halle 1968) *vs.* parallel models such as Optimality Theory (Prince & Smolensky [1993] 2004). Parallel models have difficulty correctly predicting cases of counterbleeding opacity, which generally produce faithfulness violations with no corresponding markedness repairs to motivate them. They also have difficulty with most types of counterfeeding opacity, which fail to repair a markedness violation whose repair is evident elsewhere.

A number of theoretical mechanisms have been proposed to handle these cases, including mechanisms that incorporate some degree of serialism into OT, such as sympathy (McCarthy 1999), Stratal OT (Kiparsky 2000; Bermúdez-Otero 2003; Nazarov & Pater 2017; Bermúdez-Otero 2018), candidate chain theory (McCarthy 2007) and serial markedness reduction (Jarosz 2014), as well as purely parallel mechanisms, such as constraint conjunction (Kirchner 1996), paradigm uniformity (Steriade 2000), language-specific constraints (Pater 2014) or indexed constraints (Pater 2010; Nazarov 2019, 2020, 2021). The need for such bespoke mechanisms has been seen as a point in favour of serial models, which handle these cases of opacity without issue (e.g., Vaux 2008).

Although counterfeeding and counterbleeding orderings are the best known configurations that result in opacity, the typologies of opacity enumerated in Baković (2007, 2011) and Baković & Blumenfeld (2019) show that these orderings are neither sufficient nor necessary conditions for opacity. They identify a number of cases of overapplication opacity that are not predicted by SPE-style rule ordering, and some which are only able to be described by parallel models. Thus, the characterisation of opacity as a unique challenge for parallel models is a simplification, though accurate in broad strokes.

In light of the lack of a unified account of opacity from either serial or parallel theories, Baković (2007) suggests that the field focus on Kiparsky's claim that opaque patterns are more difficult to learn than transparent ones. The basic motivation for this claim is that phonological processes that interact in an opaque fashion make it difficult to generalise about these processes: opaque forms constitute exceptions to otherwise robust generalisations. Kiparsky (1971) supports this claim by presenting a number of cases of historical change where an opaque process is reanalysed as a non-opaque one.

Subsequent research has presented evidence that opaque processes are learned as phonemic contrasts or lexicalised patterns rather than productive rules (e.g., Hooper 1976; Mielke *et al.* 2003; Sanders 2003; Sumner 2003; Bowers 2019; Zhang 2019), though evidence also exists that some opaque processes are applied productively in language games and other contexts (e.g., Donegan & Stampe 1979; Al-Mozainy 1981; Vaux 2011) as well as in behavioural experiments (Farris-Trimble & Tessier 2019).

2. Opacity in Uyghur backness harmony

Uyghur is a southeastern Turkic language spoken by over 12 million people in the Xinjiang Uyghur Autonomous Region in the People's Republic of China, neighbouring countries, such as Kazakhstan and Kyrgyzstan, and various diasporic communities (Engesæth *et al.* 2010; Nazarova & Niyaz 2013). It has SOV word order with highly agglutinative morphology that is almost exclusively suffixing.

https://doi.org/10.1017/S0952675725100213 Published online by Cambridge University Press

Table 1. *Harmonising segments in Uyghur*

	Front		Back	
	Unrounded	Rounded	Unrounded	Rounded
High	i	y̲		u̲
Mid	e	ø̲		o̲
Low	æ̲		ɑ̲	

(a) The Uyghur vowel system. Harmonising vowels are underlined

	Front	Back
Voiceless	k	q
Voiced	g	ʁ

(b) Harmonising consonants

The opaque phenomenon under consideration arises from the interaction of two independent processes: backness harmony and vowel reduction. I will introduce these processes separately before demonstrating how their interaction leads to opacity. The reader is referred to Mayer *et al.* (2022b) for a more detailed description of Uyghur phonology.

2.1. Segments involved in backness harmony

Like most Turkic languages, Uyghur has backness harmony (Lindblad 1990; Hahn 1991a, b; Abdulla *et al.* 2010; Engesæth *et al.* 2010). Harmony is most evident in alternations between suffix allomorphs, where, broadly speaking, segments in the suffix must agree in backness with the rightmost vowel of the roots they attach to. It may also be observed to a lesser extent in static root forms (particularly native Turkic roots), but extensive borrowing has led to many disharmonic roots.

Segments that participate in backness harmony are shown in Table 1. The underlined vowels in Table 1a act as harmony triggers (i.e., they determine the backness of suffixes attached to roots containing them), while the non-underlined vowels are transparent to harmony. The harmonising consonants in Table 1b may also serve as harmony triggers, though they tend to be weaker than vowels. This article will focus primarily on harmony driven by vowels. In addition to serving as triggers of harmony, the harmonising vowels and consonants both emerge as the outcome of harmony in harmonising suffixes.

2.2. A description of Uyghur backness harmony

The examples of harmony below include the locative suffix /-DA/ (surface forms: [-tɑ], [-dɑ], [-tæ] and [-dæ]), the plural suffix /-lAr/ (surface forms: [-lɑr] and [-lær]), or the dative suffix /-GA/ (surface forms: [-qɑ], [-ʁɑ], [-kæ] and [-gæ]). I assume that /A/ is unspecified for the feature [back], /D/ for [voice] and /G/ for both (Archangeli 1988). Voicing alternations in the initial segment are caused by voice assimilation, and are orthogonal to harmony.

The basic characterisation of backness harmony is that suffixes must agree in backness with the rightmost harmonising root vowel, front /y ø æ/ or back /u o ɑ/.

(2) *Simple front harmonising forms*

a. ty̲r-dæ̲ ‘type-LOC’
b. pæ̲n-læ̲r ‘science-PL’
c. munbæ̲r-gæ̲ ‘podium-DAT’

(3) *Simple back harmonising forms*

a. pu̲l-ʁɑ̲ ‘money-DAT’
b. to̲p-lɑ̲r ‘ball-PL’
c. ætrɑ̲p-tɑ̲ ‘surroundings-LOC’

https://doi.org/10.1017/S0952675725100213 Published online by Cambridge University Press

The vowels /i e/ are *transparent* to harmony. They do not serve as harmony triggers, but allow the harmonic value of preceding segments to 'pass through' them.

(4) *Front roots with transparent vowels*

a.	mæstʃit-tæ	'mosque-LOC'
b.	ymid-lær	'hope-PL'
c.	mømin-gæ	'believer-DAT'

(5) *Back roots with transparent vowels*

a.	student-lar	'student-PL'
b.	uniwersitet-tɑ	'university-LOC'
c.	ɑmil-ʁɑ	'element-DAT'

Roots without any harmonising segments typically take back suffixes, but some take front suffixes (see McCollum 2021; Mayer *et al.* 2022a).

(6) *Neutral roots that take back suffixes*

a.	sir-lar	'secret-PL'
b.	din-ʁɑ	'religion-DAT'
c.	hejt-tɑ	'festival-LOC'
d.	peʔil-lar	'verb-PL'
e.	tip-qɑ	'type-DAT'

(7) *Neutral roots that take front suffixes*

a.	biz-gæ	'us-DAT'
b.	bilim-gæ	'knowledge-DAT'
c.	welisipit-lær	'bicycle-PL'

Back suffixes appear to be the unmarked class in Uyghur. There has been a general diachronic shift in the population of neutral roots towards back suffixes (Lindblad 1990), and recent loanwords that lack harmonising segments typically take back suffixes (Mayer *et al.* 2022b).

2.3. Vowel reduction

The second process that contributes to opacity in the Uyghur harmony system is *vowel reduction* or *raising*, which raises the low vowels /ɑ æ/ to [i] in medial open syllables in derived environments.[2]

(8) *Reduction of* /ɑ/

a.	bɑlɑ	'child'	bɑli-ʁɑ	'child-DAT'
b.	ɑpɑ	'mom'	ɑpi-si	'mom-3.POSS'
c.	ɑŋlɑ-ʃ	'listen-GER'	ɑŋli-ʁɑn	'listen-PFV'
d.	qɑrɑ-ŋ	'look-IMP'	qɑri-di	'look-3SG.PST'

[2] An analysis of Uyghur vowel reduction is beyond the scope of this article, but past OT analyses of derived environment effects in vowel reduction have relied on Comparative Markedness (Khanjian 2009; Mascaró 2009) or local constraint conjunction (Łubowicz 2002); the latter paper mentions Uyghur raising in passing (fn. 20) but does not provide a full analysis. McCollum (2020) and Mayer (2021b: Appendix E) provide analyses of vowel reduction in Uyghur, but neither attempts to account for derived environment effects.

https://doi.org/10.1017/S0952675725100213 Published online by Cambridge University Press

(9) *Reduction of* /æ/

a.	ɑpæ̲t	'disaster'	ɑpi̲t-i	'disaster-3.POSS'
b.	æte̲	'tomorrow'	æti̲-gæ	'tomorrow-DAT'
c.	søzlæ̲-ŋ	'talk-IMP'	søzli̲-gæn	'talk-PFV'
d.	kyʧæ̲-ʃ	'strive-GER'	kyʧi̲-di	'strive-3SG.PST'

The underlying form cannot in general be predicted from forms where vowel reduction could have applied, as many words have underlying /i/ in these positions, as in /tɑksi/ 'taxi' or /æsli/ 'origin'. Certain roots resist raising categorically, particularly loanwords where the final vowel was long in the source language (Nazarova & Niyaz 2013); in the current article, I focus on roots that undergo raising.

2.4. Opaque interactions between backness harmony and vowel reduction

Vowel reduction has the potential to introduce opaque behaviour into the vowel harmony system. Consider, for example, the root /ɑʁinæ/ 'friend'. The final vowel undergoes raising when it occurs in a derived word-medial open syllable:

(10) /ɑʁinæ-ni/ → [ɑʁini-ni] 'friend-ACC'

What happens when the vowel in the suffix must harmonise with the final vowel in the root, like in the form /ɑʁinæ-DA/ 'friend-LOC'? There are two possibilities: *opaque harmony*, according to the underlying form of the root, and *surface-apparent harmony*, according to the raised form of the root. We will set aside for a moment the question of which of these we actually see in Uyghur, and briefly explore some theoretical consequences of each realisation.

A rule ordering where harmony precedes raising predicts the opaque form [ɑʁinidæ]:

(11) *Harmony precedes raising*

UR	/ɑʁinæ-DA/
Harmony	ɑʁinæ̲-dæ̲
Raising	ɑʁini-dæ
SR	[ɑʁini-dæ]

This opacity is precisely the kind that classical OT has difficulty accounting for: there is an explicit markedness violation (failure to harmonise), with no apparent motivation (cf. forms like /t**ɑ**ksi-DA/ → [t**ɑ**ksid**ɑ**] 'taxi-LOC').

If raising instead precedes backness harmony, we would expect the form [ɑʁini-dɑ] with surface-apparent harmony:

(12) *Raising precedes harmony*

UR	/ɑʁinæ-DA/
Raising	ɑʁini-DA
Harmony	ɑ̲ʁini-dɑ̲
SR	[ɑʁini-dɑ]

This outcome can be predicted by both serial and parallel models.

2.5. Modelling opacity in serial and parallel models

The kind of opacity shown in (11) is straightforward to represent in serial rule-based models: the rule that drives harmony is simply ordered before the rule that drives raising. An analysis under such a model is fundamentally identical to the derivation in (11).

https://doi.org/10.1017/S0952675725100213 Published online by Cambridge University Press

This pattern poses challenges for an analysis in classical OT. I will assume a simple markedness constraint that motivates vowel harmony, which is a combination of the local and non-local AGREE constraints used by Hayes *et al.* (2009):

(13) VAGREE: Harmonising segments in a suffix must match the backness of the rightmost harmonising vowel in the stem.

The following constraints will drive raising:

(14) *UNREDUCED: Don't have lowvowels in derived medial open syllables.

(15) ID[HEIGHT]: Don't changethe height of segments in the input.

*UNREDUCED is shorthand for a more detailed analysis of the pressures that drive vowel reduction (for vowel reduction in general, see Crosswhite 2001 and de Lacy 2002; for vowel reduction in Uyghur, see McCollum 2020 and Mayer 2021b: Appendix E).

When it is relevant, I will also employ a constraint that prevents specified [back] values from being altered:

(16) ID[BACK]: Don't change the backness of segments in the input.

ID[BACK] prevents underlyingly specified vowels in roots and certain harmony-blocking suffixes from being altered. This constraint is not violated when a segment underspecified for backness in the input is assigned a backness value in the output, nor is it violated when /æ ɑ/ are raised to [i]: assuming that [i] is unspecified for backness, these processes violate DEP and MAX constraints, respectively, which are low-ranked and omitted from the tableaux below.

These constraints allow Classical OT to derive only surface harmony, as shown in (17), where the sad face indicates the candidate that should have won, and the bomb indicates the candidate that does win. Suppose that we want to derive opaque harmony for the suffixed form /ɑʔilæ-lAr/ 'family-PL'. The desired candidate [ɑʔili-lær] is harmonically bounded by the winning *[ɑʔili-lɑr], and so will never be the optimal candidate under any ranking.

(17) *Failed tableau for [aʔili-lær] 'family-PL'*

/ɑʔilæ-lAr/	*UNREDUCED	VAGREE	ID[HEIGHT]
☹ a. ɑʔili-lær		*!	*
💣 b. ɑʔili-lɑr			*
c. ɑʔilæ-lær	*!		
d. ɑʔilæ-lɑr	*!	*	

An analysis using Stratal OT succeeds in capturing this opacity (Kiparsky 2000; Bermúdez-Otero 2003, 2018). Stratal OT divides the grammar into several strata (e.g., the stem, the word and the phrase) and assigns each of these levels a separate OT grammar with differing constraint rankings. The outputs of lower strata serve as the inputs to higher strata. We can capture the opaque pattern here by proposing that vowel harmony occurs at a lower stratum than vowel reduction, as in (18). At the word stratum, shown in (18a), the constraint driving raising is ranked below its corresponding faithfulness constraint, meaning harmony applies but raising does not. The output from the word stratum serves as the input to the tableau for the phrase stratum in (18b). At this stratum, the constraint driving raising is now ranked above its corresponding faithfulness constraint, meaning raising can apply.

https://doi.org/10.1017/S0952675725100213 Published online by Cambridge University Press

(18) *Tableaux for the derivation of [aʔili-lær] 'family-PL'*

a. *Word stratum*

/ɑʔilæ-lAr/	ID[BACK]	ID[HEIGHT]	VAGREE	*UNREDUCED
i. ɑʔili-lær		*!	*	
ii. ɑʔili-lɑr		*!		
☞ iii. ɑʔilæ-lær				*
iv. ɑʔilæ-lɑr			*!	*

b. *Phrase stratum*

/ɑʔilæ-lær/	ID[BACK]	*UNREDUCED	VAGREE	ID[HEIGHT]
☞ i. ɑʔili-lær			*	*
ii. ɑʔili-lɑr	*!			*
iii. ɑʔilæ-lær		*!		
iv. ɑʔilæ-lɑr	*!	*	*	

I employ this formalism here because it is widely used in the contemporary literature, and because there is evidence in Uyghur that raising can apply at the level of the phrase, suggesting that it belongs to a higher stratum than backness harmony (specifically, it appears to be a post-lexical process; Kiparsky 1982). For example, in phrases like *Adil Hesenge berdi* 'Adil gave it to Hesen', the dative *-ge* [-gæ] may raise to *-gi* [-gi] in rapid speech (Hahn 1991b: 53).

To summarise, rule-based analyses predict opaque harmony straightforwardly, while strictly parallel analyses predict only surface-apparent harmony. Modifications to strictly parallel models that incorporate some degree of serialism, such as those listed in §2.1, also predict opaque harmony, though they differ in their attribution of the particular mechanism responsible for it. For example, while Stratal OT captures opaque patterns by positing multiple derivational strata with different constraint rankings, candidate chain theory (McCarthy 2007) does so by evaluating candidate chains (roughly analogous to derivations) rather than candidates.

2.6. Past work on opacity in Uyghur

Which of these patterns do we observe in Uyghur? Pedagogical materials do not generally discuss these cases in any detail, since roots that can generate opaque harmony are a relatively small slice of the lexicon. Those that do discuss them suggest that opacity is the correct outcome (e.g., Hahn 1991b: §4.3.5). Hahn describes this in terms of roots falling into a particular 'harmonic category', with vowel reduction processes 'disguising' the most salient clue to this category: the final vowel of the root.

As is typical for opaque phenomena (Kiparsky 1971, 1973; Mielke *et al.* 2003), the rule ordering that produces opaque harmony reflects the relative diachronic development of each process. Backness harmony is an ancient property of Turkic languages (e.g., Clauson 1972), while raising is a newer phenomenon in Uyghur; Chagatay, the closest direct ancestor to Uyghur, appears to have had no such raising process (Bodrogligeti 2001). Opaque harmony thus maintains historical patterns of root backness at the cost of surface disharmony.

The exception to this is that certain derivational suffixes in Uyghur, such as the diminutive /-ʧæ/ and the adjectival suffix /-ɑnæ/, have been described as triggering surface-apparent harmony in raised forms (Hahn 1991b; Halle *et al.* 2000; Vaux 2000; Hall & Ozburn 2019):[3]

[3] /-ʧæ/-final stems are sometimes repaired to be internally harmonic: e.g., the root /bɑʁʧɑ/ 'park' (lit. 'orchard-DIM') is frequently produced as [bɑʁʧɑ].

https://doi.org/10.1017/S0952675725100213 Published online by Cambridge University Press

(19) *Surface-apparent harmony in raised* /-ʧæ/

/bɑʁ-ʧæ-DA/	→	[bɑ̲ʁʧid̲ɑ̲]	'park-DIM-LOC'
cf. /bɑʁ-ʧæ-m-DA/	→	[bɑʁʧæ̲mdæ̲]	'park-DIM-1SG.POSS-LOC'

We will return to these suffixes below.

Uyghur speakers I have worked with agree that opaque harmony is the correct outcome. However, in addition to the suffixes above, speakers have identified certain forms where surface-apparent harmony is mandatory (e.g., /ærzɑn-i-GA/ → [ærzinigæ] 'cheap (ones)-3.POSS-DAT'), or where both surface-apparent and opaque harmony are acceptable (e.g., /æzɑn-i-GA/ → [æziniʁɑ]/[æzinigæ] 'call to prayer-3.POSS-DAT'). These forms demonstrate that although opaque harmony occurs in the vast majority of cases, we cannot always predict whether a root or stem will trigger surface-apparent or opaque harmony. These few elicited observations of variability were one of the motivations for the corpus study presented below.

Theoretical work on the interaction between vowel reduction and harmony has claimed that there is an asymmetry between vowels (Halle *et al.* 2000; Vaux 2000; Hall & Ozburn 2019): raised /æ/ is opaque and continues to behave as a front vowel trigger (with the suffixes above constituting notable exceptions), while raised /ɑ/ is transparent, behaving identically to underlying /i/. However, these claims have been based on only eight data points collected from a single speaker, and the empirical validity of these data is unclear (see Mayer 2021b: §3.3.7). One of the main goals of this article is to put claims about opacity in Uyghur on a stronger empirical footing.

In addition to answering an important empirical question, obtaining a better understanding of this pattern is valuable from a theoretical perspective: opaque patterns, such as the one in (11), are not predicted to exist by many strictly parallel phonological models, and, indeed, the Uyghur pattern has been used to argue in favour of serial models (Vaux 2000). The remainder of this article will present a large-scale corpus study that examines the empirical facts of opacity in Uyghur and explore some of the theoretical implications of its results.

3. A corpus study of opacity in Uyghur backness harmony

In order to investigate the interaction of vowel reduction and backness harmony, I performed a corpus study using three large text corpora.[4]

Uyghur uses three different orthographies depending on where it is written: Perso-Arabic, Cyrillic or Latin. In each of these, the alternations conditioned by the raising and harmony processes are represented orthographically.[5] Hence, text corpora allow us to gather large-scale empirical data on their interaction.

The first corpus was generated from the Uyghur-language Web site of Radio Free Asia (RFA; https://www.rfa.org/uyghur/). RFA is a U.S.-sponsored non-profit news organisation. The second was generated from the Web site of *Uyghur Awazi* (*Uyghur Voice*; https://uyguravazi.kazgazeta.kz/), an Uyghur-language newspaper published in Almaty, Kazakhstan. The third was generated from *Uyghur Akadémiyisi* (Uyghur Academy; https://www.akademiye.org/ug/), a legal research organisation that publishes articles on Uyghur culture and politics.

Corpora were generated from the Web sites using Web scrapers – software that, given a starting URL, instructions for how to navigate between pages, and instructions for which information to retrieve from each page, can download content from all pages on a site, or multiple sites. Such programs allow corpora to be generated from publicly available Internet resources, in formats that are useful to researchers.

[4] The code used in this article can be found at https://github.com/connormayer/uyghur_corpora.

[5] There is one specific exception to this: the editorial guidelines of Radio Free Asia, one of the Web sites from which the corpora were generated, require proper names to be written in their unraised forms, even when they undergo raising. For example, the word for 'American' is written as *Amérikalïq* even though it is pronounced [ɑmerikɨ̲-liq]. This results in a slight undercounting of raised forms in the following sections, but this does not affect the conclusions of this article.

https://doi.org/10.1017/S0952675725100213 Published online by Cambridge University Press

Table 2. *Summary of corpora*

Corpus	Article count	Word count	Date retrieved	Orthography
Uyghur Voice	24,080	4,197,550	January 2020	Latin/Cyrillic
RFA	45,777	7,986,484	July 2022	Perso-Arabic
Uyghur Academy	2,635	2,423,519	July 2022	Perso-Arabic

There are separate Web scrapers for each of the Web sites, which are linked in the GitHub repository for this article (see fn. 3). These scrapers were written by undergraduate research assistants at the Los Angeles and Irvine campuses of the University of California, in collaboration with the author.

A summary of the contents of each corpus is shown in Table 2.[6] In addition to the contents of each article, the scrapers retrieved the author, the date and the URL.

3.1. Parsing the corpora

In order to extract information about the interaction between backness harmony and vowel reduction from the corpus, I modified an existing Uyghur morphological transducer to detect the backness of suffix forms (https://github.com/apertium/apertium-uig; Littell *et al.* 2018; Washington *et al.* 2019). This transducer is part of Apertium, a free and open-source rule-based machine translation platform (https://www.apertium.org).

The transducer maps from surface forms to underlying analyses that consist of roots plus morphological tags indicating the backness of any harmonising suffixes. For example, if the input is the surface form *qizingizgha* 'to your daughter', the output analysis will be *qiz*<n><px2sg><frm><dat-b>. This indicates that the root is *qiz*, a noun <n>, and that it is suffixed with the second person singular possessive marker in its formal form *-ingiz* <px2sg><frm>,[7] followed by the dative suffix in its back form *-gha* <dat-b>.

The output of the transducer was used to count the frequency of front or back suffixes for each root. Words for which the transducer was unable to produce a valid parse were excluded from analysis. Simple text processing comparing the parsed root and surface forms was used to detect whether vowel reduction occurred and to extract phonotactic properties of the roots and tokens.

Additional details of the transducer and data processing, including numerical validation, are presented in Appendix A.

3.2. Quantitative results

3.2.1. Comparing harmonic and disharmonic roots. In this section, I consider only tokens where (a) the rightmost two harmonising elements of the root are vowels; (b) the underlying final vowel in the root is either /æ/ or /ɑ/; (c) the final vowel undergoes raising; and (d) the raised vowel is followed by at least one harmonising suffix. The rightmost two harmonising vowels were chosen as the domain of analysis because the rightmost harmonising vowel in a stem almost invariably predicts suffix backness (see Mayer 2021b: ch. 4). When this vowel is reduced to transparent [i], the second-rightmost harmonising vowel has the potential to influence suffix backness. The effect of any preceding harmonising vowels on suffix backness appears to be negligible.

Figure 1 breaks down parsed roots into four classes according to their rightmost two harmonising vowels: two back vowels (BB; 774 roots; e.g., /bɑlɑ/ 'child', /mɑʃinɑ/ 'vehicle, machine'), a back

[6]The earlier scrape date for the Uyghur Awazi corpus is due to a redesign that changed the structure of the site and pages, which requires modifications to the scraper. It was not possible to make these modifications in time for the submission of this article.

[7]Uyghur has formal and informal versions of many second person morphemes. The corresponding informal version of the second person singular possessive marker in the above example would be *-ing*.

https://doi.org/10.1017/S0952675725100213 Published online by Cambridge University Press

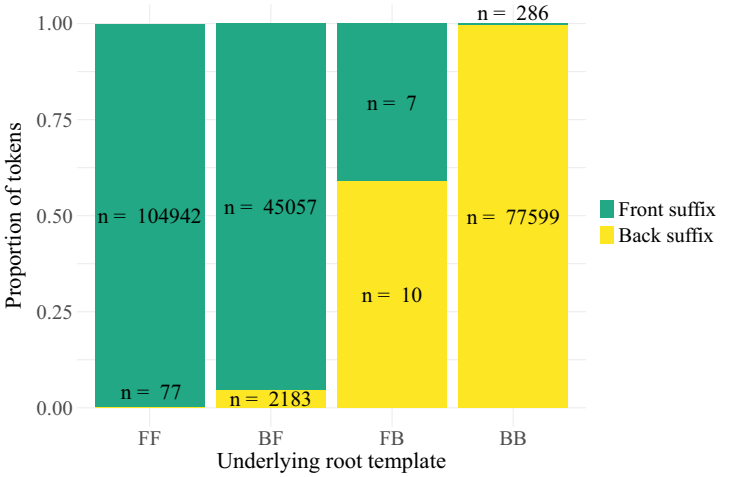


Figure 1. *Suffix harmony choice in tokens where the final root vowel raises, broken down by root class. Token counts are overlaid on each category.*

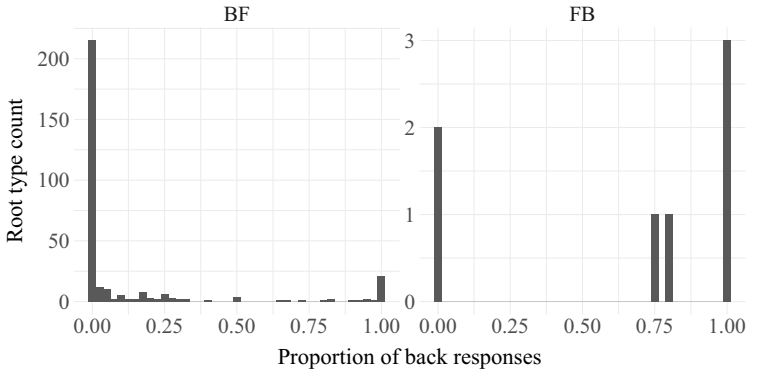


Figure 2. *Histograms showing the distribution of rates of back suffix application in BF and FB roots. Note that for raised BF roots, a back suffix constitutes surface-apparent harmony and a front suffix constitutes opaque harmony, while for raised FB roots, it is the opposite.*

vowel followed by a front vowel (BF; 311 roots; e.g., /ɑdæt/ 'custom', /ɑʔilæ/ 'family'), a front vowel followed by a back vowel (FB; 7 roots; e.g., /ærzɑn/ 'cheap', /kæsipdɑʃ/ 'colleague'), and two front vowels (FF; 528 roots; e.g., /sypæt/ 'quality', /mæsilæ/ 'problem'). The BF and FB classes have the potential to produce opaque harmony.

The FB class has very few roots and tokens compared to the others. Roots with this shape are relatively uncommon, and those that do exist tend not to undergo raising.[8] Figure 2 breaks down BF roots and FB roots by their individual rates of front *vs.* back suffixes. In both cases, we see that roots are typically categorical in whether they take back or front suffixes, while a smaller number (n = 101) show variation between the two (see Zuraw 2016).

For an example of such variation, consider the root *idare* /idɑræ/ 'office, bureau'. When used with the auxiliary verb *qilmaq* 'do', it can also mean 'to rule' or 'to govern'. This root has an overall frequency in the corpora of 1,122 per million words and occurs in its raised form *idari* [idɑri] in 74% of tokens.

[8]Three of these raising FB roots end in the derivational suffix /-dɑʃ/ '-mate', as in [xizmætdɑʃ] 'officemate'. This suffix may display idiosyncractic harmonising behaviour in a similar way to the suffixes discussed in the next section, but there is insufficient data to determine this.

https://doi.org/10.1017/S0952675725100213 Published online by Cambridge University Press

The high frequency of the raised allomorph is likely due to the contexts in which this word tends to be used: since it's common to talk about an office relating to a person or entity, the root is often realised with the 3.POSS suffix [-si], as in *ürümchi sayaset* ***idarisi*** 'Ürümchi Tourism Office'. When it occurs in its raised form with a harmonising suffix attached, it displays opaque harmony in about 89% of cases (913 tokens) and surface-apparent harmony in about 11% of cases (113 tokens). The examples in (20) show tokens of /idɑræ/ from the RFA corpus in its unsuffixed form in (20a), with opaque harmony in (20b), and with surface-apparent harmony in (20c).

(20) *Representative tokens of the root* idare *from the RFA corpus*

a. *Unsuffixed*
...döletni qanun arqiliq ***idare*** *qilish...*
'…the rule of law…' (literally 'ruling the country by means of the law')

b. *Raised, with opaque harmony*
1980-yillardin boyan merkiziy axbarat ***idariside*** *ishligen.*
'He has worked at the Central Intelligence Agency since the 1980s.'

c. *Raised, with surface-apparent harmony*
Gülnar xanim saqchi ***idarisida*** *qandaq mu'amilige uchridi?*
'What kind of treatment did Gülnar receive at the police station?'

Because the final two vowels in BB and FF roots agree in backness, opaque harmony and surface harmony predict the same surface form. These roots almost categorically take the expected suffix forms.[9] Disharmonic FB and BF roots behave similarly to BB and FF roots, respectively, but both show higher rates of surface-apparent harmony. Chi-squared tests show significantly different rates of back suffix choice between BB and FB roots ($\chi^2 = 650.47$; d.f. = 1; $p < 0.0001$) and between FF and BF roots ($\chi^2 = 4597.6$; d.f. = 1; $p < 0.0001$). Thus, the quality of the harmonising vowel preceding the raised vowel, and not just the underlying quality of the raised vowel, affects suffix choice: when the backness of the preceding vowel conflicts with the backness of the raised vowel, the suffix is more likely to agree with the preceding vowel.

3.2.2. Opacity in derivational suffixes. Recall that previous work on opacity in Uyghur has suggested that certain derivational suffixes like the dimunitive /-ʧæ/ and the adjectival suffix /-ɑnæ/ behave idiosyncratically, preferring surface-apparent harmony. Manual inspection of the corpus data revealed a similar pattern for the suffix /-nɑmæ/ 'writings of' (as in /bɑburnɑmæ/ 'the writings of Babur'). Figure 3 breaks down harmony rates in the set of BF roots by suffix.[10] The suffix /-ʧæ/ rarely displays opaque harmony, aligning with observations made in the literature. /-ɑnæ/ and /-nɑmæ/ display higher rates of opaque harmony, but not as high as the general population of BF roots.

3.3. Predicting opacity

In order to identify potential factors that contribute to rates of opacity, I fitted a Bayesian mixed-effects logistic regression model to the set of raised BF and FB tokens with at least one harmonising suffix attached (for discussion of the use of logistic regression in modelling categorical corpus data, see Speelman 2014). This is a proper subset of the tokens described in the previous section, omitting the BB and FF roots. The model was fitted in R (R Core Team 2017) using the `brms` package (Bürkner 2017).

Bayesian models treat the statistical parameters of the model as random variables and use sampling techniques to estimate the posterior distribution over parameter values given the observed data and prior beliefs about credible parameter values (see, e.g., Kruschke 2014; Nicenboim & Vasishth 2016). The model here uses the default, weakly informative priors.

[9]The small number of unexpected disharmonic suffixes could be the result of typos or misidentification of the quality of the underlying vowel.

[10]The transducer typically includes these suffixes as part of the root.

https://doi.org/10.1017/S0952675725100213 Published online by Cambridge University Press

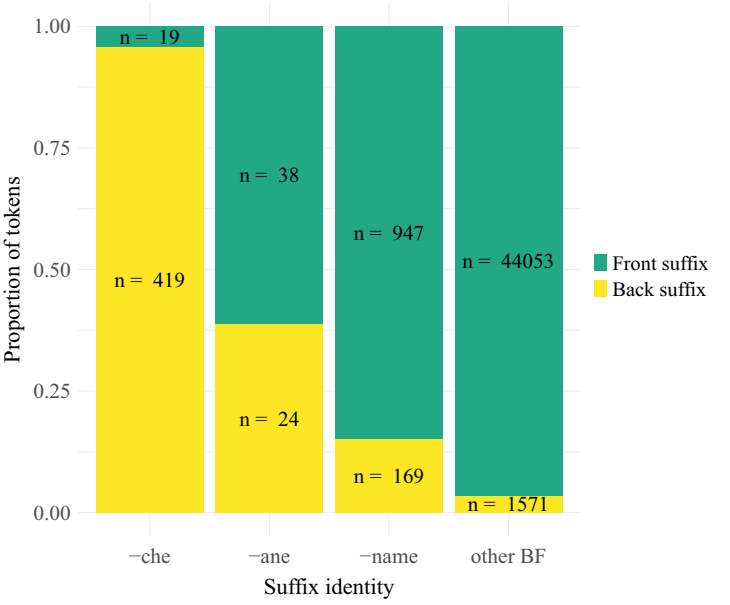


Figure 3. *Suffix choice in raised BF roots broken down by root-final derivational suffix. 'Other BF' refers to BF roots that do not end in one of the three derivational suffixes. Token counts are overlaid on each category. The tokens of /-ʧœ/ included here all have a preceding B vowel, as in /baʁʧœ/ 'park'.*

A Bayesian analysis is used for two reasons. First, the equivalent frequentist model often fails to converge with the random effects structure used here. Second, large sample sizes can lead to significant *p*-values for trivially small effects (e.g., Lin *et al.* 2013). I focus accordingly on reporting effect sizes: that is, the values of the coefficients of the fitted model. In addition to point estimates of these coefficients, Bayesian models also provide interpretable estimates of the range of credible values.[11]

The dependent variable was coded as either opaque harmony (1) or surface-apparent harmony (0). An opaque response was defined as a back suffix attached to a raised FB root, or a front suffix attached to a raised BF root.

The independent variables were selected based on previous work on Uyghur or other languages with similar vowel harmony systems:

- The *log token frequency of the root* in the corpora, normalised to count per million words. This was included as a predictor because frequency is often an important driver of phonological variability (e.g., Coetzee & Kawahara 2012; Coetzee 2016).
- The *proportion of tokens of the root that are raised*. This is defined as the number of tokens of a root containing the raised allomorph divided by the total number of tokens of that root.

  For example, the root /ɑpæt/ 'disaster' occurs 1,719 times in the corpus. Of these tokens, 544 are in forms that exhibit raising (e.g., [ɑpit-i] 'disaster-3.POSS') and 1,175 are in unraised forms (e.g., [ɑpæt-lær] 'disaster-PL'). Thus, the proportion of raised tokens for this root is $544/1,719 = 0.32$.

  This variable was included based on the observation by Hahn (1991b) that raised forms obscure the harmonic class of a root (i.e., whether it takes front or back suffixes). That is, unraised tokens provide evidence of the harmonic class of the root in the form of the final vowel, while raised tokens (particularly those with non-harmonising suffixes) do not.

[11] While confidence intervals provide a measure of dispersion for coefficient estimates in frequentist models, their interpretation is counterintuitive and frequently misunderstood (Hoekstra *et al.* 2014). The credible intervals produced by Bayesian models better align with intuitions about what such measures of dispersion should communicate.

https://doi.org/10.1017/S0952675725100213 Published online by Cambridge University Press

- The *identity of the underlying raised vowel* (F or B). This allows us to test the proposal that raised /ɑ/ is more likely to display surface-apparent harmony than raised /æ/ (Vaux 2000).
- The *distance between the rightmost two harmonising root vowels*, counted in segments. Previous work on Uyghur (Mayer 2021a) and languages with similar harmony systems (Hayes *et al.* 2009; Rebrus & Törkenczy 2017, 2021) suggests that the influence of a vocalic harmony trigger in the root decreases as greater numbers of transparent segments intervene between it and the following suffix. In a root like /ɑ**p**æt/ 'disaster', this distance is 1; in /pɑʔɑ**lij**æt/ 'activity', it is 3.
- The *distance between the root and the first harmonising suffix*. This is calculated as the number of morphological tags between the final tag of the root and the tag of the first harmonising suffix. For example, /ɑpæt-i-GA/ 'disaster-3.POSS-DAT' has a distance of 1 (the non-harmonising 3.POSS intervenes between the root and the harmonising DAT); /ɑpæt-GA/ 'disaster-DAT' has a distance of 0. Rebrus & Törkenczy (2017, 2021) found no influence of this factor on harmonising behaviour in Hungarian.
- Whether the root ends in one of the three *derivational suffixes* discussed above. This was operationalised as three dummy-coded variables that took the value 1 if the root ends in /-ʃæ/, /-ɑnæ/ and /-nɑmæ/, respectively, and 0 if not.

Random intercepts were defined for:

- A nested effect of *author* within *corpus*. This controls for different rates of opacity across sources and individual writers.[12] A nested effect is used because authors are unique within corpora.
- *Root identity*. This controls for the idiosyncratic tendencies of roots that are not captured by the dependent variables.

The results shown in Table 3 show two common summary statistics for the coefficients in Bayesian models:[13] the mean, or most credible, value of each coefficient, and the 95% credible interval (95% CI), which is the range in which the central 95% of the sampled coefficient values fall given the model and data set. Expressed in slightly different terms, the 95% CI tells us that the model estimates that there is a 95% chance that the true value of this parameter falls within this range, given the data. A 95% CI that does not include zero is interpreted as meaningful, since it indicates that the directionality of the effect is highly credible. In addition, the 95% CI provides a measurement of uncertainty about the effect size. See Appendix B for more details on the model.

These results suggest several frequency-related influences on opacity rates: more frequent roots are more likely to harmonise opaquely, but roots that frequently occur in their raised forms, where the underlying identity of the final vowel is obscured, are less likely to harmonise opaquely.

There are also phonological contributions. The model shows that the underlying identity of the raised vowel is a significant predictor of opaque harmony, with underlying back vowels being less likely to harmonise opaquely (roughly aligning with the claims in Vaux 2000). The distance between the rightmost two harmonising vowels in the root is also positively correlated with rates of opacity: as the disharmonic vowel in the root becomes further from the suffix, its influence decreases.

Finally, there is also evidence for morphological influences on opacity rates. The three derivational suffixes described above each produce greater rates of surface-apparent harmony than are seen in the general population of roots, and the specific rates vary between suffixes.

[12] Articles in the *Awazi* corpus are always published under the anonymous byline *admin.* The *Akadémiye* corpus is similar, with a small number of exceptions, resulting in a total of four unique author bylines. The RFA corpus provides greater attribution, but commonly uses partially anonymous bylines, such as *muxbirimiz Erkin* 'our reporter Erkin', resulting in a total of 154 unique author bylines. These conventions make consistent authorship attribution difficult. When coding the authorship variable used in the mixed effects logistic regression model, I use the listed author name when available. If no attribution is provided, the author is coded as 'None_<corpus_name>' to distinguish between unnamed authors in the different corpora.

[13] The coefficients here are the estimated change in the log odds of opaque harmony given a unit increase in the corresponding variable, where the log odds are defined as $\log \frac{P(\text{opaque harmony})}{1-P(\text{opaque harmony})}$. Positive values indicate an increased likelihood of opaque harmony, while negative values indicate a decreased likelihood. For example, the coefficient 0.32 for the log token count tells us that, for each unit increase in log frequency, the log odds of an opaque response relative to a surface-apparent response increase by 0.32, which corresponds to a 37% increase in the odds.

https://doi.org/10.1017/S0952675725100213 Published online by Cambridge University Press

Table 3. *Results from a mixed-effects logistic regression model whose coefficients were estimated using Bayesian inference. The 95% credible interval shows the central range in which 95% of the sample values occur. Credible intervals that do not contain zero are interpreted as a meaningful directional effect, and are marked with **

Term	Mean estimate	95% CI	
Intercept	3.46	−0.89	6.36
* Log token count	0.32	0.11	0.51
* Proportion raised	−2.25	−3.77	−0.84
* Final vowel (reference level BF)	−3.65	−6.36	−1.02
* Root vowel distance	0.80	0.41	1.23
* Suffix distance	−1.29	−1.45	−1.13
* Has -ɑnæ	−4.62	−6.77	−2.57
* Has -nɑmæ	−2.12	−3.42	−0.88
* Has -ʃæ	−9.58	−11.55	−7.73
Corpus (standard deviation)	1.89	0.16	5.86
Corpus:Author (standard deviation)	0.90	0.64	1.23
Root (standard deviation)	2.36	1.99	2.78

The number of non-harmonic suffixes directly between a root and the closest harmonic suffix was correlated with more surface-apparent harmony, which is unexpected given results from Rebrus & Törkenczy (2017, 2021). Because the majority of roots that can generate opacity are BF roots, this tendency might reflect an increased preference for the use of the default back suffix forms as distance between root and harmonising suffix increases.

4. An analysis of opacity in the corpus

Given the quantitative results described above, a model of opacity in Uyghur that is descriptively adequate (in the sense articulated by Chomsky 1965) should be able to account for the following properties:

- The majority of raised tokens harmonise opaquely, but cases of surface-apparent harmony also exist.
- The rate of opaque harmony varies across roots. This variation is correlated with a number of phonological and morphological factors, as well as frequency.

In this section, I will outline a simple parallel model that can account for these factors.

4.1. Phonological and lexical effects in backness harmony

Backness harmony is typically treated as driven by surface phonological constraints (e.g., van der Hulst 2016): whether a root takes front or back suffix allomorphs depends on which variant will minimise resulting surface disharmony by some criteria, as in the simple model presented in §2.5. I will refer to this as the *phonological component* of backness harmony.

There are cases in Uyghur, however, where surface phonological properties are not sufficient to determine harmonising behaviour (see Mayer *et al.* 2022b for a more detailed description). For example, the pair of words /sir/ ‘secret’ and /bir/ ‘one’ are nearly identical; however, /sir/ takes back suffixes (e.g., [sir-lɑr] ‘secret-PL’), while /bir/ takes front suffixes (e.g., [bir-gæ] ‘one-DAT’). Whether such roots take front or back suffixes cannot be predicted from their acoustic properties (Mayer *et al.* 2022a). Similarly,

https://doi.org/10.1017/S0952675725100213 Published online by Cambridge University Press

while the majority of roots that contain no harmonising vowels and a velar consonant /k g/ take front suffixes (e.g., [kir-gæ] 'dirt-DAT'), a smaller number take back suffixes (e.g., [gips-qɑ] 'plaster-DAT').

Hayes (2016) uses the term *zones of variation* to describe similar roots in Hungarian. Because their harmonic class is at best partially predictable from phonological properties, speakers disproportionately rely on lexical knowledge: that /sir/ takes back suffixes while /bir/ takes front suffixes must simply be memorised as a fact about each root. I will refer to this as the *lexical component* of backness harmony. A consequence is that such roots typically display higher degrees of variability in suffix choice, particularly in wug tests, where lexical knowledge is absent.

The effect of lexical information on backness harmony systems has been analysed on the basis of *Harmonic Uniformity* (Rebrus & Törkenczy 2017, 2021; Rebrus *et al.* 2023). This is a paradigm uniformity constraint which requires that the harmonic class of a root remain consistent across its extended paradigm (see, e.g., Steriade 2000 for more on paradigm uniformity). This constraint can override phonological processes that might otherwise apply. Rebrus and colleagues have provided evidence for this constraint on the basis of a variety of phenomena in the Hungarian backness harmony system. Here, we will focus on one particular consequence of Harmonic Uniformity: the harmonic class of a root should be consistent across its suffixed forms. That is, the same root should not take back forms of one harmonising suffix and front forms of another.

Speakers learn the lexical and phonological components of phonological systems simultaneously (Zuraw 2000, 2010). In the vast majority of cases in Uyghur and Hungarian backness harmony, these components favour the same suffix choices. In a smaller number of cases, such as the zones of variation described above, the phonological component is less informative and lexical knowledge plays a larger role. The opaque forms discussed in this article are cases where the phonological and lexical components actively *conflict*: in a case like /ɑpæt-i-GA/ 'disaster-3.POSS-DAT', the phonological component favours [ɑpitiʁɑ] because it displays surface-apparent harmony, while the lexical component favours the disharmonic [ɑpitigæ] because it is consistent with the paradigmatic harmonising behaviour of /ɑpæt/. In the following sections, I will demonstrate how Harmonic Uniformity can be used to predict the variable rates of opacity found in Uyghur.

4.2. Modelling gradience in opacity using Harmonic Uniformity

The models described below will use Maximum Entropy Harmonic Grammar (henceforth MaxEnt; Goldwater & Johnson 2003), a generalisation of Optimality Theory with numeric constraint weights (Pater 2009). Higher weights indicate a greater penalty for constraint violation. MaxEnt uses these weights and violation profiles to compute probability distributions over output candidates. See Appendix C for more details.

The phonological component of backness harmony will be modelled using variants of the simple VAGREE constraint introduced in §2.5. The constraints *UNRAISED and ID[HEIGHT], described in the same section, will be used to model vowel raising.

Lexical knowledge about the harmonic class of individual roots is modelled using the constraint HARMONICUNIFORMITY:

(21) HARMONICUNIFORMITY: Thebackness of a harmonising suffix must be identical to the harmonic class of the root.

I make a minor theoretical innovation here to allow this constraint to generate the variability in rates of opacity seen in Uyghur. Rebrus and collaborators divide roots into three harmonic classes: front roots, which consistently take front suffixes; back roots, which consistently take back suffixes; and vacillators, which take either. While this tripartite distinction is useful for capturing backness harmony patterns in broad strokes, it does not provide a mechanism to predict root-specific variation.

Instead, I propose that the violations of HARMONICUNIFORMITY are scaled based on certainty in the harmonic class of a root. That is, for roots where the harmonic class is certain, attaching a suffix that conflicts with its harmonic class will incur a large violation of HARMONICUNIFORMITY, while attaching a

https://doi.org/10.1017/S0952675725100213 Published online by Cambridge University Press

suffix that agrees with it will incur no penalty. For roots whose harmonic class is uncertain, violations of HarmonicUniformity will be similar between front and back suffixes, and thus phonological factors will play a greater role in deciding suffix backness. Under this conception, vacillating roots are those where certainty in the harmonic class is low.

It is natural to think of certainty in terms of a probability distribution over the harmonic classes FRONT and BACK given a root *x*. I notate this distribution as $P(\text{HC}|x)$, and occasionally use the abbreviated form $P(\text{HC})$ when the identity of the root is clear from context. Roots that categorically take front suffixes will have $P(\text{HC} = \text{FRONT}|x) \approx 1$ (and accordingly $P(\text{HC} = \text{BACK}|x) \approx 0$); roots that categorically take back suffixes will have $P(\text{HC} = \text{BACK}|x) \approx 1$ (and accordingly $P(\text{HC} = \text{FRONT}|x) \approx 0$); and maximally ambiguous roots will have $P(\text{HC} = \text{FRONT}|x) \approx P(\text{HC} = \text{BACK}|x)$. Concretely, both front and back suffix forms will violate HarmonicUniformity, but the violations of each are scaled by certainty that the root falls into the opposite harmonic class: violations of front-suffixed forms are scaled by $P(\text{HC} = \text{BACK})$ and violations of back-suffixed forms are scaled by $P(\text{HC} = \text{FRONT})$.

Let's look at an example. Suppose the root /ɑʔilæ/ 'family' has $P(\text{HC} = \text{FRONT}) = 0.99$ and the root /ɑhɑlæ/ 'resident' has $P(\text{HC} = \text{FRONT}) = 0.7$. The tableaux in (22) show the output of a simple MaxEnt model fitted only to these data points, demonstrating how scaling the violations of HarmonicUniformity by these probabilities produces variability in whether harmony is opaque or surface-apparent. For /ɑʔilæ/, in (22a), the certainty of harmonic class is so great that it overrides the violation of surface harmony. For /ɑhɑlæ/, in (22b), where there is less certainty in class membership, we see variability.

(22) *Two roots with different levels of certainty in their class membership*

a. *Tableau for the consistently opaque form /aʔilæ-lAr/ 'family-PL'*

/ɑʔilæ-lAr/	Obs. freq.	Pred. prob.	*H*	VAgree *w* = 11	HarmonicUniformity *w* = 31
ɑʔili-lær	1	≈ 1	11.3	1	$P(\text{HC} = \text{BACK}) = 0.01$
ɑʔili-lɑr	0	≈ 0	30.7	0	$P(\text{HC} = \text{FRONT}) = 0.99$

b. *Tableau for the variably opaque form /ahalæ-lAr/ 'resident-PL'*

/ɑhɑlæ-lAr/	Obs. freq.	Pred. prob.	*H*	VAgree *w* = 11	HarmonicUniformity *w* = 31
ɑhɑli-lær	0.787	0.797	20.3	1	$P(\text{HC} = \text{BACK}) = 0.3$
ɑhɑli-lɑr	0.213	0.203	21.7	0	$P(\text{HC} = \text{FRONT}) = 0.7$

Note that if $P(\text{HC} = \text{FRONT}) = P(\text{HC} = \text{BACK})$, or if the weight of HarmonicUniformity were 0, the output would be entirely determined by violations of VAgree; conversely, if the weight of VAgree were 0, the output would be determined entirely by lexical knowledge, and the predicted probability of [ɑhɑli-lær] in (22b) would be higher.

4.3. Calculating P(HC|x)

The previous section showed how varying degrees of certainty in the harmonic class produce different rates of opacity across roots. We now turn to the question of how we determine $P(\text{HC}|x)$. That is, what properties are speakers sensitive to when determining the harmonic class of a root?

I propose a simple and rather coarse model that encodes some of the factors that may determine certainty in the harmonic class of a root based on the results of the corpus study. The key piece of evidence is, of course, the *distribution* of the root: that is, do we typically see this root with front suffixes or back suffixes? This can be thought of as the driving factor behind lexical knowledge of root harmonising class.

https://doi.org/10.1017/S0952675725100213 Published online by Cambridge University Press

In addition, however, there are general properties of roots that fall into each harmonic class that can be used to infer the harmonic class of a root, even in the absence of clear distributional evidence: namely, the phonotactic properties of the root and its morphological composition, both mediated by frequency. These factors are enumerated below:

- In addition to suffixed forms, evidence for harmonic class comes from the *phonotactic* properties of the root. Some phonotactic properties of roots are highly predictive of harmonic class: if a root ends in a back vowel, you can be quite certain that it will belong to the class of back harmonisers, even if you have never encountered a suffixed form. Other properties are more weakly predictive, such as the presence of harmonising consonants in the root, as described briefly in §4.1.
- The *morphological composition* of the root is also important: certain derivational suffixes are more prone to surface-apparent harmony than others. This may relate to whether these derived forms are treated as roots in their own right, in which case opaque harmony might be expected, or as roots with disharmonic suffixes, in which case surface-apparent harmony may be preferred.
- *Prior biases*: Back harmony is the default class in Uyghur, and speakers may encode an overall preference for this class (Mayer 2021a).

The frequency-based effects observed in the corpus study connect to each of these factors. Frequent exposure to a root provides greater evidence of which suffixes it takes, as well as greater knowledge of its phonotactic properties. Roots that typically show up in raised forms (particularly with non-harmonising suffixes) do not provide as much exposure to their final vowel, and accordingly are more prone to surface-true harmony. And the relative frequency of root and derived forms has been shown to predict morphological decomposability (e.g., Hay 2001), although I do not pursue this idea further in this article. Thus, frequency plays an important role in this model, similar to other models of phonological variability (e.g., Coetzee & Kawahara 2012; Coetzee 2016).

The next section presents a modelling study that validates the claims made above.

4.4. Validating the model

To validate this proposal, I fitted six simple MaxEnt models to the set of tokens from the corpora of roots whose final two harmonising segments were BB, BF, FB and FF, and which had at least one harmonising suffix. Note that this is a broader set of tokens than used in the statistical analysis in §3: it includes tokens in contexts that do not produce raising (such as /ɑpæt-lAr/ → [ɑpætlær] 'disaster-PL', tokens that categorically fail to raise in typical raising contexts (such as /dunjɑ-si/ → [dunjɑsi] 'world-3.POSS'), and roots that are structurally ineligible for raising in any suffixed form (such as /nɑmærd/ 'disloyal', which will never raise because it has a final complex coda, or /muʔællim/ 'teacher', which will never raise because its final harmonising vowel is not in the final syllable). This data set consisted of a total of 767,761 tokens. For simplicity, tokens for each root were aggregated based on whether they had front or back suffixes. This means the models do not consider the identity or number of suffixes, merely their backness.

The rationale behind choosing these particular six models is to deconstruct the various potential influences on backness harmony and gauge which factors play the greatest role in predicting suffix choice. The models are:

1. A *surface-oriented* model of harmony which contains the VAGREE constraint defined in (13). This constraint mandates surface-apparent harmony. This model is only sensitive to the identity of the final surface harmonic vowel when determining suffix choice.
2. An *input-oriented* model that contains the constraint VAGREEUNDERLYING. This constraint mandates harmony with the underlying form of the final vowel. This model is only sensitive to the identity of the final underlying harmonic vowel when determining suffix choice.
3. A *lexical* model. This contains the HARMONICUNIFORMITY constraint defined in (21), with violations scaled according to $P(\text{HC}|x)$ for each root. This model is only sensitive to certainty in lexical harmonic class when determining suffix choice.

https://doi.org/10.1017/S0952675725100213 Published online by Cambridge University Press

4. An *input–surface* model that combines the constraints in the input-oriented and surface-oriented models. This model allows independent contributions from the identity of both the underlying and surface final harmonising vowel when determining suffix choice.
5. A *lexical–surface model* that combines the constraints in the lexical and surface-oriented models. This model allows independent contributions from the identity of the surface final harmonising vowel and certainty in lexical harmonic class when determining suffix choice. This corresponds to the analysis presented above.
6. A *lexical–input–surface* model, which combines the constraints in the lexical, surface-oriented, and input-oriented models. This model allows independent contributions from the identity of both the underlying and surface final harmonising vowels, as well as the certainty in lexical harmonic class when determining suffix choice.

All of the models above include the *UNRAISED and ID[HEIGHT] constraints defined in (14) and (15). A violation of ID[HEIGHT] was assessed for any form that exhibited raising. Violations of *UNRAISED were assessed when vowel raising did not occur in a context where it should have applied (i.e., a low vowel in a derived, word-medial open syllable). However, for roots that categorically resist raising (i.e., that never exhibited raising in the corpus), no violation was assessed.

In order to estimate $P(\text{HC} = B|x)$ for each root, I trained a mixed-effects logistic regression model to predict the likelihood of observing each root with a back suffix based on properties of the root.[14] The fixed effects in the model include a three-way interaction between final vowel identity, log root token frequency and the proportion of raised tokens of the root, as well as the pairwise interactions of these variables. The model also includes three separate binary dummy-coded predictor variables corresponding to whether the root ends with the suffix /-ʃæ/, /-ɑnæ/ or /-nɑmæ/. Finally, the model includes a random intercept for root identity. This allows root-specific deviations from population-level trends to be encoded (i.e., idiosyncratic harmonising behaviour of particular roots). I do not explicitly encode a bias towards back suffixes into the model. The separation of the fixed, population-level effects from the lexically specific random effects is crucial for the model to be able to generalise to unseen roots. More will be said about this in §5. The dependent variable in the regression was the proportion of back suffixes taken by the root.[15] The coefficients of the model fitted to the entire data set can be found in Appendix D.

Models were fitted using the R package `maxent.ot` (Mayer *et al.* 2024). In order to prevent overfitting, k-fold cross validation was used, with $k = 10$. This means that instead of fitting the grammar to the entire data set, the tokens were randomly partitioned into ten subsets. One of the subsets (10% of the data) was held out, and the model was trained on the remaining nine subsets (90% of the data). The trained model was then applied to predict the held-out subset and the log-likelihoods of the model applied to the training and test sets were recorded. This process was repeated ten times, with each subset being held out once. The same partitions were used for each model.

For the lexical models, the logistic regression model used to approximate $P(\text{HC} = B|x)$ was fitted only to the training data in each case. It is important to note, however, that the predictor variables in the model, log root token frequency and the proportion of raised tokens, were calculated based on counts from the entire corpus. This includes unsuffixed tokens and tokens without harmonising suffixes, which are not included in the current analysis. This means that the values of these variables for each root were consistent across all folds, but the coefficients of the model that dictate how these properties are related to suffix backness differed depending on the training data. The models were configured to use root-specific random effects when they made predictions for roots found in the training data, but to make predictions using population-level effects when they encountered new roots.

[14] For computational tractability, this model was not implemented in a Bayesian framework and did not include random intercepts for corpus/author.

[15] Specifically, the model was fitted to an aggregated data set consisting of root types with proportion of back suffixes as the dependent variable, weighted by the number of tokens.

https://doi.org/10.1017/S0952675725100213 Published online by Cambridge University Press

Table 4. *Mean log-likelihood (LL) for the training and test sets across each of the ten folds and the optimal value of σ. The lexical–surface model (in bold) obtains the best performance on the held-out test folds with the fewest number of constraints*

Model	Constraints	σ	Mean test LL	Test std. dev.	Mean train LL	Train std. dev.
Surface-oriented	3	10	−31,926	220	−290,574	347
Input-oriented	3	10	−16,570	127	−152,369	274
Lexical	3	10	−15,730	122	−144,397	258
Input–surface	4	20	−15,670	142	−144,272	281
Lexical–surface	**4**	**10**	**−15,097**	**142**	**−138,725**	**272**
Lexical–input–surface	5	20	−15,097	142	−138,725	272

L2 regularisation was used to prevent overfitting (Goldwater & Johnson 2003). For each model, the default value of each constraint, μ, was set to 0. σ, which determines how strongly deviations from the default value are penalized, was chosen based on a simple search over a range of values: 500, 200, 100, 50, 20, 10, 1, 0.5, 0.2, 0.1, 0.05, 0.02, 0.01 and 0.001. The results below correspond to the values of σ for each model that produced the highest log-likelihood on the held-out data.

Table 4 shows for each model the mean and standard deviation of the log-likelihoods across the training and test sets for each of the ten folds, as well as the optimal value of σ. Higher log-likelihoods indicate better model fit and lower standard deviations indicate greater model stability across folds. Because of the large data set, the choice of σ was not particularly important so long as it was large enough to allow the weights to fit the data effectively. All simulations where $\sigma \geq 1$ produced similar results.

Of the six models tested, the lexical–surface model, which is sensitive to the lexical class of the root as well as the surface properties of the word, achieves the best performance with the fewest number of constraints. Three specific comparisons are particularly insightful: (1) that this model performs better than the lexical model indicates that there is a pressure towards surface-apparent harmony; (2) that it performs better than the input–surface model indicates that suffix choice is sensitive to lexical factors, such as root frequency and the proportion of raised tokens rather than simply underlying final vowel identity; and (3) that it performs identically to the lexical–input–surface model indicates that an additional constraint mandating harmony with the backness of the underlying harmonising vowel is unnecessary given the HarmonicUniformity constraint.

Although the improvement in log-likelihood between the input–surface and lexical–surface models is perhaps modest, it is important to keep in mind that roots that have the potential to produce opaque harmony are a relatively small proportion of this data set, and the roots within that class that display variability in opacity are even fewer.

5. Discussion

Based on the results of the corpus and modelling studies, this paper suggests that the opacity observed in Uyghur backness harmony is mediated by both surface constraints on backness harmony and pressure for a root to display consistent harmonising behaviour across its extended paradigm (Rebrus & Törkenczy 2017, 2021; Rebrus *et al.* 2023). When certainty in the harmonic class of a root is high, this can override conflicts with phonological generalisations; conversely, when certainty in harmonic class is low, a greater influence of surface harmony constraints may be observed. This analysis is not only able to model root-specific rates of opacity, but also accounts for the relation of these rates to factors such as frequency and morphological composition.

This account aligns with proposals that some cases of opacity are driven by paradigm uniformity (Steriade 2000), such as Canadian raising in certain dialects of English (Hayes 2004) and several

https://doi.org/10.1017/S0952675725100213 Published online by Cambridge University Press

processes in Tiberian Hebrew (Green 2004) and Polish (Sanders 2003; Łubowicz 2003). A clearer understanding of the relation between paradigm uniformity and opacity may help to further unify at least some of the variegated types of opaque phenomena that have been observed to date (Baković 2011; Baković & Blumenfeld 2019).

An additional consequence of this treatment of opacity is that it allows the Uyghur data to be analysed using a strictly parallel model. If a paradigm uniformity account is indeed correct, the inability of parallel models to represent this opaque phenomenon as a purely phonological process may be seen as a point in their favour, rather than a failure.

In the remainder of this section, I will touch on some concerns with the current model, alternative analyses and possible directions for future research.

5.1. Generalisation

One of the central goals of phonological theory is to account for generalisation from attested data to novel forms, such as is observed in wug-test studies (Berko 1958). Although the proposal made in this paper hinges on lexical knowledge of root harmonising class, it also predicts generalisation to unattested roots based on the population-level properties of attested roots: in addition to learning the specific harmonic classes of attested items, speakers also learn the properties of roots that fall into each class. The claim made here is that, by relying on these generalisations about lexical classes, speakers can estimate the harmonic class of a new root even with only a single exposure. Indeed, the simple model of $P(\text{HC} = B|x)$ presented in the previous section predicts that a B-final root observed once in its unraised form will have a 0.9997 probability of being a back harmoniser. In forms with such obvious clues to their harmonic class, the estimated effects of frequency and the proportion of raised forms from the corpus data are subtle.

This raises the question of how the claims in this paper might be tested experimentally if tokens with surface-apparent harmony are unlikely to be encountered without a large sample. I have two thoughts on this.

First, it may be the case that the model presented here overestimates the subtlety of these effects because it is trained on a relatively restricted genre of text, namely, newspaper and academic articles. Since such products are typically reviewed by an editor, surface-apparent forms, which are non-standard, are likely to be corrected prior to publication. It may be the case that these forms are more frequent in colloquial speech or writing, and therefore perhaps more amenable to simple wug-testing.

Second, even if this is not the case, there may be more nuanced ways of testing this phenomenon. Even if Uyghur speakers are highly consistent in producing opaque harmony on novel words, they might exhibit frequency-driven differences in language processing tasks (see Ellis 2002 for an overview of such effects). For example, an unexpected suffix form on a highly frequent word might be more surprising or disruptive than one on a word they have only seen once. This effect might be tested using online behavioural measures, such as eye-tracking, reading time or electroencephalography. It may also be the case that certainty in harmonic class is mediated by recency of exposure. One might test for this by asking participants to produce a suffixed form of a root immediately after seeing its unsuffixed form *vs.* seeing it in its raised form without a harmonic suffix, or to do a similar task in a priming study where they are primed with either raised or unraised forms of a root. If certainty in harmonic class is mediated by both recency and frequency, one might expect high-frequency roots to be more resistant to disruption by recent exposure to raised forms than low-frequency roots. I leave these as interesting areas for future research, in Uyghur or in other backness harmony languages where lexical factors are relevant.

5.2. Pathological roots

When considering root-specific behaviour in the Hungarian vowel harmony system, Hayes & Londe (2006) note that an acceptable analysis must preclude the existence of pathological examples like a

https://doi.org/10.1017/S0952675725100213 Published online by Cambridge University Press

simple back root (e.g., the hypothetical /pɑb/) categorically taking front suffixes. The current model allows for this pathology, to the extent that it is possible to define $P(\text{HC} = \text{FRONT}|/\text{pɑb}/)$ in such a way that this behaviour would be predicted.

Setting aside the question of how such a pathological root could arise in Uyghur in the first place, even in a case where there is sudden massive and unequivocal distributional evidence for /pɑb/ as a front harmoniser, the model predicts that its long-term survival would be uncertain for two reasons:

1. The phonological component exerts pressure for surface harmony based on general patterns in the language.
2. A root like /pɑb/ would be highly atypical within the class of Uyghur front roots.

Both of these factors predict that /pɑb/ would be produced with back suffixes at least some of the time, particularly for speakers with low exposure to the root who rely more on knowledge of the general properties of front harmonisers than on root-specific knowledge. Over generations this variation would decrease $P(\text{HC} = \text{FRONT}|/\text{pɑb}/)$, resulting in a gradual drift towards back suffixation.

The model's predictions for the future of opacity in Uyghur are less clear. Although total surface-apparent harmony would be preferable from the perspective of the phonology, the zones of variation in the harmony system preclude a total shift away from lexical factors.

5.3. Alternative analyses using serial models

Analyses of opacity that use serial variants of classical OT grammars, such as Stratal OT, do not straightforwardly predict the kind of variability in rates of opacity seen in the corpus here. It is possible, however, to generate variation between opaque and transparent outcomes by introducing weighted constraints or variable constraint ranking into these models.[16] For example, Stratal OT can be made probabilistic by specifying a MaxEnt grammar for each stratum rather than a classical OT grammar (Nazarov & Pater 2017). In terms of the Stratal OT analysis presented in (18), if the weighting of ID[BACK] is allowed to vary probabilistically in the phrase stratum as a function of individual lexical items, this model can produce opaque outcomes when ID[BACK] is weighted higher than VAGREE and transparent outcomes when it is ranked lower. Similarly, a Serial Markedness approach (Jarosz 2014) could encode this variation by making the constraints that force extrinsic ordering of processes variably ranked (Jarosz 2016; Prickett & Jarosz 2021) in a way that is sensitive to lexical identity.

It is beyond the scope of this paper to consider these alternatives in any depth, but it will be important to compare their predictions against the account provided here. I will make two brief comments, however, on why the parallel account presented here may provide greater explanatory value.

First, in addition to the need for strata or serial constraint evaluation to model opacity and the variability therein, these serial analyses will still require some mechanism to encode the lexical aspects of harmony in Uyghur, namely, the need to memorise the harmonising class of roots that fall in zones of variation. The parallel analysis presented here is somewhat more parsimonious in that it attributes both phenomena to the same lexical factors.

Second, it is not clear to me how the relation between root frequency and constraint weighting can be motivated in these models. For example, the probabilistic Stratal OT analysis requires that the weight of ID[BACK] be scaled up for high-frequency forms and down for low-frequency forms. This runs counter to most accounts of frequency effects in phonology, where high-frequency forms are typically *less* faithful than their low-frequency counterparts (e.g., Coetzee & Kawahara 2012; Coetzee 2016). Some account of this discrepancy will be important. By contrast, the relation between frequency and opacity is clear in the model presented in this paper: increased exposure to a root provides greater certainty about its harmonic class.

[16] Thanks very much to an anonymous reviewer who pointed out these possible analyses.

https://doi.org/10.1017/S0952675725100213 Published online by Cambridge University Press

5.4. Limitations of the corpus

Because this is a corpus study, it is necessarily exploratory and does not constitute hypothesis confirmation (see Roettger (2019) for more on this distinction). Exploratory analyses are most valuable as a tool for hypothesis generation. It will be important to test the predictions made by the model proposed here in controlled, experimental contexts, as described in §5.1.

As well, because this study uses corpora where authorship cannot be uniquely determined, it does not tell us to what extent the variability we observe happens at the level of the individual or in aggregate across the population as the result of dialect variation. This is a particular concern because the corpora do not come from the same region: the *Awazi* corpus is written for Uyghurs living in Kazakhstan, while the other two corpora are written for Uyghurs living in China and the diaspora. The results of this study do, however, tell us where we might expect variation, either within or across speakers, or both. The sources of this variation can be carefully determined using more granular corpus or experimental studies.

5.5. A closing remark on corpus methodologies

In addition to its empirical and theoretical contributions, this paper demonstrates the value of taking a more holistic and comprehensive empirical approach to linguistic data collection and analysis. The Internet has allowed for the proliferation of large amounts of textual data, even for languages that, like Uyghur, have relatively small populations of users. Computational tools, such as Web scrapers and morphological transducers, can allow us to marshal the complexity inherent in such large data sets, and provide access to new types of empirical data that allow us to supplement other data sources and measure phonological patterns writ large. Such tools will become increasingly important as phonological theory continues to develop.

A. The Uyghur morphological transducer

This appendix describes the morphological transducer used to analyse the corpora in this paper.

The transducer is a modified version of the *apertium-uig transducer* (Littell *et al.* 2018; Washington *et al.* 2019). This is implemented using finite-state transducers (FSTs): specifically, within the Helsinki Finite State Technology framework (HFST; Linden *et al.* 2011). An FST is a finite-state automaton (FSA) that contains two tapes – in this case, one corresponding to underlying analyses and one to surface forms. Each transition or arc in the transducer has a symbol corresponding to each tape. Either tape may be designated as the input. The transducer reads the input and takes the appropriate transitions between states. The symbols on the transitions corresponding to the output tape are written to an output buffer. If the transducer reaches a valid output state after consuming the entire input, then the contents of the output buffer are returned.

Any SPE-style system that uses sequences of rewrite rules to map from underlying analyses to surface forms can be implemented as an FST (Johnson 1972; Kaplan & Kay 1994). In practice, this poses several problems, the most serious of which is that although the mapping from an underlying analysis to a surface form is deterministic given a set of rules, the inverse is not true in general. In fact, it is possible for a given surface form to correspond to a large, or even infinite, number of underlying analyses under certain rule systems. This quickly becomes intractable for any practical implementation of a morphological transducer. The *two-level morphology* framework (Koskenniemi 1983, 1984, 1986; Beesley & Karttunen 2003), which is implemented in HFST, was designed to mitigate these issues.

Two-level morphology divides the mapping between underlying analyses and surface forms into two stages. The first stage maps between a morphological analysis and an abstract morphophonological form, which allows a minimal representation of roots and suffixes. For example, the analysis *qiz*<dat> will map to *qiz*>{G}{A} at this level, where > represents a morpheme boundary and {G} and {A} are essentially archiphonemes. It is this stage that solves the problem of overgeneration of underlying analyses: every valid underlying root must be explicitly encoded in the transducer.

https://doi.org/10.1017/S0952675725100213 Published online by Cambridge University Press

The output of the first level then serves as input to the next level, which maps abstract morphophonological forms to surface forms. In this case, the phonological rules specified in the transducer will map {G} to *gh* and {A} to *a*, producing the surface form *qizgha* 'to a girl'. (I use Latin orthography throughout this appendix rather than IPA, since it more closely reflects the input to the transducer.)

In HFST, the first stage is implemented using the LEXC formalism, while the second is implemented using the TWOLC formalism. The rules specified at these levels are compiled into FSTs, which are then compose-intersected to form a single transducer. This transducer will only accept or propose underlying roots that are specified in the lexicon. Unfortunately, this introduces a degree of brittleness, since the transducer will not recognise any forms that are not present in the lexicon, and has no means by which to 'guess' the underlying form from the surface form unless augmented with additional tools.

The transducer can also map in the opposite direction: between surface forms and underlying analyses that consist of roots plus morphological tags. To repeat the example used in §3.1, if the input is the surface form *qizingizgha* 'to your daughter', the output analysis will be *qiz*<n><px2sg><frm><dat>. This indicates that the root is *qiz*, a noun <n>, and that it is suffixed with the second person singular possessive marker in its formal form <px2sg><frm>, followed by the dative suffix <dat>.

A.1. Modifying the transducer

I modified this transducer to detect the harmonic quality of suffixes when mapping from surface to underlying forms.[17] While a form like *qizingizgha* maps to *qiz*<n><px2sg><frm><dat> under the original transducer, it maps to *qiz*<n><px2sg><frm><dat-b> under this modified system, indicating that the dative suffix surfaces in one of its back harmonising forms (*-qa* or *-gha*) rather than one of its front harmonising forms (*-ke* or *-ge*).

The modified transducer splits each tag corresponding to a harmonising suffix into three different forms corresponding to front variants (e.g., <dat-f>), back variants (e.g., <dat-b>) and ambiguous variants (e.g., <dat>). These tags are mapped to more restricted, though still abstract, morphological forms in the first stages. For example, <dat-f> will map to {Gf}{Af}, while <dat-b> will map to {Gb}{Ab}.

The second stage has been modified to map the newly introduced archiphonemes at the first stage to a restricted set of surface forms with corresponding backness. For example, it maps {Gf} and {Af} only to front allophones, and {Gb} and {Ab} only to back allophones. In addition, the restrictions that the phonological component of the transducer imposes on harmony have been lifted. The original transducer, for example, would reject a form like **at-ler*, 'horses', for being disharmonic. The modified transducer will simply interpret this as an instance of the front form of the plural suffix.

In a few cases, multiple parses that correspond to identical surface parses were removed from the transducer for simplicity. Take the word *doktur* 'doctor' as an example. This noun may be parsed as a nominal (appropriate in cases like *doktur chong* 'the doctor is old'), as a copular form (appropriate in cases like *Adil doktur* 'Adil is a doctor'), and so on. Because such distinctions are not relevant for the current project, all but the nominal parse was removed.

Finally, the vowel raising processes described in §2.3 can obscure the harmonising quality of suffixes; for example, the surface realization of /dost-lAr-m/ 'friend-PL-ISG.POSS' is [dostlirim], which does not allow the backness of the plural morpheme to be determined. In such cases, the modified transducer does not attempt to guess the backness of the suffix (i.e., to report either <pl-f> or <pl-b>), but will instead remain agnostic, simply reporting <pl>. This paper only uses tokens where the harmonic value of at least one suffix is unambiguous.

A.2. Interpreting and sanitising the transducer output

Applying the transducer to the corpus produces one or more possible parses for each word that the transducer is able to recognise. The transducer was able to successfully analyse about 2.6 million of

[17] https://github.com/connormayer/apertium-uig/tree/vowel_harmony.

https://doi.org/10.1017/S0952675725100213 Published online by Cambridge University Press

the 4.2 million words in the *Uyghur Awazi* corpus (61%), about 7 million of the 8 million words in the RFA corpus (87%), and about 1.8 million of the 2.4 million words in the Uyghur Academy Corpus (75%). The typical reason the transducer would fail to parse a word is because the root is not included in the transducer's list of valid roots; this is a key limitation of rule-based parsers. Tokens for which the transducer produced no parse were discarded.

Additional filtering was done to sanitise the data produced by the transducer. One challenge that arises is how to deal with multiple analyses in cases where they provide conflicting information about the root. The surface form *orgini,* for example, could correspond to underlying /orgɑn-i/ 'organ-3.POSS' or to /or-GAn-i/ 'harvest-PFV-3.POSS' (though the latter is unlikely because it corresponds to the front form of /-GAn/ attached to an unambiguously back root). I take a maximally conservative approach and discard all tokens for which such ambiguous parses are possible.

In addition, a number of suffixes are harmony blockers. Such suffixes, like the progressive suffix /-wɑt/ in (23), block harmony by failing to harmonise, and impose their own harmonic values on the following suffixes. Such suffixes are often historically derived from multi-word constructions (Mayer *et al.* 2022b).

(23) *Harmony-blocking* /-wɑt/

/søzlæ-wɑt-GAn/	→	[søzlæwɑtqɑn]	'speak-PROG-PFV'
cf. /søzlæ-GAn/	→	[søzligæn]	'speak-PFV'

I discard tokens containing such suffixes.

Finally, a number of spurious parses were omitted based on manual inspection of the results. Particular attention was paid to tokens reflecting surface-apparent harmony to ensure these were not overcounted.

B. Additional details on statistical models

The Bayesian models were fit using the default parameters of the `brm` function. Figure B1 shows samples from the posterior for each coefficient.

The R package DHARMa (Hartig 2022) was used to apply standard residual diagnostics to the model (Figure B2). The diagnostics show that the model satisfies tests for outliers, uniformity and zero

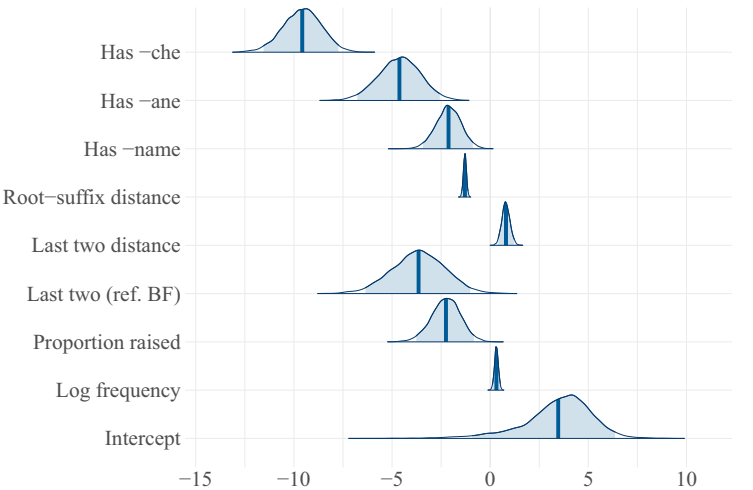


Figure B1. *Plot of samples from the posterior for each model parameter. The dark areas are the mean values, and the shaded areas are the 95% credible intervals.*

https://doi.org/10.1017/S0952675725100213 Published online by Cambridge University Press

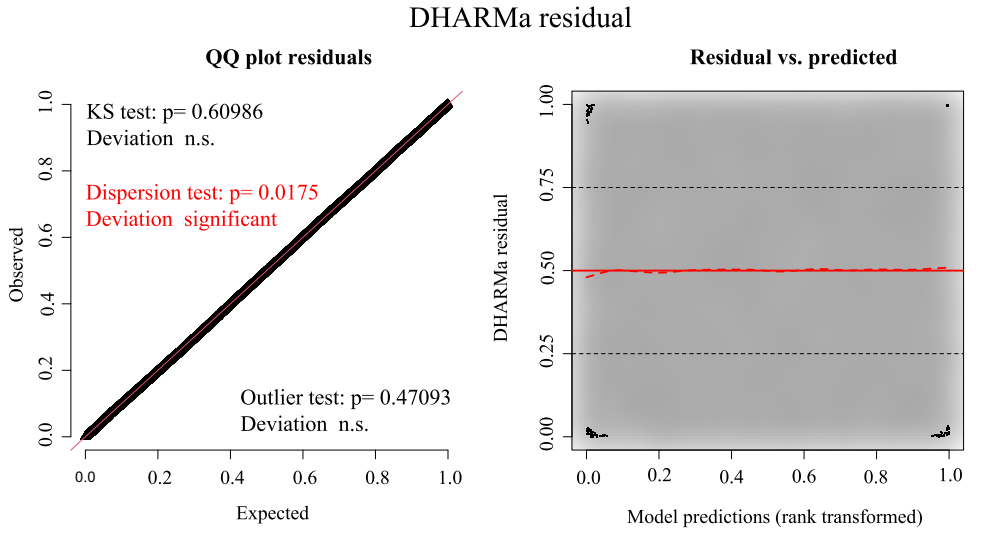


Figure B2. *Output of DHARMa model checks run on the Bayesian logistic regression model.*

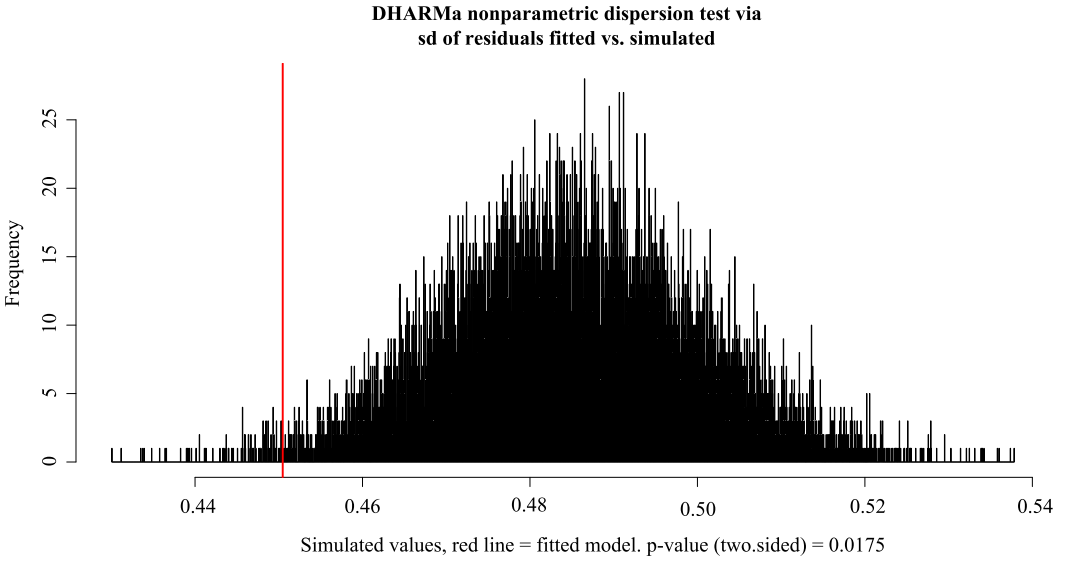


Figure B3. *Histogram comparing simulated model residual standard deviations against the residual standard deviation from the fitted model.*

inflation (not pictured), but shows significant underdispersion. This means that the residual variance in the data is smaller than expected under the fitted model (see Figure B3). The significance of this test tells us that, because of the large number of data points, we can be fairly certain that the model is underdispersed. However, this is not a concern for two reasons. First, the dispersion parameter is about 0.93, indicating that there is 7% lower variance in the observed data than predicted by the model. This is quite a small deviation. Second, underdispersion produces more conservative parameter estimates from the model (Hartig 2022), so we may still be confident in the interpretation of the model presented in the article.

https://doi.org/10.1017/S0952675725100213 Published online by Cambridge University Press

C. Maximum Entropy Optimality Theory

In a Maximum Entropy Optimality Theory grammar, each constraint is associated with a real-valued weight that represents its strength. In a grammar with N constraints, the weight of the ith constraint can be notated w_i. The function $C_i(x, y)$ returns the number of times an output candidate y for the input x violates the ith constraint. The harmony $H(x, y)$ of an output candidate y given the input x is:

$$H(x,y) = \sum_{i}^{N} w_i C_i(x,y),$$

where higher values of $H(x, y)$ are associated with more severe constraint violations. The probability of an output candidate y given input x is

$$P(y|x) = \frac{\exp(-H(x,y))}{\sum_{z\in\Omega}\exp(-H(z,y))},$$

where Ω is the set of all possible output candidates given the input x.

The likelihood of a data set under a MaxEnt model can be calculated by multiplying together the probabilities assigned to each token by the model (or summing them in log space). It is also straightforward to learn constraint weights that optimise fit to a dataset (see Hayes & Wilson 2008 for more details on this learning procedure).

D. *P*(HC = *B*|*x*) model coefficients

The coefficients of the model used to estimate $P(\text{HC} = B|x)$ in §4.4 are shown in Table D1. It is important to note that these are the coefficients learned when the model is fitted to the entire data set. These will differ slightly from the coefficients learned during the k-fold cross-validation process, though not substantially. p-values are not reported because this model is being used only for predictive purposes.

The coefficients correspond to the expected effects: the final vowel of the root has a substantial predictive effect on $P(\text{HC}|x)$, and this effect is strengthened the more frequent the root and weakened the more often it occurs in raised forms. As well, the three exceptional suffixes all increase $P(\text{HC} = F|x)$, which corresponds to surface-apparent harmony in those suffixed forms.

Table D1. *Coefficients from the mixed-effects logistic regression model for approximating $P(\text{HC} = B|x)$ when it is fitted to the entire data set*

Term	Coefficient
Intercept	8.27
Final vowel (F)	−17.07
Log token count	0.66
Proportion raised	−0.08
Has *-ane*	10.02
Has *-name*	4.96
Has *-che*	8.01
Final vowel (F) : Log token count	−0.95
Final vowel (F) : Proportion raised	0.74
Log token count : Proportion raised	−0.13
Final vowel (F) : Log token count : Proportion raised	0.20
Root (standard deviation)	5.40

https://doi.org/10.1017/S0952675725100213 Published online by Cambridge University Press

Abbreviations. 1SG.POSS = first-person singular possessive, 3.POSS = third-person possessive, 3SG.PST = third-person singular past, ACC = accusative, DAT = dative, DIM = diminutive, GER = gerund, LOC = locative, PFV = perfective, PL = plural, PROG = progressive.

Data availability statement. The corpora and code used in the statistical analysis and MaxEnt model can be found at https://github.com/connormayer/uyghur_corpora. The modified transducer used to parse the corpus can be found at https://github.com/connormayer/apertium-uig/tree/vowel_harmony.

Acknowledgements. I thank Bruce Hayes, Kie Zuraw, Adam McCollum, Tim Hunter, Eric Baković, Travis Major, Canaan Breiss, and the attendees of AMP 2020, mfm 2022, UCLA Phonology Seminar and the Berkeley Phorum for their feedback at various stages of this project; Jonathan Washington for his help with the Apertium transducer; Tyler Carson, Daniela Zokaeim and Rutvik Gandhasri for their work on the Web scrapers; and Gülnar Eziz, Mahire Yakup, Gulnisa Nazarova and other members of the Uyghur community for sharing their language and culture with me.

Funding statement. This work was supported by the Social Sciences and Humanities Research Council of Canada.

Competing interests. The author declares that there are no competing interests regarding the publication of this article.

Ethical standards. The research meets all ethical guidelines, including adherence to the legal requirements of the study country.

References

Abdulla, Arslan, Yari Ebeydulla & Abduréhim Raxman (2010). *Hazirqi zaman uyghur tili [Modern Uyghur]*. Ürümchi: Xinjiang Xelq Neshriyati [Xinjiang People's Publishing House].

Al-Mozainy, Hamza Q. (1981). *Vowel alternations in a Bedouin Hijazi Arabic dialects*. PhD dissertation, University of Texas at Austin.

Archangeli, Diana (1988). Aspects of underspecification theory. *Phonology* **5**, 183–207.

Baković, Eric (2007). A revised typology of opaque generalisations. *Phonology* **24**, 1–43.

Baković, Eric (2011). Opacity and ordering. In John A. Goldsmith, Jason Riggle & Alan C. L. Yu (eds.) *The handbook of phonological theory*, 2nd edition. Oxford: Blackwell, 40–67.

Baković, Eric & Lev Blumenfeld (2019). Rule interaction conversion operations. *Loquens* **6**, article no. e062 (14 pp.).

Beesley, Kenneth R. & Lauri Karttunen (2003). Two-level rule compiler. Ms., Xerox Research Centre Europe and Stanford University. https://web.archive.org/web/20150608132412/https://web.stanford.edu/~laurik/.book2software/twolc.pdf.

Berko, Jean (1958). The child's learning of English morphology. *Word* **14**, 150–177.

Bermúdez-Otero, Ricardo (2003). The acquisition of phonological opacity. In Jennifer Spenader, Anders Eriksson & Östen Dahl (eds.) *Proceedings on the Stockholm Workshop on Variation within Optimality Theory*. Stockholm: Department of Linguistics, Stockholm University, 25–36.

Bermúdez-Otero, Ricardo (2018). Stratal phonology. In S. J. Hannahs & Anna R.K. Bosch (eds.) *The Routledge handbook of phonological theory*. Abingdon: Routledge, 100–134.

Bodrogligeti, András J. E. (2001). *A grammar of Chagatay*. Munich: Lincom Europa.

Bowers, Dustin (2019). The Nishnaabemwin restructuring controversy: new empirical evidence. *Phonology* **36**, 187–224.

Bürkner, Paul-Christian (2017). brms: an R package for Bayesian multilevel models using Stan. *Journal of Statistical Software* **80**, 1–28.

Chomsky, Noam (1965). *Aspects of the theory of syntax*. Cambridge, MA: MIT Press.

Chomsky, Noam & Morris Halle (1968). *The sound pattern of English*. New York: Harper & Row.

Clauson, Gerard (1972). *An etymological dictionary of pre-thirteenth-century Turkish*. Oxford: Clarendon Press.

Coetzee, Andries W. (2016). A comprehensive model of phonological variation: grammatical and non-grammatical. *Phonology* **33**, 211–246.

Coetzee, Andries W. & Shigeto Kawahara (2012). Frequency biases in phonological variation. *NLLT* **31**, 47–89.

Crosswhite, Katherine (2001). *Vowel reduction in Optimality Theory*. New York: Routledge.

Donegan, Patricia J. & David Stampe (1979). The study of natural phonology. In Daniel A. Dinnsen (ed.) *Current approaches to phonological theory*. Bloomington, IN: Indiana University Press, 126–173.

Ellis, Nick C. (2002). Frequency effects in language processing: a review with implications for theories of implicit and explicit language acquisition. *Studies in Second Language Acquisition* **24**, 143–188.

Engesæth, Tarjei, Mahire Yakup & Arienne Dwyer (2010). *Teklimakandin salam: hazirqi zaman Uyghur tili qollanmisi / Greetings from the Teklimakan: a handbook of Modern Uyghur*. Lawrence, KS: University of Kansas Scholarworks.

Farris-Trimble, Ashley & Anne-Michelle Tessier (2019). The effect of allophonic processes on word recognition: eye-tracking evidence from Canadian raising. *Lg* **95**, e136–e160.

Goldwater, Sharon & Mark Johnson (2003). Learning OT constraint rankings using a maximum entropy model. In Jennifer Spenader, Anders Eriksson & Östen Dahl (eds.) *Proceedings of the Stockholm Workshop on Variation within Optimality Theory*. Stockholm: Stockholm University, Department of Linguistics, 111–120.

Green, Anthony D. (2004). Opacity in Tiberian Hebrew: morphology, not phonology. *ZAS Papers in Linguistics* **37**, 37–70.

https://doi.org/10.1017/S0952675725100213 Published online by Cambridge University Press

Hahn, Reinhard F. (1991a). Diachronic aspects of regular disharmony in Modern Uyghur. In William G. Boltz & Michael C. Shapiro (eds.) *Studies in the historical phonology of Asian languages*. Amsterdam: Benjamins, 68–101.
Hahn, Reinhard F. (1991b). *Spoken Uyghur*. Seattle: University of Washington Press.
Hall, Daniel Currie & Avery Ozburn (2019). When is derived [i] transparent? A subtractive approach to Uyghur vowel harmony. *NELS* **49**, 89–102.
Halle, Morris, Bert Vaux & Andrew Wolfe (2000). On feature spreading and the representation of place of articulation. *LI* **31**, 387–444.
Hartig, Florian (2022). DHARMa: residual diagnostics for hierarchical (multi-level / mixed) regression models. R package, version 0.4.6. https://cran.r-project.org/web/packages/DHARMa/.
Hay, Jennifer (2001). Lexical frequency in morphology: is everything relative? *Linguistics* **39**, 1041–1070.
Hayes, Bruce (2004). Phonological acquisition in Optimality Theory: the early stages. In René Kager, Joe Pater & Wim Zonneveld (eds.) *Constraints in phonological acquisition*. Cambridge: Cambridge University Press, 158–203.
Hayes, Bruce (2016). Comparative phonotactics. *CLS* **50**, 265–285.
Hayes, Bruce & Zsuzsa Londe (2006). Stochastic phonological knowledge: the case of Hungarian vowel harmony. *Phonology* **23**, 59–104.
Hayes, Bruce & Colin Wilson (2008). A maximum entropy model of phonotactics and phonotactic learning. *LI* **39**, 379–440.
Hayes, Bruce, Kie Zuraw, Peter Siptar & Zsuzsa Londe (2009). Natural and unnatural constraints in Hungarian vowel harmony. *Lg* **85**, 822–863.
Hoekstra, Rink, Richard D. Morey, Jeffrey N. Rouder & Eric-Jan Wagenmakers (2014). Robust misinterpretation of confidence intervals. *Psychonomic Bulletin & Review* **21**, 1157–1164.
Hooper, Joan B. (1976). *An introduction to natural generative phonology*. New York: Academic Press.
van der Hulst, Harry (2016). Vowel harmony. In Mark Aronoff (ed.) *Oxford research encyclopedia of linguistics*. Oxford: Oxford University Press. not paginated. https://doi.org/10.1093/acrefore/9780199384655.013.38.
Jarosz, Gaja (2014). Serial markedness reduction. In John Kingston, Claire Moore-Cantwell, Joe Pater & Robert Staubs (eds.) *Proceedings of the 2013 Annual Meeting on Phonology*. Washington, DC: Linguistic Society of America, 12 pp.
Jarosz, Gaja (2016). Learning opaque and transparent interactions in Harmonic Serialism. In Gunnar Ólafur Hansson, Ashley Farris-Trimble, Kevin McMullin & Douglas Pulleyblank (eds.) *Proceedings of the 2015 Annual Meeting on Phonology*. Washington, DC: Linguistic Society of America, 12 pp.
Johnson, C. Douglas (1972). *Formal aspects of phonological description*. The Hague: Mouton.
Kaplan, Ronald & Martin Kay (1994). Regular models of phonological rule systems. *Computational Linguistics* **20**, 331–378.
Khanjian, Hrayr (2009). Stress dependent vowel reduction. *BLS* **35**, 178–189.
Kiparsky, Paul (1971). Historical linguistics. In William Orr Dingwall (ed.) *A survey of linguistic science*. College Park, MD: Linguistics Program, University of Maryland, 576–642.
Kiparsky, Paul (1973). Abstractness, opacity, and global rules. In Osamu Fujimura (ed.) *Three dimensions of linguistic theory*. Tokyo: TEC, 57–86.
Kiparsky, Paul (1982). Lexical morphology and phonology. In Linguistic Society of Korea (ed.) *Linguistics in the morning calm: selected papers from SICOL-1981*. Seoul: Hanshin, 3–91.
Kiparsky, Paul (2000). Opacity and cyclicity. *The Linguistic Review* **17**, 351–367.
Kirchner, Robert (1996). Synchronic chain shifts in Optimality Theory. *LI* **27**, 341–350.
Koskenniemi, Kimmo (1983). *Two-level morphology: a general computational model for word-form recognition and production*. Publication 11, Department of General Linguistics, University of Helsinki.
Koskenniemi, Kimmo (1984). A general computational model for word-form recognition and production. In Anna Sågvall Hein (ed.) *Proceedings of the 4th Nordic Conference of Computational Linguistics (NODALIDA 1983)*. Uppsala: Centrum för datorlingvistik, Uppsala University, 145–154.
Koskenniemi, Kimmo (1986). Compilation of automata from morphological two-level rules. In Fred Karsson (ed.) *Papers from the 5th Nordic Conference of Computational Linguistics* Helsinki: Department of General Linguistics, University of Helsinki, 143–149.
Kruschke, John (2014). *Doing Bayesian data analysis: a tutorial with R, JAGS, and Stan*. New York: Academic Press.
de Lacy, Paul (2002). *The formal expression of markedness*. PhD dissertation, University of Massachusetts Amherst.
Lin, Mingfeng, Henry C. Lucas Jr. & Galit Shmueli (2013). Too big to fail: large samples and the *p*-value problem. *Information Systems Research* **24**, 906–917.
Lindblad, Vern M. (1990). *Neutralization in Uyghur*. Master's thesis, University of Washington.
Linden, Krister, Miikka Silfverberg, Erik Axelson, Sam Hardwick & Tommi Pirinen (2011). HFST: framework for compiling and applying morphologies. In Cerstin Mahlow & Michael Pietrowski (eds.) *Systems and frameworks for computational morphology*, number 380 in Communications in Computer and Information Science. Heidelberg: Springer, 67–85.
Littell, Patrick, Tian Tian, Ruochen Xu, Zaid Sheikh, David Mortensen, Lori Levin, Francis Tyers, Hiroaki Hayashi, Graham Horwood, Steve Sloto, Emily Tagtow, Alan Black, Yiming Yang, Teruko Mitamura & Eduard Hovy (2018). The ARIEL-CMU situation frame detection pipeline for LoReHLT16: a model translation approach. *Machine Translation* **32**, 105–126.
Łubowicz, Anna (2002). Derived environment effects in Optimality Theory. *Lingua* **112**, 243–280.
Łubowicz, Anna (2003). *Contrast preservation in phonological mappings*. PhD dissertation, University of Massachusetts Amherst.

https://doi.org/10.1017/S0952675725100213 Published online by Cambridge University Press

Mascaró, Joan (2009). Comparative markedness and derived environments. *Theoretical Linguistics* **29**, 113–122.
Mayer, Connor (2021a). Capturing gradience in long-distance phonology using probabilistic tier-based strictly local grammars. In Allyson Ettinger, Ellie Pavlick & Brandon Prickett (eds.) *Proceedings of the Society for Computation in Linguistics 2021*. Kerrville, TX: Association for Computational Linguistics, 39–50.
Mayer, Connor (2021b). *Issues in Uyghur backness harmony: corpus, experimental, and computational studies*. PhD dissertation, University of California, Los Angeles.
Mayer, Connor, Travis Major & Mahire Yakup (2022a). Are neutral roots in Uyghur really neutral? evaluating a covert phonemic contrast. In Peter Jurgec, Liisa Duncan, Emily Elfner, Yoonjung Kang, Alexei Kochetov, Brittney K. O'Neill, Avery Ozburn, Keren Rice, Nathan Sanders, Jessamyn Schertz, Nate Shaftoe & Lisa Sullivan (eds.) *Proceedings of the 2021 Annual Meeting on Phonology*. Washington, DC: Linguistic Society of America, 12 pp.
Mayer, Connor, Adam McCollum & Gülnar Eziz (2022b). Issues in Uyghur phonology. *Language and Linguistics Compass* **16**, article no. e12478, 25 pp.
Mayer, Connor, Adeline Tan & Kie Zuraw (2024). Introducing maxent.ot: an R package for maximum entropy constraint grammars. *Phonological Data and Analysis* **6**, article no. 4, 44 pp.
McCarthy, John J. (1999). Sympathy and phonological opacity. *Phonology* **16**, 331–339.
McCarthy, John J. (2007). *Hidden generalizations: phonological opacity in Optimality Theory*. London: Equinox.
McCollum, Adam G. (2020). Sonority-driven stress and vowel reduction in Uyghur. In Hyunah Baek, Chikako Takahashi & Alex Hong-Lun Yeung (eds.) *Proceedings of the 2019 Annual Meeting on Phonology*. Washington, DC: Linguistic Society of America, 12 pp.
McCollum, Adam G. (2021). Transparency, locality, and contrast in Uyghur backness harmony. *Laboratory Phonology* **12**, article no. 10, 35 pp.
Mielke, Jeffrey, Michael Armstrong & Elizabeth Hume (2003). Looking through opacity. *Theoretical Linguistics* **29**, 123–139.
Nazarov, Aleksei (2019). Formalizing the connection between opaque and exceptionful generalizations. *Toronto Working Papers in Linguistics* **41**, 17 pp.
Nazarov, Aleksei (2020). Bedouin Arabic multiple opacity with indexed constraints in parallel OT. In Hyunah Baek, Chikako Takahashi & Alex Hong-Lun Yeung (eds.) *Proceedings of the 2019 Annual Meeting on Phonology*. Washington, DC: Linguistic Society of America, 12 pp.
Nazarov, Aleksei (2021). Learnability of indexed constraint analyses of phonological opacity. In Allyson Ettinger, Ellie Pavlick & Brandon Prickett (eds.) *Proceedings of the society for computation in linguistics 2021*. Kerrville, TX: Association for Computational Linguistics, 158–166.
Nazarov, Aleksei & Joe Pater (2017). Learning opacity in stratal maximum entropy grammar. *Phonology* **34**, 299–324.
Nazarova, Gulnisa & Kurban Niyaz (2013). *Uyghur: an elementary textbook*. Washington, DC: Georgetown University Press.
Nicenboim, Bruno & Shravan Vasishth (2016). Statistical methods for linguistic research: foundational ideas—Part II. *Language and Linguistics Compass* **10**, 591–613.
Pater, Joe (2009). Weighted constraints in generative linguistics. *Cognitive Science* **33**, 999–1035.
Pater, Joe (2010). Morpheme-specific phonology: constraint indexation and inconsistency resolution. In Steve Parker (ed.) *Phonological argumentation: essays on evidence and motivation*. London: Equinox, 123–154.
Pater, Joe (2014). Canadian raising with language-specific weighted constraints. *Lg* **90**, 230–240.
Prickett, Brandon & Gaja Jarosz (2021). Modeling the acquisition of phonological interactions: biases and generalization. In Ryan Bennett, Richard Bibbs, Mykel Loren Brinkerhoff, Max J. Kaplan, Stephanie Rich, Amanda Rysling, Nicholas Van Handel & Maya Wax Cavallaro (eds.) *Proceedings of the 2020 Annual Meeting on Phonology*. Washington, DC: Linguistic Society of America, 12 pp.
Prince, Alan & Paul Smolensky ([1993] 2004). *Optimality Theory: constraint interaction in generative grammar*. Oxford: Blackwell. Originally published in 1993 as technical report no. 2 of the Rutgers Center for Cognitive Science.
R Core Team (2017). R: a language and environment for statistical computing. https://www.R-project.org/.
Rebrus, Péter, Péter Szigetvári & Miklós Törkenczy (2023). How morphological is Hungarian vowel harmony? In Noah Elkins, Bruce Hayes, Jinyoung Jo & Jian-Leat Siah (eds.) *Proceedings of the 2022 Annual Meeting on Phonology*. Washington, DC: Linguistic Society of America, 9 pp.
Rebrus, Péter & Miklós Törkenczy (2017). Co-patterns, subpatterns and conflicting generalizations in Hungarian vowel harmony. In Harry van der Hulst & Anikó Lipták (eds.) *Approaches to Hungarian volume 15: papers from the 2015 Leiden conference*, Amsterdam: Benjamins, 135–156.
Rebrus, Péter & Miklós Törkenczy (2021). Harmonic Uniformity and Hungarian front/back harmony. *Acta Linguistica Academica* **68**, 175–206.
Roettger, Timo B. (2019). Researcher degrees of freedom in phonetic research. *Laboratory Phonology* **10**, article no. 1, 27 pp.
Sanders, Robert Nathaniel (2003). *Opacity and sound change in the Polish lexicon*. PhD dissertation, University of California, Santa Cruz.
Speelman, Dirk (2014). Logistic regression: a confirmatory technique for comparisons in corpus linguistics. In Dylan Glynn & Justyna A. Robinson (eds.) *Corpus methods for semantics: quantitative studies in polysemy and synonymy*. Amsterdam: Benjamins, 487–533.

https://doi.org/10.1017/S0952675725100213 Published online by Cambridge University Press

Steriade, Donca (2000). Morphemic invariance and the phonology–phonetics boundary. In Michael B. Broe & Janet B. Pierrehumbert (eds.) *Papers in laboratory phonology V: acquisition and the lexicon*. Cambridge: Cambridge University Press, 313–334.

Sumner, Megan (2003). *Testing the abstractness of phonological representations in Modern Hebrew weak verbs*. PhD dissertation, State University of New York at Stony Brook.

Vaux, Bert (2000). Disharmony and derived transparency in Uyghur vowel harmony. *NELS* **30**, 671–698.

Vaux, Bert (2008). Why the phonological component must be serial and rule-based. In Bert Vaux & Andrew Nevins (eds.) *Rules, constraints, and phonological phenomena*. Oxford: Oxford University Press, 20–60.

Vaux, Bert (2011). Language games. In John A. Goldsmith, Jason Riggle & Alan C. L. Yu (eds.) *The handbook of phonological theory*, 2nd edition. Oxford: Blackwell, 722–750.

Washington, Jonathan, Ilnar Salimzianov, Francis M. Tyers, Memduh Gökirmak, Sardana Ivanova & Oğzhan Kuyrukçu (2019). Free/open-source technologies for Turkic languages developed in the Apertium project. In L. Sh. Kubedinova (ed.) *Sed'maya Mezhdunarodnaya Konferentsiya po Komp'yuternoy Obrabotke Tyurkskikh Yazykov TurkLang 2019 [Seventh International Conference on Turkic Language Processing (TurkLang 2019)]*. Kazan: Publishing House of the Academy of Sciences of the Republic of Tatarstan, 30–71.

Zhang, Jie (2019). Speakers treat transparent and opaque alternation patterns differently: evidence from Chinese tone sandhi. *WCCFL* **36**, 22–40.

Zuraw, Kie (2000). *Patterned exceptions in phonology*. PhD dissertation, University of California, Los Angeles.

Zuraw, Kie (2010). A model of lexical variation and the grammar with application to Tagalog nasal substitution. *NLLT* **28**, 417–472.

Zuraw, Kie (2016). Polarized variation. *Catalan Journal of Linguistics* **15**, 145–171.

https://doi.org/10.1017/S0952675725100213 Published online by Cambridge University Press