

ARTICLE

# Non-concatenative morphological domains constrain phonotactics: a case study of Egyptian Arabic

Lily Xu 

Department of Linguistics, University of California, Los Angeles, Los Angeles, CA, USA  
Email: [lilyokc@ucla.edu](mailto:lilyokc@ucla.edu)

**Received:** 31 July 2023; **Revised:** 13 April 2024; **Accepted:** 21 August 2024

**Keywords:** lexicon study; wug test; Arabic; morphology–phonology interface; non-concatenative morphology; morphosyntax

## Abstract

Phonotactic patterns are commonly constrained by morphology. In English, for example, non-homorganic nasal–stop sequences are disallowed within morphemes but may occur across morpheme boundaries. This article demonstrates that similar effects of morphology on phonotactics can be found with non-concatenative morphology, even though they involve morphological domains that are more difficult to identify on the surface. Specifically, vowel alternation in a class of Egyptian Arabic verbs is affected by gradient phonotactic restrictions on consonant–vowel co-occurrence. However, such restrictions are only active in the imperfective form (e.g., [–rgaʕ] ‘return.IPFV’), not the perfective (e.g., [riɣiʕ] ‘return.PFV’). Using a lexicon study and a wug test, I show that this pattern is in fact bounded by morphological domains and is reliably generalised by speakers when deriving novel forms. I compare accounts of this effect that differ on whether they require abstract morphosyntactic representations and non-concatenative morphemes and discuss their implications.

## Contents

1. Introduction . . . . .	2
2. Vowel alternation in Egyptian Arabic . . . . .	2
2.1. The puzzle: asymmetrical consonant effects . . . . .	3
2.2. Morphological structure of the paradigm: surface correspondence vs. morphosyntax . . . . .	4
3. Lexicon study . . . . .	6
3.1. Data collection . . . . .	6
3.2. Modelling . . . . .	6
3.3. Predicting the perfective . . . . .	8
3.4. Predicting the imperfective . . . . .	10
3.5. Comparing the models . . . . .	12
3.6. Summary . . . . .	12
4. Wug test . . . . .	13
4.1. Methods . . . . .	13
4.2. Results . . . . .	15
4.3. Summary . . . . .	16
5. Analysis and discussion . . . . .	17
5.1. A morphosyntactic account . . . . .	18
5.2. Surface correspondence accounts . . . . .	20
5.3. Discussion . . . . .	21
6. Conclusion . . . . .	22



## 1. Introduction

It is common cross-linguistically for the application of a phonotactic restriction to be constrained by some morphological domain (Trubetzkoy 1939; Gouskova 2018). Many phonotactic restrictions, both local and non-local, have been found to be active within morphemes and not between morphemes. For example, in English, non-homorganic nasal–stop sequences are not allowed within a morpheme, but are possible if the sequences cross a morpheme boundary, for example, *wi[ŋ+d]*. Similarly, laryngeal co-occurrence restrictions in Bolivian Aymara are enforced morpheme-internally but not across morpheme boundaries (Gallagher *et al.* 2019). Other restrictions are sensitive to even finer distinctions in morphological domains. In Kaqchikel, geminates can be derived from adjacent consonants across the boundary between a stem and a derivational morpheme, but degemination is obligatory if the consonants are separated by the boundary between a stem and an inflectional morpheme (Bennett 2018). It is also common for phonotactic restrictions to be bounded by the word; one example of this type is vowel harmony in Finnish (Suomi *et al.* 1997).

A question that has been hitherto under-explored is whether similar patterns of morphological domain-bounded phonotactics can also occur with non-concatenative morphology and, if so, how these patterns are represented. Nonlinear morphophonological processes such as ablauting, mutation and melodic overwriting can make morphological domains difficult to identify on the surface. Therefore, the representation of such patterns can be expected to be more complicated, requiring structures that are more abstract than linearly defined morpheme or word boundaries.

In this article, I identify such a case in Egyptian Arabic, focusing on a pattern of vowel alternation in verbs used to mark tense/aspect distinctions. This alternation is non-concatenative, as illustrated in the pair **[rigiʕ]** ‘return (perfective)’ – **[-rgaʕ]** ‘return (imperfective)’. As is typical in Semitic languages like Arabic, the root contributes the consonants, and tense/aspect determines the vowels and syllable structure. With such structures, it is difficult to determine any morphological relation on the surface. Previous work has reported that these root–vowel combinations are affected by phonotactic restrictions on consonant–vowel co-occurrence, such as the gradient preference for [a] to occur next to pharyngeal consonants (e.g., Abdel-Massih *et al.* 1979). Further, there is an asymmetry: such restrictions are found in some morphological contexts and not others. Specifically, Abdel-Massih *et al.* (1979) only found them in imperfective forms but not perfective forms. Using a lexicon study and a wug test, I show that this asymmetry is best explained as a case of morphological domain effect on phonotactics. I also compare two analyses, in which these morphological domains are derived from either the morphosyntactic structure or surface-based paradigmatic correspondences. While I show that these analyses are both compatible with the data, they are distinct in their assumptions and implications for learning.

I present the empirical pattern of interest and relevant background information in §2. §3 investigates phonotactic restrictions and evidence for morphological structure based on lexicon statistics. §4 presents wug-test results that demonstrate speakers’ morphological and phonological knowledge of the paradigm. §5 compares two analyses of the Egyptian Arabic verbal paradigm that are compatible with the lexicon and wug-test data, and discusses the implications of either approach. §6 concludes the article.

## 2. Vowel alternation in Egyptian Arabic

Arabic verbs are formed exclusively from about a dozen fixed prosodic templates that carry grammatical information. Each template is called a *wazn* (plural: *awzaan*). *Awzaan* combine with consonantal roots, generally thought to contribute the core meaning, to form a word. Two examples of verbal *awzaan* in Egyptian Arabic are shown in (1). The root *s-m-ʕ* (bolded) denotes the meaning ‘hearing’ and combines with *wazn* I to form the verb ‘to hear’. *Wazn* II is marked by the medial geminate consonant and generally denotes causativity.

- (1) a. **simiʕ**      ‘to hear’      wazn I      CVCVC  
b. **sammaʕ**      ‘to cause to hear’      wazn II (causative)      CVCCVC

Wazn I is described as the non-derived, or basic, wazn as it has the simplest morphology and has no unifying morphosyntactic or semantic properties. All other awzaan can be analysed as derived from wazn I, either directly or indirectly (e.g., McCarthy 1993). Wazn I is also unique in having idiosyncratic vowel alternations between the two tense/aspect forms, whereas in other awzaan, the vowels do not alternate between the tense/aspect forms for the same root. As shown in Table 1, perfective wazn I verbs have either [a] or [i] in both syllables, while imperfective wazn I verbs have [a], [i] or [u] as the stem vowel. All possible vowel alternations between these two forms are attested, though perfective [i] – imperfective [u] is relatively rare.

2.1. The puzzle: asymmetrical consonant effects

While this vowel alternation pattern seems to be arbitrary at a first glance, works on Egyptian Arabic as well as on several other dialects report partial predictability based on phonological factors.<sup>1</sup> These phonological factors include root consonants and correspondences between the vowels of different forms. I first discuss the root consonants and then discuss the vowel correspondences in §2.2. In Hijazi Arabic, the presence of a guttural (pharyngeal, glottal or uvular) root consonant leads to a preference for [a], and a pharyngealised alveolar leads to a preference for [u] both in lexicon trends and in nonce-word experiment results (Ahyad 2019; Ahyad & Becker 2020). Similar gradient effects of gutturals have been reported in the lexicon of MSA (McCarthy 1991), and pharyngealised alveolar effects are reported for Egyptian Arabic (Abdel-Massih *et al.* 1979). Similar effects are shown to be categorical in Palestinian Arabic (Herzallah 1990).

These frequently cited consonant–vowel co-occurrence effects are phonetically motivated. Guttural consonants involve retraction of the tongue root, which creates an affinity for low vowels like [a] over high vowels like [i, u] (McCarthy 1994). Pharyngealised alveolars have the effect of lowering the F2 of surrounding vowels (Norlin 1987; Laufer & Baer 1988), explaining the preference for [u] over [i, a].

A curious commonality of these findings is the asymmetry between the two tense/aspect forms, where consonant effects on vowel choice are typically reported only in the imperfective and not in the perfective (cf. Blanc 1964, where the opposite pattern is reported in Muslim Baghdadi Arabic). This asymmetry is most likely not due to differences in the segmental contexts, since the alternating vowels are embedded between the second and third consonants in the underlined portion of both forms, CCVC vs. CVCVC. (Potential effects that uniquely come from the first consonant will be discussed in §3.) It is possible that the morphological structure of the wazn I verb paradigm provides some explanations

Table 1. Vowel alternation in wazn I verbs from Egyptian Arabic (3SG masculine form).

	Perfective	Imperfective		
a/a	kasar	-ksar	‘break’	
a/i	katab	-ktib	‘write’	
a/u	ɣarag	-ɣrug	‘go out’	
i/a	rigiʕ	-rgaʕ	‘return’	
i/i	libis	-lbis	‘dress’	
i/u	sikin	-skun	‘live’	(rare)

<sup>1</sup>One instance where a semantic property (specifically stativity) allows identification of one alternation pattern has been noted in Modern Standard Arabic (MSA; McCarthy 1991), but most of the other sources of predictability are phonological.

for this asymmetry, but no such accounts have been proposed because the paradigm structure itself is contended, as I discuss below.

## 2.2. Morphological structure of the paradigm: surface correspondence vs. morphosyntax

I discuss here two ways the morphological structure of the *wazn I* verb paradigm has been investigated in the literature. One line of work has focussed on surface-based mappings between vowels in the imperfective and perfective forms, but this approach has not led to any definitive results. It is possible to account for the vowel alternation in MSA with great accuracy if the imperfective form is taken as the base that predicts the perfective form (McOmber 1995). However, Guerssel & Lowenstamm (1996) have also shown that greater accuracy can be achieved by analysing the perfective form as the base from which imperfective form is predicted, though their analysis involves positing an abstract underlying vowel for a subclass of verbs. In either case, it is not clear if such surface correspondences are even a part of speakers' grammars. Ahyyad (2019) found that [i] in perfective forms predicts [a] in imperfective forms 90% of the time in Hijazi Arabic, but speakers completely fail to generalise this trend in wug tests. Whether speakers generalise the opposite imperfective-to-perfective pattern has not been tested before now.

Data from morphosyntactic studies on the imperfective and perfective verb forms in Arabic offers new insight into this issue. While the perfective and imperfective in Arabic are often treated as two different tense/aspect forms, syntactic and acquisition studies give evidence that the imperfective should actually be understood as an infinitive form. Thus, the imperfective occupies a smaller morphosyntactic domain, which is embedded within the morphosyntactic domain occupied by the perfective. Crucially, this suggests that the perfective is more likely to be derived from the imperfective than the other way around.

Benmamoun (1999) shows that in Arabic dialects like MSA and Egyptian Arabic, the perfective form of the verb is always used in past tense clauses (2), whereas the imperfective form occurs in a wide range of contexts, including present tense, as in (3a); future tense, as in (3b); after modals and auxiliaries, as in (3c) and (3d); and in infinitival clauses, as in (3e).

### (2) *Distribution of the perfective in Egyptian Arabic*

raʔas<sup>ʕ</sup>  
dance.PERF.3SG.M  
'He danced.'

### (3) *Distribution of the imperfective (underlined> in Egyptian Arabic*

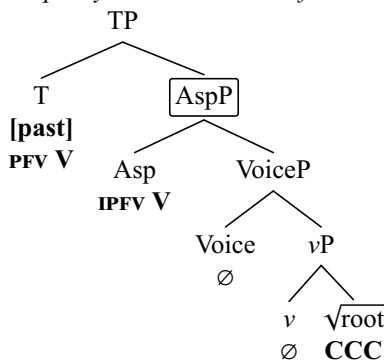
- a. bi-ju-rʔus<sup>ʕ</sup>  
PRES.INDICT-3SG.M-dance.IMP  
'He dances/is dancing.'
- b. ha-ju-rʔus<sup>ʕ</sup>  
FUT-3SG.M-dance.IMP  
'He will dance.'
- c. laazim ju-rʔus<sup>ʕ</sup>  
must 3SG.M-dance.IMP  
'He must dance.'
- d. kaan ju-rʔus<sup>ʕ</sup>  
was 3SG.M-dance.IMP  
'He was dancing/used to dance.'

- e. raah                    in-naadi    ʕaʕaan ju-rʔus<sup>ʕ</sup>  
 go.3SG.M.PERF DEF-club so.that 3SG.M-dance.IMP  
 ‘He went to the club so that he can dance.’

While the perfective verb stem is always specified for past tense, Benmamoun notes that in the sentences containing an imperfective verb, the imperfective verb stem itself is never specified for tense; tense is always conveyed by some other particle (e.g., *ha-* for future). This suggests that the ‘imperfective’ in Arabic is actually an infinitival verb or present participle. Acquisition studies offer additional support: Aljenaie (2010) found that Kuwaiti children in the age range of 1;8–3;1 use the bare imperfective stem as a non-finite form (see also Omar 1973).

In theories where morphological structure is syntactic in nature, such as Distributed Morphology (Halle & Marantz 1993), word formation happens bottom-up in the syntax. Under this framework, the Arabic imperfective would be used as the input to derive the perfective. The perfective form, always specified for past tense, must have an underlying structure that includes the tense head T and its projection. The imperfective form is never specified for tense, so it cannot contain T, meaning that it has simpler morphosyntactic structure than the perfective form. (4) shows a sketch of the underlying structure of Arabic *wazn I* verbs, adapted from Tucker (2011):

(4) *Morphosyntactic structure of Arabic wazn I verbs*



The functional head *v* combines with a consonantal root to form a verb, which is then selected by a Voice head. This structure then gets selected by an Asp head, which I treat as the position where the imperfective vowel resides. Whereas previous approaches either treat the imperfective vowel as the Voice head (Arad 2005 for Hebrew) or a composite Tense/Aspect/Voice (TAV) head (Tucker 2011 for MSA), this modification was made because vowel change does not signal active/passive alternations in Egyptian Arabic. Furthermore, since the imperfective verb always combines with other elements that carry tense information, its structure should only include AspP (boxed), not T. On the other hand, the perfective verb is always inflected for past tense. Therefore, the perfective vowel is treated as a T head specified for [past], and the structure of the perfective verb will minimally include the T head.

Following such an analysis, the vowel alternation in the perfective and imperfective are allomorphic processes limited by locality constraints, which allow only adjacent overt morphemes to affect each other in phonologically conditioned alternations (Embick 2010; Marantz 2013). Evidence of this locality constraint has been found in case studies involving both concatenative morphology (Northern Paiute, Toosarvandani 2016; Nez Perce, Deal & Wolf 2017; Greek, Paparounas 2021) and non-concatenative morphology (Hebrew, Kastner 2016, 2019). Imperfective vowel choice is conditioned by the root consonants because there is no overt intervening structure between the imperfective vowel and the root. Perfective vowel choice is not conditioned by the root consonants because there is overt intervening structure, namely the imperfective vowel. The asymmetry between imperfective and perfective verb forms in whether vowels are predictable from adjacent consonants is therefore an instance of morphologically constrained phonotactics.

In the rest of the article, I present results first from a lexicon study, then from a wug test investigating predictors of the vowel alternation in Egyptian Arabic *wazn* I verbs. Consistent with previous studies, I find that consonants play a role in the choice of vowel for imperfective but not perfective forms; these effects are successfully learned by speakers and generalised onto novel forms. Additionally, I identify vowel correspondence effects in the lexicon and show that speakers selectively generalise them in ways that align with the morphosyntactic structure of the paradigm proposed above. I argue that the vowel alternation in Egyptian Arabic is a case where non-concatenative morphological domains constrain phonotactics, just as has been found for concatenative morphological domains. However, non-concatenative morphological domains are unique in that they are by default abstract and not readily evident on the surface.

### 3. Lexicon study

In this section, I present various phonological factors that are predictive of the vowel alternation pattern based on a lexicon study of Egyptian Arabic verbs and use logistic regression models to assess them quantitatively. The lexicon study has three goals. First, I examine the consonantal effects on the vowel alternation pattern and whether they are indeed confined to the imperfective form (as proposed in Abdel-Massih *et al.* 1979). Second, I investigate predictive trends in vowel-to-vowel correspondences between the two forms. Third, I compare the regression models as a way to assess which paradigm organisation yields maximum predictability.

#### 3.1. Data collection

I compiled a corpus of 330 *wazn* I verbs in Egyptian Arabic (available on OSF at [https://osf.io/tdr6x/?view\\_only=4aec9be4f2204849ac81932103fe3dd9](https://osf.io/tdr6x/?view_only=4aec9be4f2204849ac81932103fe3dd9)). Only words with the prosodic shape CVCVC in the perfective (traditionally known as ‘sound verbs’) were included, as verbs with other prosodic shapes have different vowel alternation patterns. The words were extracted from the online dictionary *Lisaan Masry* (Green 2007) and checked with a native speaker.

Since the focus of this study is on the colloquial Arabic dialect spoken in Cairo, it is important to acknowledge that nearly all speakers of Egyptian Arabic also use MSA to some degree (Ferguson 1959; Eid 2007). To control for the differences in the morphophonology of these two varieties of Arabic, verbs that had clear features of MSA (e.g., having [a] in the imperfective agreement prefix, as in [ja-sbat] ‘he proves’, as opposed to [i] or [u] in the prefix as in Egyptian Arabic) were excluded.

The breakdown of the corpus by vowel alternation pattern is shown in Table 2. The perfective vowel is either [a] or [i], and the vowels in the two syllables (CVCVC) are always identical. The single imperfective stem vowel (-CCVC) can be [a], [i], or [u]. Note that imperfective stems always occur with an agreement prefix. The prefix vowel is generally [i] but optionally undergoes harmony when the stem vowel is [u]. All combinations of vowels in the two forms are attested, though the perfective [i] – imperfective [u] combination is very rare.

#### 3.2. Modelling

Logistic regression models with vowel and consonant predictors were fitted using the *nnet* package (Venables & Ripley 2002) in R (R Core Team 2021). Models predicting perfective vowel choices had two levels of the dependent variable ([a] or [i]), and models predicting imperfective vowel choices had three ([a], [i] or [u]), which called for multinomial logistic regression. Vowel predictors based on the vowels in the form other than the one that the model predicts were used to investigate whether vowel-to-vowel correspondence plays a role in the paradigm. The consonant predictors assess whether a consonant with a particular place of articulation is present in the word; they were used to investigate whether there is avoidance of more marked consonant–vowel combinations. Place of articulation was

**Table 2.** *Vowel alternation in wazn I verbs from Egyptian Arabic (Table 1) with counts added (a total of 330 verbs).*

	Perfective	Imperfective	Meaning	Count
a/a	kasar	ji-ksar	‘break’	71
a/i	katab	ji-ktib	‘write’	30
a/u	ḡarag	ji/ju-ḡrug	‘go out’	67
i/a	rigiʕ	ji-rgaʕ	‘return’	89
i/i	libis	ji-lbis	‘dress’	66
i/u	sikin	ji/ju-skun	‘live’	7

**Table 3.** *Consonant natural classes by place of articulation.*

Labial	Alveolar		Palatal	Velar	Uvular	Pharyngeal	Glottal
	Plain	Pharyngealised					
b, f, m	t, d, s, z, n, l	tʕ, dʕ, sʕ, zʕ, r	ʃ	k, g	q, ɣ, ʁ	ħ, ʕ	ʔ, h

chosen to be the main consonant property investigated because consonant effects on vowel alternations in other dialects generally involve place features (e.g., Ahyad & Becker 2020). The place of articulation classes are listed in Table 3. Words containing two glide consonants, j, w, were excluded from this study because they generally have different prosodic shapes and vowel patterns.

Pharyngealised alveolar consonants, also known as emphatics, are characterised by a secondary constriction in the upper pharynx (Ghazeli 1977; Laufer & Baer 1988). They exert a lowering and backing effect known as *emphasis spreading* on all vowels within the same phonological word, as well as affecting formant transitions surrounding the consonants (Lehn 1963; Norlin 1987; Watson 2002). The status of the tap /r/ in Egyptian Arabic is debated. Specifically, it has been argued that there is a contrast between a non-emphatic /r/ and an emphatic one (Youssef 2019), which causes emphasis spreading like other pharyngealised alveolars (Younes 1994; Watson 2002). My corpus based on the dictionary did not include such a contrast, and I grouped /r/ with pharyngealised alveolars.<sup>2</sup>

In the imperfective forms (–C<sub>1</sub>C<sub>2</sub>VC<sub>3</sub>), the second and the third root consonants are directly adjacent to the vowel, so one might expect stronger effects for them compared to the first consonant. To test this, I tried running models with positional consonant predictors, which specify whether a given natural class is present in each of the three positions in the consonantal root. However, because of the relatively small sample size, the models with positional predictors overfitted (as indicated by extremely large coefficient values and poor accuracy in cross-validation) and thus were not informative. I will return to this issue of consonant position when discussing the results below.

Besides quantitatively evaluating the role of vowels and consonants in predicting the vowel alternation, the models were designed with an additional purpose, namely to model different paradigm organisations. Numerous case studies (e.g., Albright 2005; Kuo 2020) have shown that speakers preferentially select a paradigm member as the base to derive other forms in ways that maximise predictability. Specifically, I evaluate two possibilities that both instantiate such unidirectional paradigm mappings but in opposite directions. The first is that speakers use the imperfective forms as the bases to derive the perfective forms. This possibility is instantiated by a model that predicts perfective vowel choices

<sup>2</sup>I ran models without /r/ tokens to verify if the effects for pharyngealised alveolars were affected by any special behaviour of /r/ and found no qualitative difference to the models reported in this article – the pharyngealised alveolars had no effect in the perfective model and significantly favour [a] over [i] in the imperfective model. These models can be found on the [OSFpage](#).



Table 4. Perfective and imperfective vowel frequencies.

Perfective V	Imperfective V			Total
	a	i	u	
a	71 (44%)	30 (31%)	67 (91%)	168 (51%)
i	89 (56%)	66 (69%)	7 (9%)	162 (49%)
Total	160 (48%)	96 (29%)	74 (22%)	

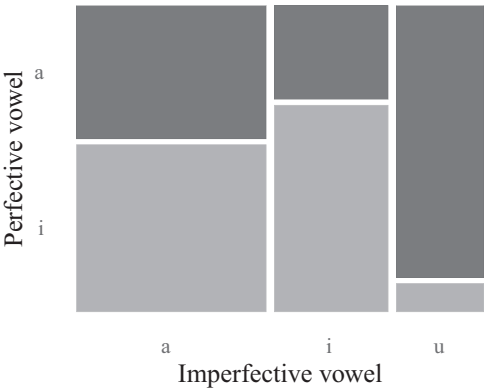


Figure 1. Breakdown of perfective vowel by imperfective vowel.

based on all and only the phonological predictors that are available in the imperfective forms, namely the root consonants and the imperfective vowels (§3.3). The second possibility is that speakers use the perfective forms as the bases to derive the imperfective forms. The model that instantiates this possibility uses root consonants and perfective vowels to predict imperfective vowel choices (§3.4). The models are then compared on the basis of predictive accuracy and explanatory power (§3.5).

3.3. Predicting the perfective

Vowel-to-vowel correspondences play a role in predicting the perfective vowel. As shown in Table 4 and plotted in Figure 1, while the overall distribution of [a]- and [i]-perfectives is even, the breakdown of perfective vowel choice for each imperfective vowel differs. There is a strong preference for [a]-perfectives (91%) when the imperfective vowel is [u]. When the imperfective vowel is [i] or [a], there are moderate preferences for [i]-perfectives (69% and 56%, respectively).


On the other hand, perfective vowel choice does not seem to be influenced by consonant place of articulation, as shown in Table 5 and Figure 2. Note that each verb in the corpus is represented three times in Table 5, once for each consonant in the root. Having a root consonant of a particular place of articulation tends not to bias the perfective vowel distribution away from the overall distribution. While pharyngealised alveolars, palatals and uvulars show a preference for [a], the effects are quite small.

Since the two vowels in a perfective form (C<sub>1</sub>V<sub>1</sub>C<sub>2</sub>V<sub>2</sub>C<sub>3</sub>) are identical in Egyptian Arabic, it is possible that only one of the vowels is represented and the other derived by copying. If this is the case, we may expect the consonant effects on vowel distribution to be positionally specific. For instance, if only V<sub>1</sub> is represented, it is possible that consonant effects are only found with the adjacent C<sub>1</sub> and C<sub>2</sub>. Table 6 presents the vowel distributions in the presence of pharyngeals and pharyngealised alveolars at each consonant position. There is no clear evidence for any positional effect. For example, perfectives



**Table 5.** *Effects of consonant natural classes on perfective vowel distribution.*

	a	i
Labial	90 (51%)	88 (49%)
Plain alveolar	128 (45%)	155 (55%)
Pharyngealised alveolar	107 (60%)	71 (40%)
Palatal	23 (66%)	12 (34%)
Velar	24 (44%)	31 (56%)
Uvular	34 (68%)	16 (32%)
Pharyngeal	59 (47%)	66 (53%)
Glottal	39 (48%)	42 (52%)
Total	51%	49%



**Figure 2.** *Effects of consonant natural classes on perfective vowel distribution.*

**Table 6.** *Effects of consonant natural classes by position in perfectives: pharyngeals and pharyngealised alveolars.*

		a	i
Pharyngeal	C <sub>1</sub>	20 (44%)	25 (56%)
	C <sub>2</sub>	17 (45%)	20 (55%)
	C <sub>3</sub>	22 (51%)	21 (49%)
	Total	59 (47%)	66 (53%)
Pharyngealised alveolar	C <sub>1</sub>	24 (53%)	19 (47%)
	C <sub>2</sub>	46 (59%)	31 (40%)
	C <sub>3</sub>	37 (63%)	21 (37%)
	Total	107 (60%)	71 (40%)

with a pharyngeal as the first root consonant have [a] 44% of the time. This number is 46% if the pharyngeal is the second consonant and 51% if it's the third.

These two observations (that perfective vowel choice is influenced by vowel-to-vowel correspondences but not by consonant–vowel co-occurrences) are borne out by the regression model, shown in Table 7. Positive coefficients indicate a preference for [i]-perfectives, whereas negative coefficients

**Table 7.** Imperfective-to-perfective model. Residual deviance: 378.35; pseudo- $R^2$ : McFadden 0.173, CoxSnell 0.213, Nagelkerke 0.284; cross-validation accuracy: 0.639. Significant factors are **bolded**.

Predictor	Coefficient	<i>p</i>
<b>Imperfective [i]</b>	0.596	0.043 *
Velar	0.467	0.175
Plain alveolar	0.255	0.275
Pharyngealised alveolar	0.184	0.436
Glottal	0.139	0.633
Labial	−0.023	0.920
Pharyngeal	−0.128	0.596
Palatal	−0.323	0.438
Uvular	−0.653	0.075 .
<b>Imperfective [u]</b>	−2.469	0.000 ***

**Table 8.** Effects of consonant natural classes on imperfective vowel distribution.

	a	i	u
Labial	73 (41%)	66 (37%)	39 (22%)
Plain alveolar	130 (46%)	104 (37%)	49 (17%)
Pharyngealised alveolar	83 (47%)	31 (17%)	64 (36%)
Palatal	14 (40%)	8 (23%)	13 (37%)
Velar	21 (38%)	20 (36%)	14 (25%)
Uvular	21 (42%)	13 (26%)	16 (32%)
Pharyngeal	87 (70%)	28 (22%)	10 (8%)
Glottal	49 (60%)	15 (19%)	17 (21%)
Total	49%	29%	22%

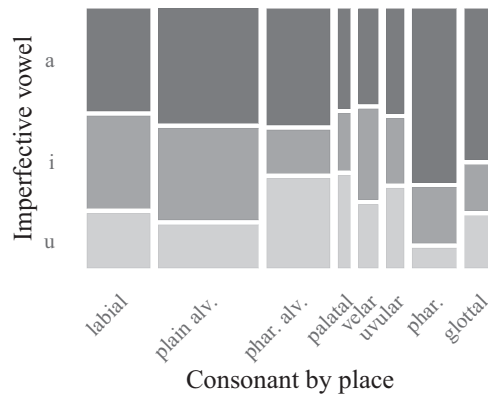
indicate a preference for [a]-perfectives. Note that there are two imperfective vowel predictors, where the effect of [i] and [u] are each compared to that of [a].

There are only two significant predictors in the above model, both of which are vowel predictors. First, imperfective [u] results in a large preference for perfective [a]. Second, there is a moderate preference for perfective [i] when there is imperfective [i]. None of the consonant predictors are significant.

3.4. Predicting the imperfective

Imperfective vowel choices also show sensitivity to perfective vowels. Notably, there is a preference for imperfective [u] when the perfective vowel is [a]. This trend is less drastic than the same preference in the opposite direction (see Table 4 and Figure 1). Perfective [i] results in a slight preference for imperfective [i] and dispreference for imperfective [u].

Additionally, there are noticeable consonant effects on the imperfective vowel distributions, as shown in Table 8 and Figure 3. Compared to the overall vowel distribution, the presence of pharyngeal and glottal root consonants is associated with more [a]-imperfectives (70% and 60%, respectively);



**Figure 3.** Effects of consonant natural classes on imperfective vowel distribution.

**Table 9.** Perfective-to-imperfective model. Residual deviance: 515.22; pseudo- $R^2$ : McFadden 0.253, CoxSnell 0.411, Nagelkerke 0.469; cross-validation accuracy: 0.606. Significant factors are **bolded**.

	Predictor	Coefficient	<i>p</i>
i vs. a	<b>Labial</b>	0.854	0.001 **
	<b>Plain alveolar</b>	0.559	0.034 *
	Velar	0.356	0.335
	Palatal	−0.014	0.979
	Uvular	−0.364	0.401
	Perfective [a]	−0.496	0.108
	<b>Pharyngealised alveolar</b>	−0.846	0.003 **
	<b>Glottal</b>	−1.315	0.000 ***
	<b>Pharyngeal</b>	−1.601	0.000 ***
u vs. a	<b>Perfective [a]</b>	2.594	0.000 ***
	Velar	0.007	0.988
	Pharyngealised alveolar	−0.118	0.715
	Palatal	−0.275	0.583
	<b>Labial</b>	−0.811	0.017 *
	<b>Plain alveolar</b>	−1.022	0.002 **
	<b>Uvular</b>	−1.043	0.020 *
	<b>Glottal</b>	−1.530	0.000 ***
	<b>Pharyngeal</b>	−2.884	0.000 ***

labials and plain alveolars with more [i]-imperfectives (37% for both); and pharyngealised alveolars with more [u]-imperfectives (36%).

A multinomial regression model predicting imperfective vowels is shown in Table 9. The model shows pairwise comparisons for [i] vs. [a] and [u] vs. [a], where positive coefficients indicate a preference for [i]- or [u]-imperfectives over [a]. Note that there is just one perfective vowel predictor, where the effect of [a] is compared to that of [i].

Consonant predictors make a substantial contribution in this model, consistent with the various consonantal effects observed above. Pharyngeals and glottals strongly favour imperfective [a] over

**Table 10.** *Effects of consonant natural classes by position in imperfectives: pharyngeals and pharyngealised alveolars.*

		a		i	u
Pharyngeal	C <sub>1</sub>	15	(33%)	21 (47%)	9 (20%)
	C <sub>2</sub>	29	(78%)	7 (19%)	1 (3%)
	C <sub>3</sub>	43	(100%)	0 (0%)	0 (0%)
	Total	87	(70%)	28 (22%)	10 (8%)
Pharyngealised alveolar	C <sub>1</sub>	23	(53%)	8 (19%)	12 (28%)
	C <sub>2</sub>	35	(45%)	15 (19%)	27 (35%)
	C <sub>3</sub>	25	(43%)	8 (14%)	25 (43%)
	Total	83	(47%)	31 (17%)	64 (36%)

[i, u], and uvulars strongly favour [a] over [u]. Pharyngealised alveolars strongly favour [a] over [i] but exhibit no preference between [a] and [u]. As discussed in §2, these effects are consistent with previous findings and have phonetic motivations.

A gross inspection of vowel distribution based on the presence of consonants in specific positions finds striking positional effects for pharyngeals. Table 10 shows that the preference for imperfective [a] is very strong when the imperfective vowel is immediately adjacent to a pharyngeal (as C<sub>2</sub> or C<sub>3</sub>) but is absent when the pharyngeal (as C<sub>1</sub>) is separated from the vowel by an intervening consonant. Interestingly, strong effects of consonant position are not found in verbs with pharyngealised alveolars, possibly because pharyngealised alveolars can influence vowel qualities across the entire phonological word, not just adjacent vowels (e.g., Watson 2002).

3.5. *Comparing the models*

The two models’ goodness-of-fit was compared in order to assess which direction of paradigm mapping yields greater predictability. The first method I used compares the accuracy of the two models with *k*-fold cross-validation (*k* = 5). The dataset was randomly divided into five parts, and each model was run on four of the parts and tested on the other. This process was repeated for all five parts, and the average model accuracy from all the trials was calculated by comparing the model predictions on the testing data in each run with the corpus. The imperfective-to-perfective model had a higher average accuracy (0.639) than the perfective-to-imperfective model (0.606), but the difference is very small.

I then assessed the two models relative to chance-level performance and found that the model predicting the imperfective is superior. The reason is that the random baseline accuracy for this model is 0.33, since there are three imperfective vowel choices, whereas the random baseline accuracy of the model predicting the perfective (between two alternatives) is 0.5. This intuition is also supported by pseudo-*R*<sup>2</sup> measures. Since the two models predict different dependent variables, they cannot be directly compared with AIC or likelihood measures, and pseudo-*R*<sup>2</sup> measures are appropriate. All three pseudo-*R*<sup>2</sup> measures (McFadden, Cox & Snell and Nagelkerke) calculated were higher for the perfective-to-imperfective model, which suggests that it is superior in terms of accounting for the variability in the data.

3.6. *Summary*

Modelling results based on lexicon data showed that both perfective and imperfective vowels are partially predictable from phonological factors. A crucial difference between the models lies in the types

of relevant phonological predictors. The imperfective vowel can be predicted primarily by consonants in ways governed by the phonotactic well-formedness of consonant–vowel sequences, with a minor role played by correspondences to the perfective vowel. The perfective vowel can be predicted based solely on the imperfective vowel. Furthermore, model comparison metrics showed that the model predicting the imperfectives from the perfectives had higher explanatory power, even though the opposite model had a slightly higher accuracy. The next section examines whether these predictive trends are generalised by native speakers in a wug test.

#### 4. Wug test

While the predictive trends for the vowel alternation in the lexicon were all gradient, a large body of research has demonstrated that speakers are capable of learning gradient generalisations for phonological patterns and have a tendency to match lexicon frequency in wug-test responses (Zuraw 2000; Ernestus & Baayen 2003; Hayes & Londe 2006). Furthermore, morphological factors such as paradigm directionality have been shown to constrain the success (or failure) of frequency matching (Jun & Albright 2017; Kuo 2020). As discussed in §2.2, Ahyad (2019) found that Hijazi Arabic speakers failed to generalise salient correspondence patterns from perfective vowels to imperfective vowels. This article builds on and extends this line of research. In this section, I will present the results of a wug-test experiment investigating whether Egyptian Arabic speakers generalise the various gradient trends found in the lexicon that are significant predictors for the vowel alternation.

The lexicon trends fall into three types: consonant effects and perfective vowel effects on imperfective vowel choice, and imperfective vowel effects on perfective vowel choice. Since the vowel effects require the use of one verb form to predict another, how speakers generalise these effects can be taken as evidence for their representation of the paradigm. If they make use of both types of vowel effects, this is evidence that they have multidirectional mappings between forms. Alternatively, they may selectively learn one direction of vowel effects.

The consonantal effects are informative of morphological structure for a different reason. Since the surrounding consonantal contexts for the alternating vowels in the perfective (CVCVC) and imperfective (CCVC) forms overlap, the occurrence of the consonant effects in only the imperfective would have to be explained by morphological and not phonological factors. Because of the obscurity of the morphological structure on the surface, however, it is also possible that speakers fail to learn the restrictedness of the consonant effects, either extending them to the perfective or ignoring them completely.

#### 4.1. Methods

**4.1.1. Participants.** Twenty-nine speakers of Egyptian Arabic from Cairo and its surrounding were recruited via various online platforms (Prolific, Facebook, etc.). The participants were presented with a brief language background questionnaire, asking for the following information: gender, age, place of birth, the dialects of Arabic that they speak, and the relative frequency of usage of these dialects in their daily life. Because the participants are likely to also speak MSA, a few phonological features were used to control for dialect when analysing the data. First, the colloquial Egyptian dialect uses the vowel [i] in agreement prefixes on imperfective verbs, whereas MSA uses [a]. Thus, a participant would be excluded if they used [a] as their prefix vowel. Participants whose dominant dialect was not Cairene Egyptian Arabic were identified by their pronunciation of the letters *qaf* <ق> and *jiim* <ج>; in Cairene Arabic, these are pronounced as [ʔ] and [g], respectively, whereas the dialects of Upper Egypt realise them as [g] and [ʒ]. Responses from six additional participants were discarded for these reasons.

**4.1.2. Materials.** A list of nonce verbs in both perfective (C<sub>1</sub>VC<sub>2</sub>VC<sub>3</sub>) and imperfective (–C<sub>1</sub>C<sub>2</sub>VC<sub>3</sub>) forms was compiled. The list contained 50 distinct roots, each occurring with the two perfective vowels

**Table 11.** Example stimuli with the roots *l-b-ḥ* and *l-b-z*.

	V	Pharyngeal C <sub>3</sub>	Non-pharyngeal C <sub>3</sub>
Perfective	a	labah	labaz
	i	libih	libiz
Imperfective	a	lbaḥ	lbaz
	i	lbiḥ	lbiz
	u	lbuḥ	lbuz

and three imperfective vowels, yielding a total of 250 nonce verbs. All possible vowel choices in both forms were included. Only a subset of consonants was included to make it easier to observe and test the effects of consonant class. The first two root consonants were always drawn from labials b, f, m, alveolar obstruents t, d, s, z, and alveolar sonorants l, n. These consonants showed a preference for imperfective [i] but did not affect the vowel distribution as strongly as other consonants, making them ideal control consonants. In keeping with OCP co-occurrence restrictions (McCarthy 1994), a root never contained two segments drawn from the same class. In half of the roots, the third consonant was drawn from the control set as well, while for the other half, the third consonant was a pharyngeal ḥ, ʕ. Pharyngeals only occurred as the third consonant because their effect on imperfective vowel was observed most strongly in this position (see Table 10).

All nonce verbs that could be created from a pair of roots differing by C<sub>3</sub> (*l-b-ḥ* and *l-b-z*) are illustrated in Table 11. Each root occurs in both forms, with all possible vowel choices. Each participant would only hear an individual root in one aspect, with one vowel choice for that aspect, pseudorandomised via Latin square. Each participant heard 50 stimuli in a randomised order, 20 of which were presented to them in the perfective (10 [a], 10 [i]) and 30 of which were presented in the imperfective (10 [a], 10 [i], 10 [u]).

The nonce words were checked with a native speaker to ensure that they were well-formed and not existing words. The words were recorded in the two frame sentences in (5) by a female native speaker. The two frame sentences contain adverbial phrases disambiguating the verb form that should occur (perfective in (5a), imperfective in (5b)). Both forms agree with a third person masculine singular subject: the imperfective has the indicative prefix [bi-] and the agreement prefix [ji-], while the corresponding perfective has no overt agreement marking.

- (5) *Frame sentences (for the nonce root l-b-ḥ <ل-ب-ح>, with perfective [a] and imperfective [i] as illustration)*
- a. *Perfective condition*  
ḡaalid labah                      ʔabli joom-een  
Khalid X.<sub>3SG.M.PFV</sub> since day-DUAL  
‘Khalid X-ed two days ago.’
- b. *Imperfective condition*  
ḡaalid bi-ji-lbiḥ                      kulli joom  
Khalid IND-3SG.M-X.IPFV every day  
‘Khalid X-es every day.’

Transitivity has been claimed to affect perfective vowel choice in Egyptian Arabic (Abdel-Massih *et al.* 1979). The lexicon compiled in this study also showed gradient effects of transitivity in Egyptian Arabic. Transitive verbs were more likely to have perfective [a] (58%) than intransitives (34%). Opposite trends were found with imperfective [a] (40% in transitives, 69% in intransitives). This was controlled for by keeping all frame sentences intransitive. Additionally, while semantic effects on vowel

choice have been reported for MSA (McCarthy 1994), Egyptian Arabic does not have similar patterns, and semantic effects are generally not expected to affect nonce-word experiment results.

**4.1.3. Procedure.** The experiment was conducted using PCIBex (Zehr & Schwarz 2018). The participants used their own devices and speakers/headphones. For each trial, participants were first presented with a nonce verb in either the perfective or the imperfective form in the corresponding frame sentence. The stimuli were presented auditorily and in Arabic orthography. Next, participants saw the frame sentence for the other aspect form and were prompted to read the sentence aloud using the appropriate verb form. The response was automatically recorded. Both the prompt and response sentences were shown with blank space in place of the nonce verb, so the only function that the orthography served was to help clarify the aspectual context.

Prior to the experiment, participants saw written instructions which explained that they would hear nonce verbs in one of the aspectual forms and that they would be asked to inflect them into the other form in their colloquial dialect. Then, two trials with real words (one for each aspectual form) were presented as examples to familiarise them with the task. To encourage participants not to use MSA, all instructions were presented in colloquial Egyptian.

(6) *Example of a trial in the experiment (translation on the left)*

Listen to the new word in the following sentence:

Khalid [perfective verb] two days ago.

Think about what the other form of the verb would be and prepare to say it in the following sentence:

Khalid [response] every day.

اسمع الكلمة الجديدة في الجملة الجاية:

خالد — قبل يومين.

فكر في ايه الصورة الثانية للفعل واجهز انكتقولها

في الجملة الجاية:

خالد — كل يوم.

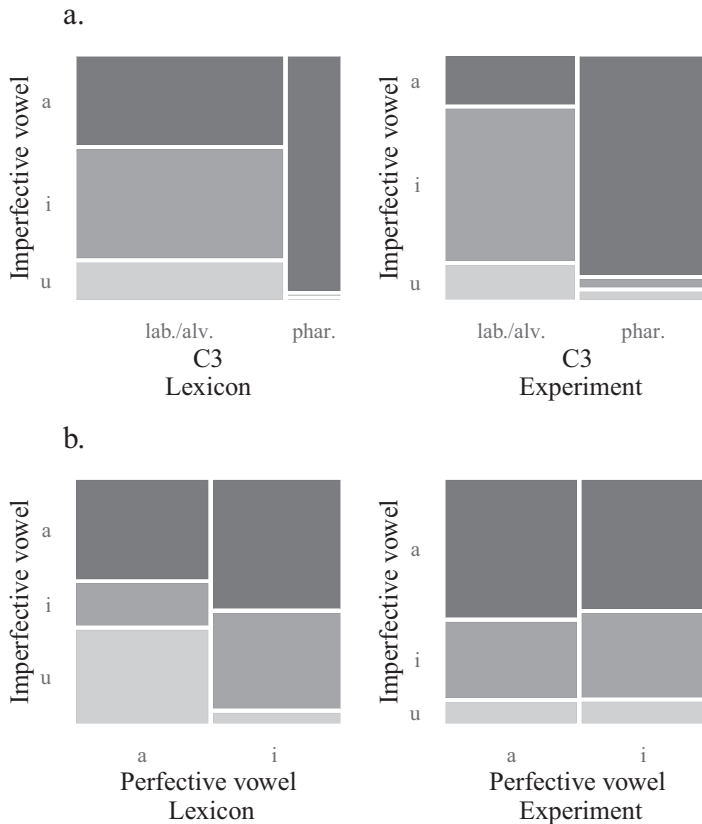
**4.1.4. Analysis.** After participants' responses had been recorded, the vowel choices in nonce verb forms were transcribed by the author, who is phonetically trained. The data were analysed using Bayesian mixed multinomial logit models in the *brms* package (Bürkner 2017) in R (R Core Team 2021). Two separate models were run for imperfective and perfective responses. The dependent variable was the vowel in the response, which had three levels ([a], [i] and [u]) in the imperfective model and two levels ([a] and [i]) in the perfective model. The models were run with [a] as the baseline, and so the imperfective results contain two pairwise comparisons, [i] vs. [a] and [u] vs. [a]. There were two fixed effects, vowel and consonant, with no interactions included. Random intercepts as well as random slopes for the fixed effects by subject and item were included. Bayesian models return a posterior probability for each effect estimate. The mean of the distribution is reported along with its 95% credible interval and the probability that the posterior distribution excludes zero.

## 4.2. Results

Results from 29 participants, with a total of 1,432 data points (18 excluded due to missing or truncated recordings), are presented below. The experimental results for imperfective responses as conditioned by root consonants ( $C_3$ ) and imperfective vowels are shown in Figure 4, next to the corresponding lexicon data. Note that the lexicon data presented here also show only the effect of  $C_3$ , and not any other consonant position. The effect of consonants in the experimental results is evident. Having a pharyngeal as  $C_3$  induces 90% [a] responses, which is almost as strong as the effect in the lexicon. On the other hand, [i] is favoured about 80% of the time in words containing only labials and plain alveolars, which is a slightly stronger preference compared to the lexicon. The model confirmed that the effect of consonant on vowel choice was credible ( $\beta = -4.79$ , 95% CrI [-5.74, -3.85],  $p(\beta < 0) = 1$ ).

As shown in Figure 4b, imperfective vowel responses did not differ visibly depending on the perfective vowel. This is confirmed by the lack of credible effect in the model ([a] vs. [i]:  $\beta = 0.30$ , 95% CrI [-0.35, 0.94],  $p(\beta > 0) = 0.83$ ; [a] vs. [u]:  $\beta = -0.39$ , 95% CrI [-1.56, 0.62],  $p(\beta < 0) = 0.76$ ).





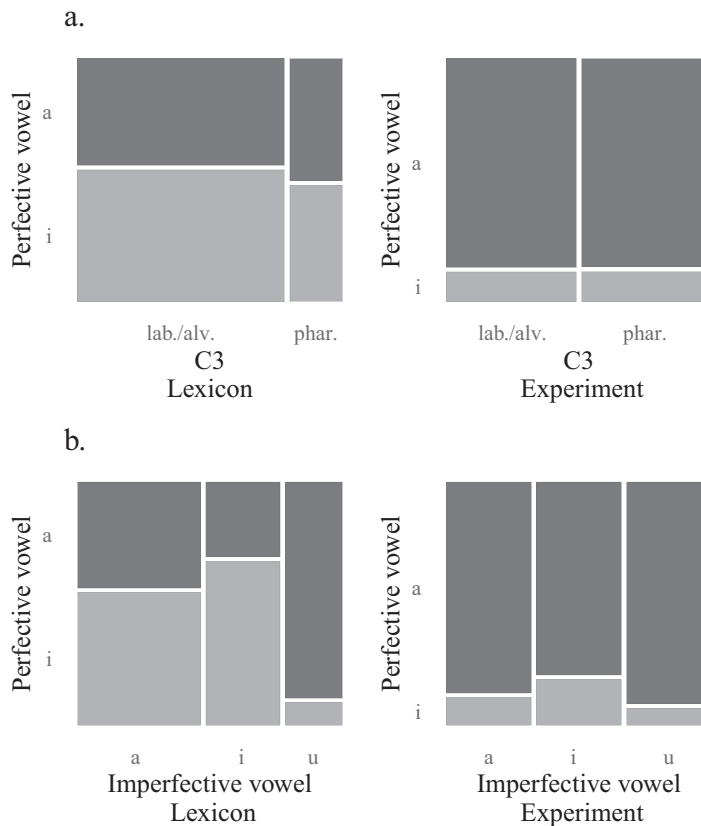
**Figure 4.** Imperfective vowel distribution by (a) root consonant  $C_3$  and (b) perfective vowel in the lexicon (left) compared to in nonce word responses (right).

This result contrasts with the lexicon data. Specifically, the preference in the lexicon for imperfective [u] with perfective [a] was not mirrored in the experimental results.

Turning to factors affecting vowel choice in the perfective, Figure 5 shows perfective vowel responses by consonant and imperfective vowel conditions in the experiment compared to the lexicon. One thing worth noting is that there is a much higher rate of [a] responses in the wug test compared to the lexicon. This preference is seen across the consonant conditions and across the imperfective [a] and [i] conditions. Consonants do not credibly affect perfective vowel choice ( $\beta = -0.44$ , 95% CrI [-1.32, 0.29],  $p(\beta < 0) = 0.86$ ), which is consistent with the lexicon study, where no effect was found either. In contrast, a credible effect of the imperfective vowel was found. Specifically, the lexicon preference for perfective [a] with imperfective [u] was generalised by participants ( $\beta = -0.96$ , 95% CrI [-1.95, -0.12],  $p(\beta < 0) = 0.99$ ). There is one discrepancy with the lexicon, in that the proportion of [a] responses for imperfective [u] is not as different compared to the proportion in other imperfective vowel conditions. Within the imperfective [u] condition, however, the proportion of [a] responses is comparable to that in the lexicon. Another, weaker lexicon trend, in which perfective [i] caused a small preference for imperfective [i], was also marginal in the experimental results ( $\beta = 0.50$ , 95% CrI [-0.21, 1.13],  $p(\beta > 0) = 0.93$ ).

#### 4.3. Summary

The experimental results show that Egyptian Arabic speakers generalise some predictive trends when deciding on vowel choices in novel verb forms, but crucially do not generalise all trends that can



**Figure 5.** Perfective vowel distribution by (a) root consonant  $C_3$  and (b) imperfective vowel in the lexicon (left) compared to in nonce word responses (right).

be found in the lexicon. Specifically, when predicting imperfective vowels, consonant effects were generalised, but perfective vowel effects were not, even though predictive trends of both types were found in the lexicon. When predicting perfective vowels, imperfective vowel effects in the lexicon were generalised. Root consonants did not help predict perfective vowels in the lexicon and were also not used by speakers.

While vowel-to-vowel correspondences in both paradigm directions were found in the lexicon, speakers only made use of imperfective-to-perfective correspondences. This suggests that speakers derive perfective forms based on imperfective forms and not the other way around; this is discussed in more detail in §5.

Notably, the asymmetry of consonant–vowel phonotactics was faithfully learned by speakers. Whereas consonants played no role in perfective vowel choice, pharyngeals strongly predicted [a] in the imperfective, and labials and plain alveolars strongly predicted [i]. Since the overlapping consonantal contexts for the alternating vowels in the perfective and imperfective forms mean that this pattern has no clear phonological explanation, the explanation must come from the morphological structure of the wazn I verb paradigm.

## 5. Analysis and discussion

In this section, I argue that the consonant–vowel co-occurrence restrictions in the Egyptian Arabic wazn I verb paradigm are a case of morphological domain-bounded phonotactics. These restrictions are active within a smaller morphological domain, where the imperfective verb forms reside. Specifically,

pharyngeals cause a preference for [a], and labials and plain alveolars cause a preference for [i]. These restrictions are not active in the perfective verb forms, which involve a larger morphological domain. This effect was found in lexicon statistics and was reliably generalised by speakers to nonce words.

The phonotactic restrictions on consonant–vowel co-occurrences exhibit an asymmetry that can only be explained by abstract morphological domains. If only the surface phonological forms were evaluated when predicting vowel choice, preference for less marked consonant–vowel sequences like [a] next to pharyngeal consonants would apply to both imperfective (CCVC) and perfective (CVCVC) forms, since the consonants surrounding the alternating vowels are the same. The effects of the final root consonant, which were tested in the wug test, illustrate this point. Nonce imperfective responses like [-lbaḥ] were strongly preferred to [-lbiḥ] or [-lbuḥ], but no preference was found for nonce perfective responses like [labah] (containing the same underlined CV sequence) relative to [libiḥ]. This pattern is thus comparable to other documented cases of morphological domain-bounded phonotactics involving concatenative morphology. For example, non-local sequences of plain stops followed by ejective stops are restricted within-morpheme in Bolivian Aymara, for example, \*[kajp'u]; such sequences are licit, however, if they occur across morphemes, for example, [paɬ-t'a-na] (Gallagher *et al.* 2019).


The pattern in the Egyptian Arabic verb paradigms are special, however, in that any morphological domains are difficult to observe on the surface due to non-concatenative morphology, as the imperfective and perfective forms are distinguished by vowel alternations (and change in prosody). Nonetheless, their existence in speakers' representations of this paradigm is supported by how speakers generalised predictive trends of vowel correspondence. First, a subset of such trends were generalised by speakers when inflecting nonce verbs, showing evidence that speakers do learn paradigmatic relations between perfectives and imperfectives, rather than treating them as independent. Additionally, even though the lexicon study shows that vowel correspondence trends were predictive in both directions, in the wug test, speakers only used imperfective vowels to predict the perfective and not *vice versa*. It is worth noting that both the consonant–vowel co-occurrence restrictions and the vowel correspondence trends are stochastic, and vowel choices in both perfectives and imperfectives are lexically idiosyncratic to a large extent. However, there is a rich literature showing that speakers will readily generalise stochastic patterns in their lexicon to nonce words (Zuraw 2000; Ernestus & Baayen 2003, *inter alia*), and failure to do so demands an explanation.

I discuss below two ways these morphological domains can be conceptualised. The first (§5.1) is morphosyntactic: the imperfective domain is embedded within the perfective domain, which includes additional structure. The second (§5.2) is in terms of paradigmatic correspondences: the imperfectives are the bases from which perfectives are derived by additional morphophonological processes. Under either approach, the imperfective domain can be viewed as simpler or smaller.

These two accounts differ greatly in their assumptions. The morphosyntactic account assumes both a direct role of the morphosyntactic structure in morphology and also non-concatenative morphemes (e.g., consonantal roots) as building blocks. On the other hand, the surface-based analysis assumes no decomposition into non-concatenative morphemes. The base–derived relationship of imperfective-to-perfective is also not taken for granted as a direct result of the morphosyntax. Rather, language learners face a larger hypothesis space of possible paradigm structures, which includes, for example, the opposite perfective-to-imperfective mapping and multidirectional mappings. Below, I will first show that, perhaps unsurprisingly, the morphosyntactic analysis successfully accounts for the lexicon and wug test data. I then explore whether the assumptions inherent to the morphosyntactic account are indeed necessary by presenting alternative accounts based on paradigmatic correspondences. Lastly, I compare the two analyses based on their implications for learning.

### 5.1. A morphosyntactic account

First, I show that a morphosyntactic analysis of the Egyptian Arabic verbal paradigm aligns well with the lexicon and wug-test results and provides an explanation for the asymmetrical effects of consonant–vowel co-occurrences. The verbal morphosyntactic structure in (4) is repeated in (7):

(7) *Morphosyntactic structure of Arabic wazn I verbs (repeated from (4))*

I follow previous morphosyntactic analyses of Arabic verbal structure such as Tucker (2011) in positing separate morphemic status for the consonantal root, the imperfective vowel and the perfective vowel. The imperfective vowel is treated as an Asp head, and the perfective vowel as a T(ense) head, consistent with the proposal that the imperfective is a tenseless, infinitival form, whereas the perfective is inflected for past tense (Benmamoun 1999). According to the structure in (7), the consonantal root first combines with the phonologically null functional head  $v$ , and then combines with the imperfective vowel to form the imperfective. Prosodic constraints have been shown to be successful at deriving the correct surface linearisation, with the consonantal and vocalic morphemes interleaved (Tucker 2011); I abstract away from them here.

The predictions of this analysis align perfectly with the wug test results. Since the perfective is derived from the imperfective and not *vice versa*, we expect that speakers should generalise lexicon trends based on the imperfective vowel when deriving perfective forms but do not expect generalisation in the opposite direction. These predictions are borne out.

Likewise, the absence of consonant–vowel interactions in the perfective form follows from independently proposed syntactic locality constraints, which disallow allomorph selection between any two elements that are separated by other overt material in the morphosyntactic structure (Embick 2010). Since the consonantal root merges with the imperfective vowel first, it is possible for phonological interactions between consonants and vowels to influence the choice of imperfective vowel. Further, since the perfective vowel is structurally adjacent to the imperfective vowel but not to the consonantal root, it follows that only imperfective vowel predictors contribute to predicting perfective vowel choice.

The vowel choice in both the imperfectives and the perfectives is understood as suppletive allomorphy in this analysis. Specifically, the consonant–vowel co-occurrence restrictions are phonological conditioning factors for allomorphic selection of the imperfective vowel, which is the Asp morpheme in (7). There are several different ways to implement such phonological conditioning factors. I briefly discuss two possible implementations here. One possibility is that such information is encoded as subcategorisation constraints, such that allomorphs of affixes select stems based on the stems' underlying phonological properties (Paster 2005). An illustration of one such subcategorisation rule is shown in (8a). This rule states that the Asp morpheme should be realised as [a] if it precedes a root morpheme that contains a pharyngeal. Another possibility is that selectional restrictions of allomorphs are learned as phonotactic grammars over the subsets of the lexicon which take the various allomorphs (Gouskova *et al.* 2015; Becker & Gouskova 2016). Phonotactic properties of the stems themselves and those of the affixed output forms can both play a role. Example rules are shown in (8b), following Gouskova *et al.* (2015). Specifically, the [a] allomorph of the Asp morpheme may be indexed to a sublexicon of consonantal roots that contain pharyngeal consonants, as in (8b-i), or to a sublexicon of imperfective forms that contain sequences of low vowel and pharyngeal, as in (8b-ii).

(8) *Example rules for imperfective vowel (Asp) allomorphy*a. *Subcategorisation*

$$\text{Asp} \leftrightarrow a / \_ \dots \{\text{h}, \text{ʔ}\}$$

b. *Sublexicon*

$$\text{i. Asp} \leftrightarrow a / \_ [+A] \mid \{\text{rg}\text{ʔ}_{+A}, \text{df}\text{ʔ}_{+A}, \text{sm}\text{h}_{+A} \dots\}$$

$$\text{ii. Asp} \leftrightarrow a / \_ [+A] \mid \{\text{rga}\text{ʔ}_{+A}, \text{dfa}\text{ʔ}_{+A}, \text{smah}_{+A} \dots\}$$

Similarly, the selection of perfective vowels is treated as suppletive allomorphy of the past tense morpheme. However, phonological conditioning of this morpheme's realisation can only refer to the identity of the imperfective vowel, which is local in the morphosyntactic structure, and not to the consonantal root.

There are two other factors that could potentially condition the perfective vowel allomorphy, though neither was supported by the data presented here. One is that [i] may be emerging as the default perfective vowel in Egyptian Arabic. This observation was made by my native speaker consultant based on the fact that a few words borrowed from MSA have undergone a change in the perfective vowel from [a] to [i] after becoming common in the colloquial variety. The second factor is transitivity: there is a slight preference in the lexicon for intransitive verbs to have [i] in the perfective (§4.1.2). However, in the wug test, [a] was predominantly the preferred choice for perfective vowel, even though the frame sentences force an intransitive interpretation of the verb.

As mentioned above, the phonological conditioning factors for both the imperfective and perfective vowel allomorphy are stochastic and exceptionful. As a result, allomorphy rules that refer to idiosyncratically indexed sets of lexical items are needed in addition to the phonologically based rules.

I illustrate how this analysis plays out with the following example. To form the imperfective form of 'return', [-rgaʕ], the consonantal root *r-g-ʕ* first selects [a] as the imperfective vowel, as opposed to [i] or [u], motivated by the avoidance of a marked phonotactic structure (namely, a high vowel next to a pharyngeal consonant). This happens below the AspP level.<sup>3</sup> The imperfective form can then be used as the input to form the perfective form [rigiʕ]. Because of locality restrictions, phonological information about the lowest embedded morpheme, the consonantal root, is no longer accessible. As the probabilistic correspondences between imperfective and perfective vowels found in this study do not include one that favours perfective [i] when the imperfective has [a], selection of [i] as the perfective vowel is based on a lexical indexation rule that is not phonologically conditioned.

## 5.2. *Surface correspondence accounts*

Having shown that an analysis which assumes both non-concatenative morphemes and morphosyntactic structure successfully account for how Egyptian Arabic speakers represent the wazn I verb paradigm, I now explore the extent to which these representations are necessary. To this end, I discuss how morphological theories based on surface correspondence within paradigms may explain the same data.

One idea that emerges from surface correspondence theories is that learners approach morphological paradigms by selecting one paradigm member as the base and using it to derive other members of the paradigm. Crucially, speakers choose a base that maximises informativity (Albright 2002), that is, one that can allow them to derive other forms most reliably. This kind of base choice can be independent of morphosyntax. For example, in Yiddish nouns, plurals have been shown to be more informative bases than singulars, even though singulars have simpler morphosyntactic structure, and patterns of paradigm levelling are consistent with the hypothesis that plurals are the base (Albright 2008).

There is, however, a mismatch between what would be the most informative base choice and what Egyptian Arabic speakers opted to use. §3.5 compared two regression models which instantiate the

<sup>3</sup>Note that the imperfective form is not realised on the surface as CCVC. Agreement affixes are expected to merge at higher levels of the structure, yielding forms such as [ji-rgaʕ] 'he returns'. Alternatively, if no overt agreement affix is available, as in the case of imperatives, vowel and glottal stop epenthesis will apply to yield surface forms like [ʔirgaʕ] 'return!'

two paradigmatic directions. The model that used the perfective as the base to predict the imperfective resulted in more improvement from chance, as shown by pseudo- $R^2$  metrics. In other words, there is more predictability to be gained from learning the trends that predict the imperfectives from the perfectives. This was the opposite of what speakers actually did in the wug test.

Additionally, vowel correspondence trends in the paradigm are redundant in that they contribute significantly to vowel predictability in both the perfective and the imperfective, and such redundancy in theory maximises predictability if speakers can use all of the trends (Bochner 1993). However, speakers selectively generalised trends in one paradigmatic direction (imperfective to perfective) but not the other.

We must therefore explain why Egyptian Arabic speakers represent the wazn I verb paradigm in a way that does not use paradigmatic relations most informatively. An analysis in which morphosyntax plays a direct role has been explored in §5.1. Another alternative explanation may be found in acquisition. Imperfectives may be preferentially chosen as the base because they are more common than perfectives in the early input and early production of Arabic-acquiring children (Aljenaie 2010). However, previous work on correspondences within paradigms has not made clear claims about how the acquisition timeline affects base selection, and more work on the early acquisition of Arabic verbs is necessary to address these questions.

A surface-based account where speakers select the imperfectives as the base to derive the perfectives has two implications that differ from the morphosyntax based analysis. First, an anonymous reviewer points out that analysing the imperfective as the derivational base has the consequence that the consonant–vowel co-occurrence restrictions active in the imperfective may be understood as phonotactic restrictions learned from stems (Anderson 1992). In contrast, under the morphosyntactic analysis, these restrictions are allomorphic selection criteria specific to the morphological context of the imperfective. Since this article focuses solely on the wazn I verb paradigm, I do not have the data to determine whether these consonant–vowel restrictions are morphologically specific or are part of more general phonotactic knowledge. Further work, such as lexicon studies on other morphologically simple words in Egyptian Arabic or patterns of loanword adaptation, would be informative on this issue.

Second, the morphosyntactic analysis makes a strong prediction about the asymmetrical presence of the consonant–vowel restrictions: given the simpler morphosyntactic structure of the imperfective, the absence of these restrictions in the perfective follows directly from syntactic locality constraints. This is different from approaches that allow different phonological grammars for different levels of morphological constructions or different lexical classes, such as Stratal OT (Bermúdez-Otero 1999; Kiparsky 2000), Cophonology (Anttila 1997; Inkelas 1998), and Cophonologies by Phase (Sande *et al.* 2020). Under these frameworks, this morphologically constrained phonotactic pattern would be modelled by demoting constraints against marked consonant–vowel sequences when computing the perfective form as opposed to the imperfective form. However, there is no *a priori* reason that the markedness constraints should be demoted rather than promoted when evaluating the perfective. Existing work on other Arabic dialects generally reports that the consonant–vowel co-occurrence restrictions are found solely in the imperfectives, as in Egyptian Arabic, with the exception of Baghdadi Arabic (Blanc 1964). Comparative work and detailed investigation on dialects like Baghdadi is needed to decide between these approaches.

### 5.3. Discussion

While morphological domain effects on phonotactic patterns are a common phenomenon typologically, they have been argued to be challenging to learn. Specifically, without knowledge of morphological domains, licit occurrences of a phonotactic sequence across morphological boundaries would become noise, making it more difficult for a learner to correctly reject these sequences within the relevant domain. Precisely for this reason, Gallagher *et al.* (2019) show that a computational phonotactic learner performs better when trained on data with morphological segmentation than when trained on data



without this information. Perhaps paradoxically, however, such patterns have also been argued to aid in early morphological learning, since the sequences that violate phonotactic restrictions occur exclusively at morphological boundaries and thus can serve as cues for discovering these boundaries (Trubetzkoy 1939). The usefulness of these cues in segmentation has been demonstrated using computational models employing some version of a heuristic whereby word or morpheme boundaries are inserted to break up low-probability sequences. They have been shown to outperform models that only employ other distributional cues for segmentation, such as transitional probability (Daland 2009; Adriaans & Kager 2010, etc.). Morphological domain effects on phonotactics thus provide a good testing ground for the relationship between phonotactic and morphological learning.

Previous discussions of such patterns have focused on concatenative morphology. A crucial feature of non-concatenative morphology is that morphological domains are difficult to distinguish at the surface level. This article was only able to find evidence supporting the complicated morphological structure through a detailed study of the predictors for the vowel alternation pattern in a lexicon study and a behavioural experiment. Furthermore, there are contending accounts of the nature of the morphological knowledge needed to describe this pattern. The data presented in this article are compatible with two distinct analyses that differ in the representations they assume and offer distinct accounts of the nature of the morphological domain effect. In turn, they have different implications for learning.

The first analysis presented assumes the representation of abstract morphosyntactic structure. One potential acquisition trajectory for this account is that children have to be able to associate verb forms with their underlying morphosyntactic structure before learning that this structure plays a role in phonotactics. Prior to having the morphosyntactic knowledge, discovery of phonotactic patterns is harder because cross-morphemic occurrences of the restricted sequences introduce noise. This mirrors the problem proposed for patterns in concatenative morphology. On the other hand, morphologically constrained phonotactics could potentially also aid in the discovery of morphosyntactic structure. Locality constraints may be crucial, in that the absence of phonological effects on allomorphy serves as a cue for intervening morphosyntactic structure. However, since morphological domain effects on phonotactics are all restricted to specific allomorphic conditioning factors in such accounts, it is still likely that they are learned late because of the relatively small amount of data available.

The second assumption of the morphosyntactic account is non-concatenative morphemes, specifically the consonantal root and vocalic morphemes (McCarthy 1979). In particular, the distinct functions that these morphemes have been argued to carry lend themselves well to highly compositional frameworks like Distributed Morphology (e.g., Arad 2005). The involvement of the consonantal root in forming the imperfective form and its inaccessibility – due to syntactic locality constraints – in subsequent morphological processes are crucial for this analysis to account for the absence of consonant effects in predicting the perfective vowel in Egyptian Arabic. This article, however, has not explored any alternative account which dissociates the role of non-concatenative morphemes and that of morphosyntactic structure.

The account based on paradigmatic correspondences requires neither of the two above assumptions. However, the challenge for this account is to explain why learners choose imperfective-to-perfective over other possible paradigm organisations. As discussed above, the criterion of maximising informativity prefers the perfective as the base instead, so some alternative explanation is called for. While the acquisition literature on Arabic suggests that the imperfective is learned earlier (Aljenaie 2010), the extent to which acquisition order determines base selection has not been explored extensively in the literature.

## 6. Conclusion

This article presents a case study on a vowel alternation pattern in Egyptian Arabic which demonstrates that morphological domains may condition the application of phonotactic restrictions, even in cases where non-concatenative morphological processes obscure the domains on the surface. Specifically,



phonotactic restrictions on consonant–vowel co-occurrences are found in one verb form but not the other; this is because the phonotactic restrictions are confined within some morphological domain. Such cases of morphological domain effects on phonotactics that involve non-concatenative morphology may require different representations, and they raise novel yet comparable learning challenges compared to cases involving concatenative morphology. Continued investigations on similar patterns using lexicon studies and behavioural experiments, as well as acquisition studies and learning models, will shed further light on the interplay of phonotactics and morphology.

**Data availability statement.** The data that support the findings of this study and all supplemental material are openly available on OSF at [https://osf.io/tdr6x/?view\\_only=4aec9be4f2204849ac81932103fe3dd9](https://osf.io/tdr6x/?view_only=4aec9be4f2204849ac81932103fe3dd9).

**Acknowledgements.** I would like to express my sincere gratitude to Claire Moore-Cantwell, Bruce Hayes, Hilda Koopman, Megha Sundara, Canaan Breiss, Stefan Keine, Elizabeth Sola-Llonch, the audience at UCLA Phonology Seminar, AMP 2021, ASAL 35 and IATL 37, and four anonymous reviewers for their valuable comments. I'm extremely grateful to my consultant, Fatema Shokr, and to Mary Bishara, Abeer Abbas and Brady Ryan for help in creating the experiment materials and recruiting. I would also like to thank Michelle Chan, Paige Escobar, Serena Lee, Jennifer Miyaki and Devon Whalen for help in data processing. All errors are my own.

**Competing interests.** The author declares no competing interests.

## References

- Abdel-Massih, Ernest T., Zaki N. Abdel-Malek & El-Said M. Badawi (1979). *A comprehensive study of Egyptian Arabic*, volume 3. Ann Arbor, MI: MPublishing.
- Adriaans, Frans & René Kager (2010). Adding generalization to statistical learning: the induction of phonotactics from continuous speech. *Journal of Memory and Language* 62, 311–331.
- Ahyad, Honaida Yousuf (2019). *Vowel distribution in the Hijazi Arabic root*. PhD dissertation, Stony Brook University.
- Ahyad, Honaida Yousuf & Michael Becker (2020). Vowel unpredictability in Hijazi Arabic monosyllabic verbs. *Glossa* 5, Article no. 32 (18 pp.).
- Albright, Adam (2002). *The identification of bases in morphological paradigms*. PhD dissertation, University of California, Los Angeles.
- Albright, Adam (2005). The morphological basis of paradigm leveling. In Laura J. Downing, T. Alan Hall & Renate Raffelsiefen (eds.) *Paradigms in phonological theory*. Oxford: Oxford University Press, 17–43.
- Albright, Adam (2008). Inflectional paradigms have bases too: arguments from Yiddish. In Asaf Bachrach & Andrew Nevins (eds.) *Inflectional identity*. Oxford: Oxford University Press, 271–312.
- Aljenaie, Khawla (2010). Verbal inflection in the acquisition of Kuwaiti Arabic. *Journal of Child Language* 37, 841–863.
- Anderson, Stephen R. (1992). *A-morphous morphology*. Cambridge: Cambridge University Press.
- Anttila, Arto (1997). Deriving variation from grammar. In Frans Hinskens, Roeland Hout & W. Leo Wetzels (eds.) *Variation, change and phonological theory*. Amsterdam: Benjamins, 35–68.
- Arad, Maya (2005). *Roots and patterns: Hebrew morpho-syntax*. Dordrecht: Springer.
- Becker, Michael & Maria Gouskova (2016). Source-oriented generalizations as grammar inference in Russian vowel deletion. *LI* 47, 391–425.
- Benmamoun, Elabbas (1999). Arabic morphology: the central role of the imperfective. *Lingua* 102, 175–201.
- Bennett, Ryan (2018). Recursive prosodic words in Kaqchikel (Mayan). *Glossa* 3, Article no. 67 (33 pp.).
- Bermúdez-Otero, Ricardo (1999). *Constraint interaction in language change: quantity in English and Germanic*. PhD dissertation, University of Manchester and Universidad de Santiago de Compostela.
- Blanc, Haim (1964). *Communal dialects in Baghdad*. Cambridge, MA: Harvard University Press.
- Bochner, Harry (1993). *Simplicity in generative morphology*. Berlin: De Gruyter Mouton.
- Bürkner, Paul-Christian (2017). brms: an R package for Bayesian multilevel models using Stan. *Journal of Statistical Software* 80, 1–28.
- Daland, Robert (2009). *Word segmentation, word recognition, and word learning: a computational model of first language acquisition*. PhD dissertation, Northwestern University.
- Deal, Amy Rose & Matthew Wolf (2017). Outwards-sensitive phonologically-conditioned allomorphy in Nez Perce. In Vera Gribanova & Stephanie S. Shih (eds.) *The morphosyntax–phonology connection: locality and directionality at the interface*. Oxford: Oxford University Press, 29–60.
- Eid, Mushira (2007). Arabic on the media: hybridity and styles. In Jonathan Owens (ed.) *Approaches to Arabic linguistics*. Leiden: Brill, 403–434.
- Embick, David (2010). *Localism versus globalism in morphology and phonology*. Cambridge, MA: MIT Press.
- Ernestus, Mirjam & R. Harald Baayen (2003). Predicting the unpredictable: interpreting neutralized segments in Dutch. *Lg* 79, 5–38.

- Ferguson, Charles A. (1959). Diglossia. *Word* **15**, 325–340.
- Gallagher, Gillian, Maria Gouskova & Gladys Camacho Rios (2019). Phonotactic restrictions and morphology in Aymara. *Glossa* **4**, Article no. 29 (38 pp.).
- Ghazeli, Salem (1977). *Back consonants and backing coarticulation in Arabic*. PhD dissertation, University of Texas at Austin.
- Gouskova, Maria (2018). Morphology and phonotactics. In Mark Aronoff (ed.) *Oxford research encyclopedia of linguistics*. Oxford: Oxford University Press.
- Gouskova, Maria, Luiza Newlin-Lukowicz & Sofya Kasyanenko (2015). Selectional restrictions as phonotactics over sublexicons. *Lingua* **167**, 41–81.
- Green, Mike (2007). Lisan masry Egyptian Arabic dictionary. Published online at <https://www.lisaanmasry.org/>.
- Guerssel, Mohand & Jean Lowenstamm (1996). Ablaut in Classical Arabic measure I active verbal forms. In Jacqueline Lecarme, Jean Lowenstamm & Ur Shlonsky (eds.) *Studies in Afroasiatic grammar*. The Hague: Holland Academic Graphics, 123–134.
- Halle, Morris & Alec Marantz (1993). Distributed Morphology and the pieces of inflection. In Kenneth Hale & Samuel Jay Keyser (eds.) *The view from Building 20: essays in linguistics in honor of Sylvain Bromberger*. Cambridge, MA: MIT Press, 111–176.
- Hayes, Bruce & Zsuzsa Czirák Londe (2006). Stochastic phonological knowledge: the case of Hungarian vowel harmony. *Phonology* **23**, 59–104.
- Herzallah, Rukayyah S. (1990). *Aspects of Palestinian Arabic phonology: a nonlinear approach*. PhD dissertation, Cornell University.
- Inkelas, Sharon (1998). The theoretical status of morphologically conditioned phonology: a case study from dominance. In Geert Booij & Jaap Marle (eds.) *Yearbook of morphology 1997*. Amsterdam: Springer, 121–155.
- Jun, Jung-ho & Adam Albright (2017). Speakers' knowledge of alternations is asymmetrical: evidence from Seoul Korean verb paradigms. *JL* **53**, 567–611.
- Kastner, Itamar (2016). *Form and meaning in the Hebrew verb*. PhD dissertation, New York University.
- Kastner, Itamar (2019). Templatic morphology as an emergent property: roots and functional heads in Hebrew. *NLLT* **37**, 571–619.
- Kiparsky, Paul (2000). Opacity and cyclicity. *The Linguistic Review* **17**, 351–367.
- Kuo, Jennifer (2020). *Evidence for base-driven alternation in Tgdaya Seediq*. Master's thesis, University of California, Los Angeles.
- Lauffer, Asher & Thomas Baer (1988). The emphatic and pharyngeal sounds in Hebrew and in Arabic. *Language and Speech* **31**, 181–205.
- Lehn, Walter (1963). Emphasis in Cairo Arabic. *Lg* **39**, 29–39.
- Marantz, Alec (2013). Locality domains for contextual allomorphy across the interfaces. In Alec Marantz & Ora Matushansky (eds.) *Distributed Morphology today: morphemes for Morris Halle*. Cambridge, MA: MIT Press, 95–115.
- McCarthy, John J. (1979). *Formal problems in Semitic phonology and morphology*. PhD dissertation, Massachusetts Institute of Technology.
- McCarthy, John J. (1991). Semitic gutturals and distinctive feature theory. *University of Massachusetts Occasional Papers in Linguistics* **14**, 29–50.
- McCarthy, John J. (1993). Template form in prosodic morphology. In Laurel Smith Stvan (ed.) *FLSM III: papers from the third annual meeting of the Formal Linguistics Society of Midamerica*. Bloomington, IN: Indiana University Linguistics Club, 187–218.
- McCarthy, John J. (1994). The phonetics and phonology of Semitic pharyngeals. In Patricia Keating (ed.) *Papers in Laboratory Phonology III: phonological structure and phonetic form*. Cambridge: Cambridge University Press, 191–233.
- McOmber, Michael L. (1995). Morpheme edges and Arabic infixation. In Mushira Eid (ed.) *Perspectives on Arabic linguistics VII: papers from the Seventh Annual Symposium on Arabic Linguistics*. Amsterdam: Benjamins, 173–188.
- Norlin, Kjell (1987). *A phonetic study of emphasis and vowels in Egyptian Arabic*. PhD dissertation, Lund University.
- Omar, Margaret K. (1973). *The acquisition of Egyptian Arabic as a native language*. Berlin: De Gruyter Mouton.
- Paparonas, Lefteris (2021). Default by intervention: allomorphy and locality in the Modern Greek verb. *Proceedings of the Linguistic Society of America* **6**, 499–513.
- Paster, Mary (2005). Subcategorization vs. output optimization in syllable-counting allomorphy. *WCCFL* **24**, 326–333.
- R Core Team (2021). *R: a language and environment for statistical computing*. Vienna: R Foundation for Statistical Computing. Published online at <https://www.R-project.org/>.
- Sande, Hannah, Peter Jenks & Sharon Inkelas (2020). Cophonologies by ph(r)ase. *NLLT* **38**, 1211–1261.
- Suomi, Kari, James M. McQueen & Anne Cutler (1997). Vowel harmony and speech segmentation in Finnish. *Journal of Memory and Language* **36**, 422–444.
- Toosarvandani, Maziar (2016). Vocabulary insertion and locality: verb suppletion in Northern Paiute. *NELS* **46**, 247–257.
- Trubetzkoy, Nikolai S. (1939). Grundzüge der Phonologie. *Travaux du Cercle linguistique de Prague* **7**.
- Tucker, Matthew A. (2011). The morphosyntax of the Arabic verb: toward a unified syntax-prosody. In Nicholas LaCara, Anie Thompson & Matthew A. Tucker (eds.) *Morphology at Santa Cruz: papers in honor of Jorge Hankamer*. Santa Cruz, CA: SlugPubs, 177–211.
- Venables, William & Brian Ripley (2002). *Modern applied statistics with S*. 4th edition. New York: Springer.

- Watson, Janet C. E. (2002). *The phonology and morphology of Arabic*. Oxford: Oxford University Press.
- Younes, Munther (1994). On emphasis and /r/ in Arabic. In Mushira Eid, Vicente Cantarino & Keith Walters (eds.) *Perspectives on Arabic linguistics VI: papers from the sixth Annual Symposium on Arabic*. Amsterdam: Benjamins, 215–233.
- Youssef, Islam (2019). The phonology and micro-typology of Arabic R. *Glossa* 4, Article no. 131 (36 pp.).
- Zehr, Jeremy & Florian Schwarz (2018). PennController for IBEX. Published online at <https://doc.pcibex.net>.
- Zuraw, Kie (2000). *Patterned exceptions in phonology*. PhD dissertation, University of California, Los Angeles.