First-order definable phonological structure DRAFT

Adam Jardine

October 20, 2017

Abstract

This paper posits the hypothesis that phonological structure is first-order definable from surface strings, and shows how this hypothesis meaningfully restricts the typology with respect to surface correspondence-based analyses of long-distance consonant agreement. This is accomplished through the technique of logical transductions, which make it possible to compare the relative expressivity of constraints operating over distinct structures, and thus to connect phonological structure with computational characterizations of phonological patterns. Comparisons are also made to previous explanations of long-distance consonant agreement under formal language theory and surface correspondence in Optimality Theory.

Keywords: computational phonology, logic, surface correspondence, long-distance consonant agreement

1 Introduction

Computational analyses of phonological patterns have resulted in restrictive characterizations of phonology in terms of string patterns, but a yet unanswered question is what role phonological structure plays in this complexity. This paper offers an answer to this question, both by introducing a method by which we can compare the relative complexity of distinct structures, and by positing an upper bound on this complexity. Specifically, the following motivates and explores the hypothesis that *phonological structure is first-order (FO-)definable from strings*.

As a test case, this paper examines surface correspondence (Hansson, 2001, 2010; Rose and Walker, 2004; Bennett, 2013) as a mechanism for explaining *long-distance consonant agreement* (LDCA). An example of LDCA is in Navajo, in which all sibilants in a word must agree in anteriority (relevant sounds are highlighted in bold).

(1) Navajo sibilant harmony (Sapir and Hoijer, 1967)

a. **ʃ**í-t**ʃ**ʼíːh 'my nose' ***s**í-t**ʃ**ʼíːh b. $\widehat{\text{ʃì-tʃ}}$ʼíːh

c. **s**i-**z**id 'my scar' ***ʃ**i-**z**id d. $\widehat{\text{si-z}}$id

LDCA, as in Navajo, holds over intermediate material without perceptibly affecting it, a fact that has been analyzed through a correspondence relation that directly relates the interacting segments, as depicted visually in (1b) and (d). This paper shows that the correspondence relations that capture attested LDCA patterns are definable in a fragment of FO, and also that the FO-definability hypothesis excludes logically possible, unattested correspondence relations, including at least one that can be defined in Optimality Theory (Prince and Smolensky, 1993, 2004).

FO-definitions of phonological structure are made through the technique of logical transductions (Courcelle, 1994; Engelfriet and Hoogeboom, 2001; Courcelle et al., 2012). Briefly, logical transductions define an output structure in the logical language of the input structure. This paper introduces this technique and shows how, in terms of a *theory* of phonological structure, FO-definability constrains the expressivity of representations. More specifically, (FO-definable) constraints written over FO-definable structure are not any more expressive than FO-definable constraints written over strings. This is significant because FO-definable sets of strings appear to be the upper bound on the complexity of phonology (Graf, 2010a,b; Rogers et al., 2013). Concretely, it is shown that agreement constraints that enforce assimilation are FO-definable, and thus do not go beyond this complexity threshold.

This approach appears different from that of OT, which builds structure using violable constraints, as well as from approaches based on computational characterizations of string patterns, which use a minimal amount of structure. However it is shown below how a logical approach to structure can incorporate insights from both. Future work can articulate a more restrictive theory based on both the substantive constraints that are the hallmarks of OT analyses and the computational analyses of formal language-theoretic analyses of phonology.

This paper is structured as follows. §2 reviews examples of LDCA and language-specific conditions on correspondence. §3 motivates and introduces FO definitions of phonological structure, and §4 demonstrates that correspondence relations for attested LDCA patterns are FO definable, and how the proposal excludes some logically possible but unattested correspondence relations. §5 compares this approach to OT and existing formal language-theoretic approaches to phonology, §6 discusses paths for future work, and §7 concludes.

2 The empirical focus

This section reviews important typological generalizations about the primary empirical domain of interest, long-distance consonant agreement (LDCA), and how they motivate a surface correspondence relation that directly connects agreeing consonants (Hansson, 2001, 2010; Rose and Walker, 2004; Bennett, 2013).

Importantly, the following views the relevant patterns in terms of *phonotactics*, or surface well-formedness. This is not how these patterns are always studied; LDCA is often viewed as a *process* in which the featural makeup of segments in an input underlying form is changed to produce an output surface form, and similarly stress patterns are often viewed as the application of some stress-assigning process to an input string of unstressed syllables. However, there are several reasons why we can and should look at these patterns, and their concomitant structure, in terms of surface well-formedness. First, it is considerably simpler enterprise—we can analyze the surface structure without making additional assumptions about how it got there. Second, phonotactic patterns can be viewed as *sets* (Heinz, 2010), and the computational properties of sets that are relevant to phonology are better understood than those for processes, which are *maps* (see Chandlee (2014) for discussion). Third, and perhaps most importantly, anything we learn about phonotactics is a generalization that can and should form a constraint on our theory of processes. We can do this directly by studying logical transformations, a point discussed in more detail in §6. In sum, well-formedness is a natural starting point for the hypothesis advanced in this paper, but it is by no means an end goal.

2.1 Long-distance consonant agreement and surface correspondence

LDCA can be characterized as agreement between two or more consonants separated by at least one vowel without the intervening segments being affected (Hansson, 2001, 2010; Rose and Walker, 2004). For example, in Navajo, all sibilants in a word assimilate in anteriority to the rightmost word (Sapir and Hoijer, 1967). The relevant segments are highlighted in bold.

(2) Navajo sibilant harmony (Sapir and Hoijer, 1967)

a.	**ʃ**i-lí̃ːʔ	‘my horse’	d.	**s**ì-ʔá	‘a round object lies’
b.	**ʃ**í-**tʃ**í̃ːh	‘my nose’	e.	**ʃ**ì-teː**ʒ**	‘they (dual) are lying’
c.	**s**i-**z**id	‘my scar’	f.	da**s**doːli**s**	‘he (4th) has his foot raised’

In (2a) through (2e), [s] and [ʃ] alternate depending on the anteriority of a sibilant in the root, if present. The examples (2e) through (f) show that this dependency between sibilants holds across arbitrary distances; words like (2f) [dasdolis] 'he (4th) has his foot raised' are attested but hypothetical words like *[daʃdolis] or *[dasdoliʃ] are ungrammatical, even though the disagreeing sibilants are in nonadjacent syllables and are separated by the five-segment sequence [doli].

Aside from the long-distance interaction, the Navajo pattern has several characteristics typical of LDCA patterns (Hansson, 2001, 2010). First is that agreement affects *similar* segments—all sibilants in a word are affected. Second is that intervening material is not phonologically affected (Hansson, 2001, 2010; Rose and Walker, 2004). From these properties, Walker (2000, 2001), Hansson (2001, 2010), Rose and Walker (2004), and subsequent researchers concluded that the interaction between consonants in LDCA is based on a direct *correspondence* relation (McCarthy and Prince, 1995, 1999) between the agreeing segments, and is not mediated through changes to the intervening segments. On this view, in Navajo all sibilants correspond, and the agreement generalization is thus that corresponding segments must agree in anteriority. This is depicted in (3), with the correspondence relation represented visually as a curved line between segments. (Correspondence is usually indicated through indices, but representing it visually will be useful when considering explicit models of phonological structure in §3.)

(3) a. **ʃ**ì-teː**ʒ** b. ***s**ì-teː**ʒ** c. *[+ant] [–ant]

In (3a) and (3b), a correspondence relation connects the sibilants [ʃ] and [ʒ] and [s] and [ʒ], respectively. Why (3b) is ill-formed can then be stated concisely as depicted in (3c): a configuration in which a [+anterior] segment and a [–anterior] segment are in correspondence is forbidden. (Note that while the [+anterior] feature is written before the [–anterior] one, they are unordered.) Such generalizations are usually formalized as IDENT-CC faithfulness constraints between correspondents, while (3c) is written like a markedness constraint, but they are equivalent in that they identify forbidden configurations of correspondents. Thus, the fact that sibilants agree in Navajo without any perceptual effect on the intervening segments is captured by this direct correspondence relation between sibilants.

The surface correspondence relation has been defined differently by different authors. This paper will assume Bennett (2013)'s definition, as it is the most explicit. Bennett (2013) argues that surface correspondence is an *equivalence relation*, which means that it is reflexive (i.e., that every segment is in correspondence with itself), symmetric (i.e., unordered), and transitive. Not all other theories of surface correspondence agree on these points. For example, the definition of correspondence in Walker (2001) is not symmetric. As discussed at the relevant points below, changing this assumption does not bear on the FO-definability of correspondence.

2.2 Language-specific conditions on correspondence

More importantly, whereas the properties of an equivalence relation are assumed to be *universal* properties of correspondence, any theory of correspondence must admit *language-specific* conditions on correspondence. One such condition, the featural similarity of the corresponding segments, has already been discussed. The following reviews several LDCA patterns that, assuming that LDCA is the result of a correspondence relation, show that surface correspondence can also be conditioned on distance, intervening segments, and on the prosodic role of the correspondents.

As such, the following focuses on the correspondence relation itself and abstracts away from faithfulness considerations that factor into OT theories of correspondence. Thus, the following shall not discuss featural identity, root identity, and directionality to the extent that they affect the surface realization of segments and not the correspondence relation itself (Rose and Walker, 2004; Hansson, 2001, 2010). While this may seem unusual, it will allow us to directly analyze the structural properties of surface correspondence in §4, and compare and contrast this to how surface correspondence is generated in OT in §5.1. Finally, as correspondence is (usually) assumed not to be directly observable, there can be a tremendous amount of ambiguity as to when correspondence is actually present. The following focuses on situations in which correspondence is directly related to agreement: that is, a correspondence relation is present between two distinct segments if and only if there is LDCA between them. This is a reasonable assumption for focusing on the core cases of LDCA; an OT analysis that does not assume this will be discussed in §5.1.

First, it is not always the case that LDCA is unbounded (Hansson, 2010; McMullin and Hansson, 2016). In Koyra (Koorete; Hayward, 1982), only sibilants separated by at most one vowel harmonize in anteriority. In (4a) below, the sibilant in the affixes /-us-/ 'causative', /-osːo-/ '3 masc. sing. perfective', and /-esːe-/ is realized as [–anterior] after affixes ending in a [–anterior] sibilant. However, as shown in (4b), the affix sibilants are not affected by a [–anterior] sibilant occurring further to the left.

(4) Koyra (Hayward, 1982)

a.

/goː**tʃ**-u**s**-/	goː**tʃ**-u**ʃ**	'cause to pull'
/ʔor**dʒ**-u**s**-/	ʔor**dʒ**-u**ʃ**	'to make big, increase (tr.)'
/**dʒ**a**ʃ**-u**s**-e**s**ːe/	**dʒ**a**ʃ**-u**ʃ**-e**ʃ**ːe	'let him/them frighten someone'

b.

/**ʃ**od-u**s**-/	**ʃ**od-u**s**	'cause to uproot'	***ʃ**od-u**ʃ**
/**tʃ**aːn-u**s**-/	**tʃ**aːn-u**s**-	'cause to load'	***tʃ**aːn-u**ʃ**-
/**ʃ**od-d-o**s**ːo-/	**ʃ**odː-o**s**ːo	'he uprooted'	***ʃ**odː-o**ʃ**ːo

Importantly, in a sequence of syllable-adjacent sibilants, all harmonize together. In (4a) /**dʒ**a**ʃ**-u**s**-e**s**ːe/ 'let him/them frighten someone', the assimilation of the /s/ in /-us-/ to the preceding /ʃ/ in /dʒaʃ-/ feeds assimilation in the following /sː/ in /-esːe/; thus all sibilants harmonize, producing [**dʒ**a**ʃ**-u**ʃ**-e**ʃː**e], not *[**dʒ**a**ʃ**-u**ʃ**-e**sː**e].

Assuming a correspondence relation is responsible for this LDCA pattern, the well-formed and ill-formed correspondence structures in Koyra are as in (5).

(5) a. goː**tʃ**-u**ʃ** d. *goː**tʃ**-u**s** 'cause to pull' (4a)

b. **dʒ**a**ʃ**-u**ʃ**-e**ʃː**e e. ***dʒ**a**ʃ**-u**s**-e**sː**e 'let him/them frighten someone' (4a)

c. **ʃ**od-ud**s** f. ***ʃ**od-u**ʃ** 'cause to uproot' (4b)

Well-formed correspondence structures in Koyra are as in (5a) through (c), contrasted with the ill-formed structures in (5d) through (5f). Syllable-adjacent sibilants must correspond (5a) (c.f. (5d)); more generally, due to transitivity, all sibilants in a string of syllable-adjacent sibilants must correspond (5b) (c.f. (5d)). However, any sibilants that are separated by one or more non-sibilant consonants cannot correspond ((5c) and (5f)). Note that, as in the ill-formed (5d) versus the well-formed (5c), when the correspondence relation is absent is just as important to the well-formedness generalization as when it is present.

Koyra is not the only case in which intervening material can affect LDCA. In colloquial Slovenian (henceforth simply 'Slovenian'; Jurgec, 2011) coronal obstruents [t,d] block [±anterior] harmony but coronal sonorants [n,ɾ] do not. (As in Kinyarwanda, Slovenian sibilant harmony is optional except in the blocking cases, in which it is entirely blocked.)

(6) Slovenian sibilant harmony (Jurgec, 2011)

a. **s**pi 'sleeps' **ʃ**pi-**ʃ** '(you) sleep'
zapɔɾ 'prison' **ʒ**apɔɾ-ni**ʃ**ki 'prison (adj)'

b. **s**i**t** 'full' na-**s**i**t**-i**ʃ** '(you) feed' *na-**ʃ**i**t**-i**ʃ**
zi**d**a 'build' **z**i**d**a-**ʃ** '(you) build' ***ʒ**i**d**a-**ʃ**

The condition on correspondence in these languages is thus that it is blocked when some class of segments intervenes. In Slovenian, the well- and ill-formed correspondence relations can be posited as follows.

(7) a. po**ʒ**abi-**ʃ** c. * po**z**abi-**ʃ** '(you) forget' (6a)

b. **z**i**d**a-**ʃ** d. * **ʒ**i**d**a-**ʃ** '(you) build' (6b)

To account for LDCA in Slovenian with surface correspondence, the relation has to hold between two sibilants when no non-sibilant coronal obstruents intervene, as in (7a) and (c). Conversely, it cannot hold when coronal stops intervene (7b) and (d). Similar patterns are also attested in Kinyarwanda (Walker and Mpiranya, 2006; Walker et al., 2008; Bennett, 2013) and several dialects of Berber (Elmedlaoui, 1995; Hansson, 2010).[1]

Finally, nasal harmony in Kikongo (Ao, 1991; Odden, 1994) shows that syllable structure can also factor into conditions on correspondence (Rose and Walker, 2004). In Kikongo, a suffix /l/ becomes [n] following a stem containing an onset nasal. The following illustrates this with the applicative suffix /il/.

(8) Kikongo (Odden, 1994; Rose and Walker, 2004)

a.	ku-toot-il-a	'to harvest for'	b.	ku-ki**n**-i**n**-a	'to dance for'
	sakid-il-a	'congratulate for'		ku-ki**n**-is-i**n**-a	'to make dance for'

Rose and Walker (2004) highlight that coda nasals do not trigger the harmony (9a), but they are transparent to it (9b). The following illustrates this with the active perfective suffix /ele/. (Realization of the suffix vowels depends on vowel harmony, and an unrelated process changes /l/ to [d] before [i].)

(9) Neutrality of Kikongo NCs (Rose and Walker, 2004)

a.	ba**nt**ik-i**d**i	'begun'	*ba**nt**ik-i**n**i	b.	tu-**mant**-i**n**i	'we climbed'
	ke**mb**-e**l**e	'we hunted'	*ke**mb**-e**n**e		tu-**meŋg**-i**n**i	'we hated'

In terms of correspondence, while onset nasals form a relation with sonorant consonants and voiced stops, coda nasals must be transparent to the correspondence relation. This is diagrammed explicitly in (10). In (10), periods mark syllable boundaries to make it clear when a consonant is in coda or onset position.

(10) a. ku.ki.**n**i.**n**a c. *ku.ki.**n**i.**l**a 'to dance for' (8b)

b. ba**n**.ti.ki.**d**i d. *ba**n**.ti.ki.**n**i 'begun' (9a)

As shown in (10a), onset nasals form correspondence relations with other sonorant consonants and voiced stops, and it is ill-formed when they do not (10c). In contrast, coda nasals do not form correspondence relations with other sonorant consonants and voiced stops, and it is ill-formed when they do (10d).

To briefly summarize, correspondence can be conditioned not only on the featural makeup of the correspondents, but can also be conditioned on intervening material, and on the prosodic role of the correspondents.

[1] Bennett (2013)'s analysis of Kinyarwandan retroflex harmony relies on intermediate, local correspondence chains that are unbroken by noncoronal obstruents, which are not specified for [±retroflex], but blocked by coronals that are necessarily specified as [−retroflex]. Such an analysis is untenable for Slovenian, in which intervening [+anterior] coronals [n] and [r] do not block harmony (as in (6a) [ʒapɔr-niʃki] 'prison (adj)').

2.3 Some unattested correspondence conditions

Bennett (2013) argues persuasively that surface correspondence should be an equivalence relation, but it is certainly not the case that every possible equivalence relation is an attested surface correspondence relation. The following highlights this with two logically possible, but unattested, sets of conditions on correspondence, taken from McMullin and Hansson (2016).

The first is a correspondence relation as in (11), in which the correspondence relation picks out even and odd pairs of syllable-adjacent sibilants.

(11) a. **s**a**s**aʃa b. **ʃ**a**ʃ**a**s**a**s**a c. ***s**a**s**a**s**a

In (11), in a sequence of three syllable-adjacent sibilants, only two will pair off, leaving the third to only correspond with itself (11a) (cf. (11c)). In an a sequence of four sibilants, the correspondence relation splits them up into adjacent pairs (11b). In general, in sequences of even sibilants, correspondence will divide them into a sets of adjacent pairs, and in sequences of odd sibilants all except for one will be divided into pairs. This is pattern is unattested, but as McMullin and Hansson (2016) discuss and as shown below in §5.1, it can be generated in OT when conditions on syllable-adjacency compete with the definition of correspondence as transitive.

The second logically possible, but unattested correspondence pattern is what McMullin and Hansson (2016) call ‘agreement-by-proxy’, in which two dissimilar segments do not correspond unless a third segment, similar to each, is present in the word. The example they give involves two similarity conditions on correspondence: one in which obstruents of the same place correspond (as in (12a) below) and all continuant obstruents correspond (as in (12c) below). In the following example, the LDCA pattern is regressive assimilation to a following [+voice] obstruent.

(12) a. **z**a**d**a, **g**a**x**a b. **s**a**g**a c. **s**a**x**a d. **z**a**g**a**x**a, ***s**a**g**a**x**a

In the absence of other obstruents, obstruents of different place do not correspond, as in (12b). However, if we assume the transitivity of correspondence, then if we add another consonant that agrees in place with one obstruent and continuancy in another, then all three must correspond, as in (12d). Thus, this third obstruent that sits midway in similarity between the other two obstruents serves as a ‘proxy’ that conditions their correspondence. This is a logically possible equivalence relation one could define (and, as pointed out in McMullin and Hansson (2016) and described in §5.1, one can define in OT), but is not actually attested empirically.

Thus, a theory of correspondence should meaningfully distinguish between the attested conditions on correspondence in §2.2 and the unattested conditions just discussed. The remainder of the paper argues that theory that states phonological structure is FO-definable best accounts for this typological variation.

3 Computational characterizations of phonological structure

That phonological structure should be FO-definable from strings is directly related to restrictive computational characterizations of phonotactic patterns. This section surveys results that show that phonotactics are FO-definable sets of strings, and discusses how to connect these results directly to statements about phonological structure. An important consequence of this, highlighted in §3.3, is that FO-definable structure does not increase expressivity beyond FO-definable sets of strings.

3.1 Computational characterizations of phonology

Recent work has marshaled evidence for theories of phonotactic patterns based on their computational properties (Heinz, 2009, 2010; Graf, 2010a,b; Heinz et al., 2011; Rogers et al., 2013; Lai, 2015; McMullin and Hansson, 2016). We can view a phonotactic pattern as a (potentially infinite) *set* of strings, and then study the expressive power required to describe that set (Heinz, 2010). Formal language theory, a subfield of theoretical computer science, then provides method for organizing sets of strings into classes with well-studied properties (overviews can be found in textbooks such as Hopcroft et al., 2006). A theory of phonotactics based in these properties thus makes hard claims about bounds cognitive complexity of phonotactic patterns (Rogers and Pullum, 2011; Rogers et al., 2013), and how they are learned (Heinz, 2009, 2010; Lai, 2015; McMullin and Hansson, 2015; Jardine and Heinz, 2016; Jardine and McMullin, 2017).

For example, the Navajo sibilant harmony pattern described in §2.1 can be thought of the infinite set S_{Nav} in (13) of strings of segments that conform to the Navajo sibilant harmony pattern.

(13) $S_{\text{Nav}} = \{$ʃilí̃ːʔ, sìʔá, ʃiteːʒ, dasdoːlis, dasdodododoːlis, ... $\}$

The strings not in S_{Nav} are thus the strings that do *not* conform to the Navajo pattern, as exemplified in (14).

(14) ʃiteːz $\notin S_{\text{Nav}}$, sìteːʒ $\notin S_{\text{Nav}}$, dasdodododoːliʃ $\notin S_{\text{Nav}}$, ...

The question is: what computational power is required for a procedure that distinguishes the strings in S_{Nav} from the strings not in S_{Nav}? The classifications in formal language theory are based on the complexity of these procedures.

A procedure for checking membership in S_{Nav} can be very simple; if for any two segments in the string one is in the set {s, z} and the other is in the set {ʃ, ʒ}, then the string is not in S_{Nav}. (This is simplifying somewhat; to consider a larger range of sibilants, one simply needs to expand the sets of [+ant] and [–ant] sibilants.)

This is so simple as to seem obvious, or perhaps even trivial. However, the importance of the simplicity of checking membership in S_{Nav} reveals itself when put in the context of all the logically possible procedures for checking membership in a set of strings. For example, consider a set whose membership-checking procedure counts the number of sibilants in a string and says the string is in the set when the number is even and is not in the set when the number is odd. Let S_{Even} denote this set of strings; examples are given in (15).

(15) S_{Even} = {ʃiteːʒ, sìteːʒ, ʃiteːz, dasdoːlis, daʃdoːlis, daʃdoʃdaʃdos, ...}
ʃilı̋ːʔ $\notin S_{\text{Even}}$, sìʔá $\notin S_{\text{Even}}$, daʃdoʃdaʃ $\notin S_{\text{Even}}$, ...

For S_{Even}, the strings 'ʃiteːʒ,' 'sìteːʒ,' and 'ʃiteːz' are all members of the set, because they all contain an even number of sibilants, regardless of whether or not those sibilants agree. Conversely, 'sìʔá' and 'daʃdoʃdaʃ' are *not* members of the set because they have an odd number of sibilants.

It can be shown that the membership-decision procedure for S_{Even} is categorically more complex than that for S_{Nav}. In precise terms, S_{Nav} is in the *star-free* (SF) complexity class of string sets, which is a proper subset of the *regular* (REG) class (McNaughton and Papert, 1971). The set S_{Even} is properly in REG, meaning it is in REG but not SF. The reason why will be discussed momentarily.

The other important fact is that S_{Even} is, from the perspective of phonotactics, bizarre. No pattern in either the typologies of Hansson (2001, 2010) or Rose and Walker (2004) makes reference to anything like the number of a particular type of consonant in the word. Thus, a hypothesis that phonotactics are SF correctly excludes S_{Even} from its predicted typology based on the type of computations required to judge well-formedness of a string with respect to the pattern.

Studies of phonotactics have found them to be at most SF (Heinz, 2009, 2010; Graf, 2010a,b; Heinz et al., 2011; Rogers et al., 2013; McMullin and Hansson, 2016) with two potential exceptions (see Fn. 8). In fact, these studies have argued for more restrictive complexity classes for phonotactics (see §5.2), but this paper will focus on the SF class for two reasons. One, as to be discussed in §5.2, SF more readily captures the full range of attested LDCA patterns. Two, and more importantly for the purposes of this paper, the SF class has a natural definition in terms of FO logic (McNaughton and Papert, 1971), as detailed next in §3.2. As is then shown in §3.3, FO logic allows us to define phonological structure in a way that is in line with results from computational characterizations of phonology.

3.2 First-order logic over strings

Formal language theory provides different, but converging, ways to characterize the complexity classes discussed above (e.g., finite-state machines; for an overview see Hopcroft et al. 2006). Mathematical logic is one well-established method (Büchi, 1960; McNaughton and Papert, 1971; Immerman, 1980; Thomas, 1982; Rogers, 1997; Rogers et al., 2013). This is because satisfaction in a particular logical language can model the kinds of computations that are characteristic of a complexity class. The following shows how FO logic does this for sets in the SF class.

In order to define FO logic over strings, we need to explicitly define strings using *finite model theory* (interested readers are referred to Libkin 2004 for a thorough introduction). First, we fix an *alphabet*, or finite set of symbols. This can include entire set of symbols in the IPA, but for the present purposes let us, without loss of generality, consider a simplified alphabet of {s, ʃ, a, t}. We can uniquely identify a string over this alphabet by specifying a *domain* D of positions, a *precedence relation* $<$ that forms a linear order over these positions, and for each symbol a in our alphabet, a *unary relation* indicating the subset P_a of D that specifies which symbols in D are labeled a. Thus, for an alphabet {s, ʃ, a, t} strings are of the form in (16).

(16) $\langle D; <, P_{\mathrm{s}}, P_{\mathrm{ʃ}}, P_{\mathrm{t}}, P_{\mathrm{a}} \rangle$

Such a general shape of a structure is called a *signature*. A specific instantiation of a signature is called a *model*; a model of the string 'saʃa' is given in Fig. 1.

$$\begin{array}{lrcl} \langle & D & = & \{1,2,3,4\}; \\ & < & = & \{(1,2),(1,3),(1,4), \\ & & & \;(2,3),(2,4),(3,4)\ \}, \\ & P_{\mathrm{s}} & = & \{1\}, \quad P_{\mathrm{ʃ}} \ = \ \{3\}, \\ & P_{\mathrm{a}} & = & \{2,4\}, \ P_t \ = \ \{\} \qquad \rangle \end{array}$$

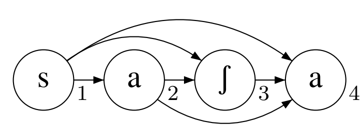

Figure 1: Mathematical and graphical description of a model of the string 'saʃa'. Elements in D are depicted as indexed circles, pairs in the order $<$ depicted as arrows, and membership in the sets P_{s}, $P_{\mathrm{ʃ}}$, P_{a}, and P_{t} is indicated by the label on each circle.

The string 'saʃa' has five elements in its domain, indexed in Fig. 1 using the natural numbers $\{1,2,3,4,5\}$. The precedence relation $<$ is thus the usual order over these numbers, and the sets P_{s}, $P_{\text{ʃ}}$, P_{a}, and P_{t} indicate which elements in D are labeled with each symbol. For example, the elements 2 and 4, which correspond to the second and fourth positions in ‘saʃa’, are both in the set P_{a}. Note that P_{t} is empty, as no positions in ‘saʃa’ are labeled ‘t’.

Every string one can create out of the alphabet {s, ʃ, a, t} thus can be represented with a model of the signature in (16). As an additional example, a model for the string ‘ʃataʃ’ is given in Fig. 2.

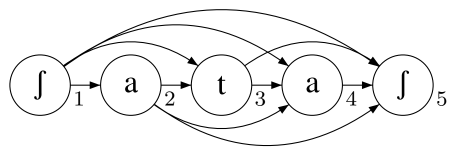

Figure 2: Graphical representation of a model of the string ‘ʃataʃ.’

Given a fixed signature (i.e., given a fixed alphabet), we can then define a FO logic for strings using the relations in that signature. Let x, y, z, etc., be *variables* and then let $x < y$ and $P_a(x)$ for each a in the alphabet be *atomic predicates*. These variables will be assigned values in D in a model; thus, $x < y$ will evaluate to true if and only if x and y are assigned to a pair in $<$ in that model. For example, $x < y$ is true in Fig. 1 when x is assigned to 1 and y is assigned to 5 (because $(1, 5)$ is a pair in $<$), but not when x is assigned to 3 and y is assigned to 2 (because $(3, 2)$ is not in $<$). Similarly, $P_a(x)$ is true in a model only when x is assigned to an element in P_a in that model. For example, $P_{\text{s}}(x)$ is true in Fig. 1 when x is assigned to 1, but not when it is assigned to 2. We also add an atomic predicate $x = y$ that evaluates to true when x and y are assigned to the same element in D.

Atomic predicates are also *predicates*. We define the full set of predicates recursively using the standard Boolean connectives and quantifiers $\forall$ and $\exists$ as follows. For predicates φ and ψ, $\neg\varphi$, $\varphi \wedge \psi$, $\varphi \vee \psi$, $\varphi \rightarrow \psi$, and $\varphi \leftrightarrow \psi$ are also predicates, and we assume the usual evaluation of their respective truth values (e.g., $\varphi \wedge \psi$ is true in a model when both φ *and* ψ are true in the model). Parentheses will be used to disambiguate; e.g., $(\varphi \wedge \psi) \rightarrow \rho$ is distinct from $\varphi \wedge (\psi \rightarrow \rho)$. Let the notation $\varphi(x)$ mean that x is the only variable in the predicate φ not already bound by a quantifier. Then $(\forall x)[\varphi(x)]$ and $(\exists x)[\varphi(x)]$ are also predicates. Quantifiers control the assignment of variables: $(\forall x)[\varphi(x)]$ is true when for *all* assignments of x to elements in the model's domain D, $\varphi(x)$ is true; $(\exists x)[\varphi(x)]$ is true when for *at least one* assignment of x to an element in the model's domain D, $\varphi(x)$ is true. For example, $(\forall x)[P_{\text{s}}(x)]$ is false in Fig. 1 because while $P_{\text{s}}(x)$ is true when x is assigned to 1, it is false when x is assigned to 2. In contrast, this means that $(\exists x)[P_{\text{s}}(x)]$ is true in Fig. 1, because $P_{\text{s}}(x)$ is true for at least one assignment of x. As shorthand, predicates with nested, adjacent quantifiers like $(\exists x)[(\exists y)[\varphi(x, y)]]$ will be written $(\exists x, y)[\varphi(x, y)]$.

FO logic is the set of all predicates defined as above; an example is in (17).

(17) $\varphi_{\text{Nav}} = (\forall x, y)[\neg(P_{\text{ʃ}}(x) \wedge P_{\text{s}}(y))]$

This predicate can be described in English as follows: for all x and y, it cannot be the case that x is a 'ʃ' and y is a 's'. If we try evaluating this predicate in the model in Fig. 1, then it will evaluate to false. When x is assigned to 3 and y is assigned to 1, then the predicate inside the quantification fails: 3 is in 'ʃ' and 1 is in 's'. So the innermost conjunction is true, which makes its negation false. Thus, there is an assignment of x and y such that the predicate inside the quantification is false, and so the entire quantification is false. (Note that there is no requirement that x must be assigned to an element that precedes the element y is assigned to.) In contrast, for the model in Fig. 2, for every assignment of x and y to pairs in D, the inner predicate will be true, and so the entire predicate will be true.

Predicates whose variables are all bound by a quantifier describe sets of strings. Consider φ_{Nav}: for any string which contains both an 'ʃ', and a 's', φ_{Nav} will evaluate to false in its model, and for any string without both an 'ʃ' and a 's', φ_{Nav} will evaluate to true in this model. A predicate in which all variables are bound by a quantifier is called a *sentence*. A classic result of formal language theory is that FO sentences describe exactly the SF sets (McNaughton and Papert, 1971); thus, no FO sentence can describe the properly REG set S_{Even}.

Lemma 1 *There is no FO sentence that describes* S_{Even}.

The proof of Lemma 1 is technical, but it is a well-established fact that FO logic cannot check whether there are an even or odd number of elements in a model. (For a discussion aimed at linguists, see Graf 2010b). However, we can get the intuition why not by looking at a logic that *can* describe S_{Even}: *monadic second-order* (MSO) logic, which describes exactly REG (Büchi, 1960). MSO logic is defined in the same way as FO logic, but it extends the FO definitions with *set variables* X, Y, Z, etc., that are assigned to arbitrary sets of elements in the domain. Informally, we can define S_{Even} by a MSO sentence that defines a set X representing the even sibilants in a word as counted from the left edge, and ensures that the rightmost sibilant is in X.

These logical characterizations say something meaningful about the computations required for the two sets. The set X used for defining even sibilants is *abstract* in the sense that it is not one of the sets P_a in the signature. In contrast, FO cannot make any such reference to abstract sets. If, as phonologists, we consider strings of segments, and take the labels in our sets P_a to represent *perceptible* information in the string, then an FO description necessarily only refers to perceptible information in the string (including order). Thus, a FO theory of phonological structure has the cognitive interpretation that it must be derived from perceptible information in the surface string.[2] We now turn to how to articulate such a theory.

[2]Morphological and prosodic information complicates this picture somewhat. LDCA, and thus correspondence,

3.3 First-order definitions of phonological structure

We can reconcile theories of phonological structure with computational characterizations of phonology based in strings through logic, which also gives us a powerful way to talk about the relationship between different kinds of structure. This comes from the theory of *logical transformations* from one structure into another (Courcelle, 1994; Engelfriet and Hoogeboom, 2001; Courcelle et al., 2012), which are based on logical *interpretations* (Enderton, 1972; Hodges, 1997) of the output structure in terms of the logic of the input structure. (This is thus not unlike the logical descriptions of phonological structures in Declarative Phonology (Scobbie et al., 1996).)

The figure in Fig. 3 illustrates the question for surface correspondence. As discussed in §2.1, surface correspondence is an additional relation $\mathcal{R}$ between surface segments (Walker, 2001; Rose and Walker, 2004; Hansson, 2001; Bennett, 2013). What do we gain by adding this relation? In other words, what kind of information does the correspondence relation represent? We can determine this precisely by defining, a logical transformation, the translation indicated in Fig. 3 from strings in a signature like that in (16) to a new signature with an additional binary relation $\mathcal{R}$.

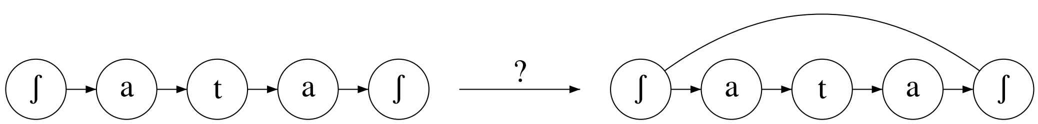

$\langle D; <, P_{\text{s}}, P_{\text{ʃ}}, P_{\text{t}}, P_{\text{a}} \rangle$ $\qquad$ $\langle D; <, \mathcal{R}, P_{\text{s}}, P_{\text{ʃ}}, P_{\text{t}}, P_{\text{a}} \rangle$

Figure 3: Comparison of a string model (left) to a string plus surface correspondence model (right). The correspondence relation, denoted $\mathcal{R}$ in the model on the right, is depicted by the unordered curved line. The precedence relation in each model has been abbreviated for visual clarity.

The first step is to understand that we can define new predicates in our FO logic without increasing its expressive power. To illustrate, we can define predicates expressing featural information using the basic string signature from (16). For example, the following defines what it means to be a sibilant in that signature:

(18) $[\text{+stri}](x) \stackrel{\text{def}}{=} P_{\text{s}}(x) \vee P_{\text{ʃ}}(x)$

The equation in (18) defines a new predicate $[\text{+stri}](x)$ that is true when x is assigned to either a 's' or 'ʃ'in the string model. Because the right-hand side of the equation is a FO

is often confined to a stem- or root-domain (Hansson, 2001, 2010; Rose and Walker, 2004; Bennett, 2013). Domain information is not always (if ever) directly perceptible, but we can argue that that information comes from the morphology and so can be included into the signature without it representing the same kind of 'abstract' information as do the arbitrary sets of MSO. See also Fn. 5 and §6.

predicate as defined above, [+stri](x) doesn't extend the power of our logic—it just serves as a shorthand for the predicate on the right-hand side.

This is exactly the technique we can use to define one structure in terms of another. The theory of logical transformations states that we can define a mapping from the set of *input* structures in one signature to the set of *output* structures in another—e.g., the mapping depicted in Fig. 3—by defining each relation in the signature of the output in terms of the logic of the input (Courcelle, 1994; Engelfriet and Hoogeboom, 2001; Courcelle et al., 2012). If this can be done in the FO logic of the input signature, then FO sentences in the output signature are equivalent to FO sentences in the input signature. In other words, the output signature is *FO-equivalent* to the input signature.

To see how this works for surface correspondence, consider the correspondence relation in Navajo, as represented in (3), in which any two sibilants must be in correspondence. First, we use [+stri](x) from (18) to define the following binary predicate that is true for any two elements that are both stridents.

(19) [+stri]$(x, y) \overset{\text{def}}{=}$ [+stri]$(x) \wedge$ [+stri](y)

The Navajo correspondence relation can thus be defined as in (20).[3]

(20) $x\mathcal{R}y \overset{\text{def}}{=} x = y \vee$ [+stri](x, y)

Interpreted as building structure on a string model, the definition in (20) adds a correspondence relation between every segment and itself (because, recall, we are assuming surface correspondence is an equivalence relation) and between every pair of stridents. Examples are given in Fig. 4.

Fig. 4a repeats the mapping from Fig. 3, only with some more detail. In the model on the left side, a correspondence relation $\mathcal{R}$ has been created in exactly the pairs for which $x\mathcal{R}y$ in (20) is true—including loops indicating that each element corresponds with itself. Fig. 4b shows how this correspondence relation will be built with more than one pair of sibilants; the definition for $x\mathcal{R}y$ in (20) will in this situation pair each sibilant with every other sibilant. Finally, Fig. 4c highlights that even non-agreeing sibilants will be in correspondence, because (as per the usual definition of surface correspondence) the definition for $x\mathcal{R}y$ in (20) makes no reference to agreement. Thus, for the Navajo correspondence relation, the question posited in Fig. 3 has been answered: $\mathcal{R}$ is FO-definable, as witnessed by (20). Note that the function of this definition is very similar to a CORR[+stri] constraint that enforces correspondents between stridents (e.g. as in Bennett, 2013).

[3] To make this an asymmetric correspondence relation a la Walker (2001), we simply add another condition: $x\mathcal{R}y \overset{\text{def}}{=} x = y \vee ($[+stri]$(x, y) \wedge x < y)$.

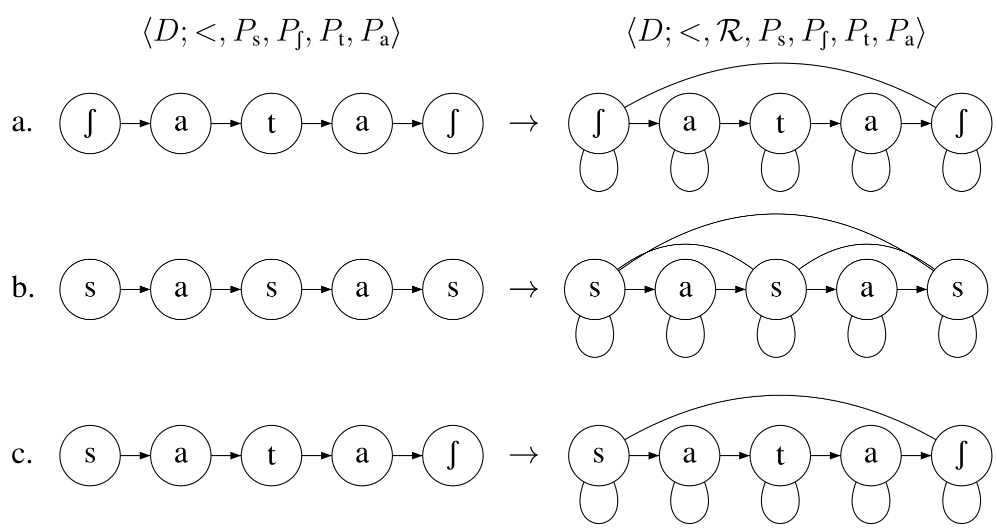

Figure 4: Examples of adding surface correspondence to strings as defined in (20).

As noted in (3c), agreement is then enforced through constraints that forbid a particular configuration of corresponding segments. The constraint in (3c) for Navajo is repeated below, using our explicit graphical notation for correspondence models.[4]

(21) * 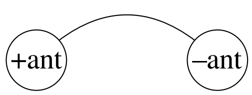

The sentence in (21) is a FO definition of a standard IDENT-CC[±ant] constraint (c.f. Bennett, 2013): if two segments correspond, then they both must either be [+anterior] or [–anterior] (if they are specified for that feature at all). The constraint in (21) is a *banned substructure constraint*: it marks as ill-formed any structure containing it. The structure in (21) references [±ant] features, which are not technically in the model, but we have already seen how such features can be interpreted in terms of the model. Thus, for example, Fig. 4c *does* contain (21), as it includes a [+ant] and a [–ant] sibilant that are in correspondence, and so thus is ill-formed. In contrast, Fig. 4a does not contain (21), and so is thus well-formed.

Banned substructure constraints, when interpreted as inviolable constraints over the surface form, are computationally very simple kinds of constraints, and are strictly less expressive than FO sentences (Rogers et al., 2013; Jardine and Heinz, in press; Jardine, 2017a). To illustrate, a FO sentence equivalent to (21) is given in (22).

(22) $\neg(\exists x, y)\,[\, x\mathcal{R}y \wedge [\text{+ant}](x) \wedge [\text{–ant}](y)]$

[4]To create a directional version of this constraint, a la Rose and Walker (2004)'s ID-$\text{C}_\text{L}\text{C}_\text{R}$, we simply additionally draw an arrow indicating a $<$ relation between the first node and second.

Importantly, the above sentence is written in the FO logic of the Navajo correspondence signature $\langle D; <, \mathcal{R}, P_{\mathrm{s}}, P_{\mathrm{ʃ}}, P_{\mathrm{t}}, P_{\mathrm{a}}\rangle$. It describes sets of models in that signature. However, because we know that $\mathcal{R}$ is definable in the FO logic of the string signature $\langle D; <, P_{\mathrm{s}}, P_{\mathrm{ʃ}}, P_{\mathrm{t}}, P_{\mathrm{a}}\rangle$, *we know that this sentence has a translation in the FO logic of the string signature*. This is a *guaranteed* result: because $\mathcal{R}$ has a FO definition, for any constraint we can write as a sentence in the FO language of the $\mathcal{R}$ model, there is an equivalent FO sentence in the string model.

Thus, the following assumes that agreement patterns are the result of banned substructure constraints defined over language-specific correspondence relations. As the correspondence relation contains some FO-definable information, this increases the power of banned substructure constraints as compared to those over strings (see §5.2). However, the following will show that we can account for the attested language-specific conditions on correspondence using a small set of similar predicates. Thus, through these constraints, phonological structure allows access to the formal expressive power of FO in a constrained way.

Once the right correspondence relation is in place, all of the agreement patterns discussed in this paper can be described by constraints like (21)/(22). The main goal, then, is to show that these language-specific correspondence conditions are indeed FO-definable. The following thus abstracts away from the actual agreement constraints and thus focuses on FO definitions of the correspondence relations.

4 The first-order definability of correspondence

The following analyzes the language-specific constraints on correspondence in terms of FO logic. In doing so, we discover an even stronger characterization: all of these conditions are expressible with universal ($\forall$) quantification. We also discover that conditions on correspondence are essentially just conjunctions and disjunctions of FO predicates that follow a few basic templates. §4.2 will then prove for the unattested pairing correspondence that there is no FO definition, and conjecture that the restriction to universal quantification excludes agreement-by-proxy.

4.1 Attested correspondence conditions

The previous section has already shown that the basic correspondence condition on Navajo—simply that two segments are in correspondents if they are both sibilants—is FO-definable. The following shows that the rest of the attested conditions on correspondence in §2.2 are FO-definable.

We begin with Koyra, which exemplified a correspondence relation conditioned on distance. Examples are repeated below in (23) from (5).

(23) a. goːtʃ-uʃ d. *goːtʃ-us 'cause to pull' (4a)

b. dʒaʃ-uʃ-eʃːe e. *dʒaʃ-us-esːe 'let him/them frighten someone' (4a)

c. ʃod-us f. *ʃod-uʃ 'cause to uproot' (4b)

Recall that correspondence in Koyra works as follows: syllable-adjacent sibilants must be in correspondence, as in (23a) versus (23d); sequences of syllable-adjacent sibilants are all in correspondence, as in (23b) versus (23e); and non-syllable adjacent sibilants do not correspond, as in (23c) versus (23f).

We can define this correspondence, drawing from McMullin and Hansson (2016)'s characterization of it as *transvocalic* agreement, as transvocalic correspondence. That is, sibilants only separated by vowels, and other sibilants, are in correspondence. The first step in defining such a property is to identify when an element intervenes between two other elements. This is captured by (24).

(24) $\texttt{inter}(x, y, z) \stackrel{\text{def}}{=} (x < z \wedge z < y) \vee (y < z \wedge z < x)$

The predicate $\texttt{inter}(x, y, z)$ is true when z intervenes between x and y—note that it is specifically designed so that the order of x and y does not matter. (This aids in capturing a symmetric correspondence relation.) The property of being transvocalic can then be captured by the following FO predicate.

(25) $\texttt{transV}(x, y) \stackrel{\text{def}}{=} (\forall z)\left[\big(\texttt{inter}(x, y, z) \wedge \text{C}(z)\big) \rightarrow \text{[+stri]}(z)\right]$

The predicate $\texttt{transV}(x, y)$ is thus true when for any consonant z intervening between x and y, that consonant must be a sibilant. The correspondence relation for Koyra is thus as in (26).

(26) $x\mathcal{R}y \stackrel{\text{def}}{=} x = y \vee \big(\text{[+stri]}(x, y) \wedge \texttt{transV}(x, y)\big)$

The definition in (26) is just the condition on (20) with $\texttt{transV}(x, y)$ as an additional condition conjoined to the similarity condition $\text{[+stri]}(x, y)$. Examples are given below in Fig. 5. As they are redundant, the reflexive pairs in the correspondence relation will be henceforth omitted from the diagrams.

The contrast between Fig. 5a and Fig. 5b captures the core point of the Koyra correspondence pattern: sibilants in adjacent syllables correspond, but sibilants separated by more than one syllable not containing a sibilant do not. Specifically, correspondence is blocked in Fig. 5b by $\texttt{transV}(x, y)$, as this predicate is not true for the two sibilants in Fig. 5b. However, sibilants in nonadjacent syllables can correspond as long as the intervening syllables also only contain sibilants, as in Fig. 5a. Thus, correspondence condition on syllable-adjacency can be captured in FO.

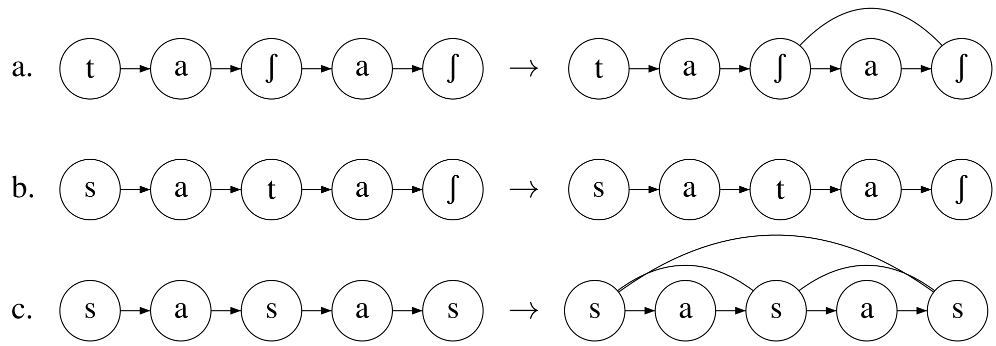

Figure 5: Examples of adding surface correspondence to strings for Koyra as defined in (26). Loops indicating the reflexive pairs in the correspondence relation have been omitted from the diagram.

Correspondence that is blocked by intervening material can also be defined in a very similar way. Recall that in Slovenian, sibilant harmony is blocked by intervening non-retroflex coronals and non-sibilant coronals, respectively. The correspondence pattern for Slovenian is repeated below in (27) from (7).

(27) a. po**ʒ**abi-**ʃ** c. * po**z**abi-**ʃ** '(you) forget' (6a)

b. **z**i**d**a-**ʃ** d. * **ʒ**i**d**a-**ʃ** '(you) build' (6b)

In Slovenian, sibilants must correspond unless a non-sibilant coronal obstruent intervenes, as in (27b) versus (27d). The following predicate tests for an intervening non-sibilant coronal, $\texttt{transS}(x,y)$ in a similar way to $\texttt{transV}(x,y)$ for Koyra.

(28) $\texttt{transS}(x,y) \stackrel{\text{def}}{=} (\forall z)\left[\big(\texttt{inter}(x,y,z) \wedge [\text{–son}](z) \wedge [\text{cor}](z)\big) \rightarrow [\text{+stri}](z)\right]$

The correspondence relation for Slovenian is thus as in (29).

(29) $x\mathcal{R}y \stackrel{\text{def}}{=} x = y \vee \big([\text{+stri}](x,y) \wedge \texttt{transS}(x,y)\big)$

Examples are given in Fig. 6. Note that they include an extra symbol 'n' in their signature, to illustrate that a sonorant coronal does not block correspondence.

Thus, conditions on intervening material can be expressed by FO predicates that differ only in specification of intervening material.[5]

Finally, the prosodic role of potential correspondents also can play a role in the conditions on correspondence, as we saw in Kikongo. As illustrated in (30) below, repeated from (10),

[5] Bennett (2013)'s family of CC-EDGE constraints which prohibit correspondence across particular domain boundaries can be defined in FO in the same way if the boundary is included in the string signature.

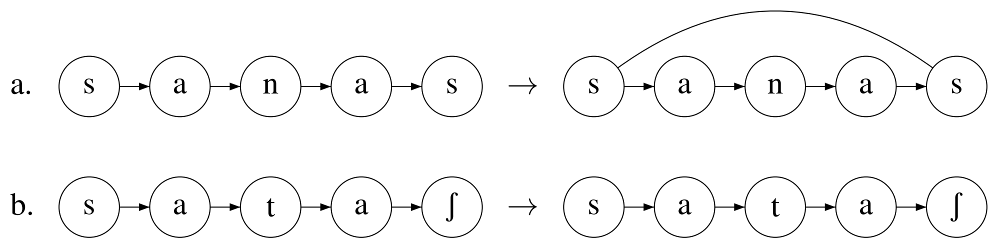

Figure 6: Examples of adding surface correspondence to strings for Slovenian as defined in (29).

only *onset* nasals correspond; coda nasals (as in (30b) and (d)) do not correspond with any other segments.

(30) a. ku.ki.**n**i.**n**a c. *ku.ki.**n**i.**l**a 'to dance for' (8b)
b. ba**n**.ti.ki.**d**i d. *ba**n**.ti.ki.**n**i 'begun' (9a)

To do this, we define predicates that identify codas and onsets, and create a similarity predicate based on syllable role. In Kikongo, coda consonants are those for whom any immediately following segment is a consonant.

(31) $\texttt{coda}(x) \stackrel{\text{def}}{=} (\forall y,z)[\quad \mathrm{C}(x) \wedge \big((x<y \wedge \neg(x<z \wedge z<y)) \rightarrow \mathrm{C}(y)\big) \quad]$

In (31), x is a consonant and for any y that immediately follows x, y is also a consonant. That y *immediately* follows x is captured by the predicate $(x<y \wedge \neg(x<z \wedge z<y))$, which ensures that no z intervenes between x and y. The predicate defining onsets is similar, demanding the following element to be a vowel.[6]

(32) $\texttt{onset}(x) \stackrel{\text{def}}{=} (\forall y,z)[\quad \mathrm{C}(x) \wedge \big((x<y \wedge \neg(x<z \wedge z<y) \ \rightarrow \mathrm{V}(y)\big) \quad]$

Using these two predicates we can define a new predicate $\texttt{SRole}(x,y)$ (after Rose and Walker (2004)'s constraint), which states that x and y are either both an onset or both a coda.

(33) $\texttt{SRole}(x,y) \stackrel{\text{def}}{=} (\texttt{onset}(x) \wedge \texttt{onset}(y)) \vee (\texttt{coda}(x) \wedge \texttt{coda}(y))$

[6]Both of these definitions are vacuously true when no segments follow x; that is, when x is the last segment in the word. This is of course correct for $\texttt{coda}(x)$ but not $\texttt{onset}(x)$. This can be remedied by adding word boundaries to the signature (see also fn. 5), in which case $\texttt{coda}(x)$ would need to be modified to check that a following element is either a consonant or a word boundary.

In FO, the definition of the correspondence relation thus states that correspondents have the same syllable role. Logically, this is added via conjunction to the similarity condition in Kikongo. Note that in Kikongo correspondence, nasals must correspond both with voiced stops and liquids. To do this, the definition below states that x and y are both either [+voiced, –continuant] (which links nasals and voiced stops) or [+sonorant, +consonantal] (which links nasals and liquids).The full definition for the correspondence relation thus conjoins this condition to the syllable role condition:

(34) $x\mathcal{R}y \overset{\text{def}}{=} x = y \vee \Big(\big(\text{[+voi,–cont]}(x,y) \vee \text{[+son,+cons]}(x,y)\big) \wedge \texttt{SRole}(x,y)\Big)$

Thus, correspondents must match the featural similarity and also either both be codas or neither be codas. (Technically, this conflates onsets and syllable nuclei, but as Kikongo arguably does not have syllabic consonants we can ignore this distinction for the sake of simplicity. Word-initial NC sequences exist but these are best analyzed as single segments (Ao, 1991), and regardless, they fall outside of the domain of agreement (Rose and Walker, 2004)). Examples are given in Fig. 7.

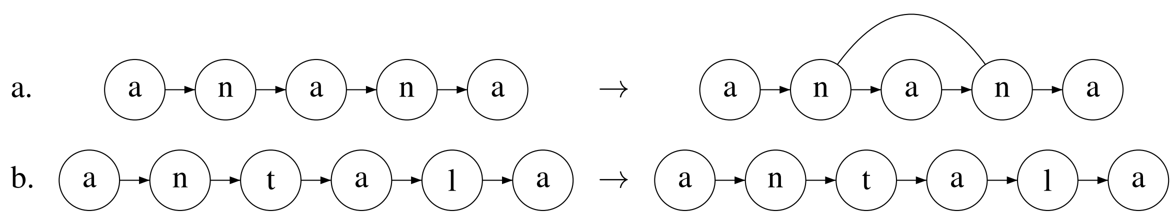

Figure 7: Examples of adding surface correspondence to strings for Kikongo as defined in (34).

To summarize, the attested conditions on correspondence outlined in §2.2—similarity, syllable-adjacency, dependence on intervening material, and dependence on syllable roles—are all FO-definable. Two points will be of interest in the remainder of the paper. The first is that all of these definitions required only use of universal (∀) quantification (if that). It is known that use of a single quantifier is strictly less powerful than the use of alternating universal and existential quantifiers (Place and Zeitoun, 2015).[7] We can thus tentatively posit an even stronger hypothesis for phonological structure: it can only be defined using universal quantification. This will factor into the following discussion of agreement-by-proxy. Second, all of these definitions involved conjunctions and disjunctions of conditions on the basic sim-

[7] Strictly speaking, results about complexity of quantifier alternation only hold in *prenex normal form*, in which all quantifiers appear at the very beginning of the predicate. This is because that, using $\rightarrow$ or $\neg$, one can define $\forall$ from $\exists$. There are some embedded $\forall$ quantifiers in the above predicates (specifically in $\texttt{transV}(x,y)$, etc.), but since these are not under the scope of $\neg$ or $\rightarrow$ we know there is an equivalent $\forall$-only predicate in prenex normal form.

ilarity constraint: correspondents must be similar in some way, *and* they satisfy some other constraint.

4.2 Unattested correspondence conditions

As explained in §2.3, one unattested, yet logically possible correspondence relation partitions correspondents into disjoint, adjacent correspondence pairs, as repeated below in (35).

(35) a. **sasa**ʃa b. ʃaʃa**sasa** c. ***sasasa**

Regardless, it can be shown that there is no FO definition for $x\mathcal{R}y$ under this version of correspondence.

Theorem 1 *A correspondence relation $x\mathcal{R}y$ which partitions the set of correspondents in a string into either disjoint pairs or singletons is not FO-definable.*

Proof: (Sketch.) By contradiction. Say $x\mathcal{R}y$ is definable in the FO-logic of string models. Then so is the sentence

$$(\forall x \exists y)[x\mathcal{R}y \wedge \neg x = y]$$

However, this sentence describes exactly the set of strings in which any sequence of syllable-adjacent sibilants must be even (note that the structure in (11b) satisfies this sentence but the structure in (11a) does not). It can be shown from Lemma 1 that this is not FO-definable. (Technically, S_{Even} counts sibilants over the entire word, while the above statement only looks at sequences of adjacent sibilants, but both depend on whether or not the sequence contains an even number of sibilants.) Because the rest of the sentence is built out of FO predicates, $x\mathcal{R}y$ must not be FO-definable. □

The other unattested pattern from §2.3, the 'agreement-by-proxy' pattern, is FO-definable but appears to require existential quantification. Recall that in this pattern, obstruents that differ in both place and continuancy do not correspond, unless there is a third continuant obstruent that agrees in continuance with one and place in the other. The examples from (12) are repeated below in (36).

(36) a. **z**a**d**a, **g**a**x**a b. **saga** c. **s**a**x**a d. **z**a**g**a**x**a ***s**a**g**a**x**a

First, we need to define the correspondence relation to occur between obstruents of the same place (as in (36a)) and continuant obstruents (as in (36c)). The following defines what it means to have the same place feature.

(37) $[\alpha\text{place}](x,y) \overset{\text{def}}{=} \big([\text{cor}](x) \wedge [\text{cor}](y)]\big) \vee \big([\text{lab}](x) \wedge [\text{lab}](y)]\big) \vee \big([\text{vel}](x) \wedge [\text{vel}](y)]\big)$

To capture a correspondence relation that connects segments that satisfy either similarity condition, we can use logical or, as in (34) for Kikongo. However, that is not enough: two segments that fail either similarity condition can still correspond if there is a 'proxy', as in (36d). This also must be included in the disjunction of the similarity conditions, as in (38).

(38) $x\mathcal{R}y \overset{\text{def}}{=} \; x = y \vee [\text{–son,}\alpha\text{place}](x,y) \vee [\text{–son,+cont}](x,y) \vee$
$(\exists z)[[\text{–son,+cont}](y,z) \wedge [\text{–son,}\alpha\text{place}](x,z)]$

The last conjunct in (38) checks for the existence of a for any x and y that fail the place and continuancy similarity conditions, they can still correspond if there is still some z that agrees in continuancy with x and in place with y. However, formally, this is a different kind of predicate than the ones we have used so far: it includes the quantifier $\exists$. This is a formal expression of the intuition of why the agreement-by-proxy correspondence relation is bizarre: instead of checking conditions only on potential correspondents, it searches for the existence of some third proxy that shares properties with these correspondents.

By restricting ourselves to definitions that only use $\forall$ quantification, we likely can exclude such correspondence conditions from the predicted typology. This requires that the following conjecture is true.

Conjecture 1 *There is no predicate that is logically equivalent to* (38) *that only includes* $\forall$ *quantifiers.*

A proof of this conjecture would be technical, and shall be left for future work.

5 Comparison to previous approaches

This paper has advocated for a view of phonological structure in which it is explicitly defined from surface strings, and has shown that restricting these definitions to a fragment of FO logic provides a restrictive, yet sufficient theory of phonological patterns. The following compares this to two other current approaches to LDCA: correspondence theory in OT and formal language-theoretic complexity classes.

5.1 Comparison to Optimality-Theoretic definitions of correspondence

At first blush, a stark contrast between the definitions for correspondence provided here and the general framework of OT is that OT is predicated on the idea of *violable* constraints, while the logical definitions given in this paper are *inviolable*, as in they are defined to apply across-the-board. However, upon closer inspection, the OT constraints that build correspondence function categorically; for the attested patterns, they are either true for the entire set of surface forms or they are not.

In OT, language-specific conditions on correspondence are built out of the interaction of of violable CORR constraints, which enforce similar segments to correspond, and other constraints. Thus, for example, to capture correspondence in Navajo, where all sibilants correspond, one would rank a constraint CORR[+stri] highly. Note, however, that counting of the violations is not necessary: all surviving candidates in an OT analysis for Navajo satisfy CORR[+stri] categorically (as all sibilants correspond and thus agree). Thus, CORR[+stri] ends up being equivalent to the following sentence.

(39) $(\forall x, y)[\text{[+stri]}(x, y) \rightarrow x\mathcal{R}y]$

That is, all of the candidates that satisfy CORR[+stri] also satisfy (39). Note that (39) is also true of any output structure of the transduction adding correspondence relations to strings according to the definition of correspondence for Navajo in (20).

This is true of other conditions on correspondence. Rose and Walker (2004)'s analysis of Kikongo uses a constraint SROLE which states that all correspondents must have the same syllable role. In Kikongo, this constraint is adhered to categorically by all attested candidates, just as pairs of correspondents in the structures resulting from definition (34) for correspondence in Kikongo obey the predicate $\texttt{SRole}(x, y)$.

In contrast, gradient evaluation of correspondents constraints can lead to pathologies. McMullin and Hansson (2016) show how the unattested pattern discussed in (2.3) in which the correspondence relation partitions sibilants into syllable-adjacent pairs. This occurs when the constraint PROXIMITY (Rose and Walker, 2004) is used under the assumption that correspondence is an equivalence relation. PROXIMITY, was originally defined as follows by Rose and Walker (2004) to account for syllable-adjacency conditions on correspondence.

(40) PROXIMITY (Rose and Walker, 2004, p.494) Correspondent segments are located in adjacent syllables.

More precisely, PROXIMITY assigns a violation for each pair of correspondents that are not in adjacent syllables. When this is highly ranked, with any candidate with a sequence of syllable-adjacent sibilants, the winning candidate will be the one in which they are split into pairs, as the candidate with the transitive correspondence relation (the one that is actually attested in Koyra) violates PROXIMITY.

(41)

	/ʃasasasa/	Prox	Corr[+stri]	CC-Id[–ant]
	a. ʃasasasa		*** *!* *	
✗	b. ʃaʃaʃaʃa	*!* *		
☞	c. ʃaʃasasa		** **	

d.* ʃaʃaʃaʃa ** *

Given such a ranking, the candidate that splits correspondents into adjacent pairs, (41c), wins over the candidate in which all sibilants correspond, (41b), because the latter violates Prox. This candidate also wins out against the candidate (41a), which violates Corr[+stri] more times. A potential candidate which obtains the correct assimilation pattern by a 'chain' of local correspondence relations, (41d), is not generated by Gen, assuming surface correspondence must be transitive.

Bennett (2013)'s solution for problems like this one is to replace Proximity with a CC-SylAdj constraint that is defined almost identically to the $\mathtt{transS}(x, y)$ predicate from §4.1: it checks for sequences of adjacent correspondents that are not broken with a non-correspondent. However, in this case as well, the winning candidate will always satisfy CC-SylAdj absolutely—counting of violations no longer factors into the conditions on correspondence.

However, the other unattested correspondence pattern can be generated in OT without counting violations. As McMullin and Hansson (2016) explain, 'agreement-by-proxy' is a pathology that arises in OT when two similar correspondence constraints, combined with the assumption of transitivity, allows for a third 'proxy' segment to serve as a bridge between two other segments that would not otherwise correspond. Normally, two highly-ranked correspondence constraints Corr[αF] and Corr[βG] will simply result in a surface condition such that two segments correspond if they both are [αF] *or* if they both are [βG] (c.f. the disjunction in definition (10) for correspondence in Kikongo). However, Corr[αF] and Corr[βG] target overlapping sets of segments and transitivity is assumed, then the presence of a segment targeted by both will cause any segment targeted by Corr[αF] to correspond with any segment targeted by Corr[βG]. In this case, characterizing the surface condition on correspondence involves not just simple disjunction, but the addition of an $\exists$ clause (as is necessary in §4.2's definition (38) of this pathological correspondence relation).

Thus, the interaction of correspondence conditions in the attested patterns boils down to simple conjunction and disjunction, but this is not true for these unattested patterns. Thus, while one step towards squaring OT analyses with the results in this paper is to ensure that constraints are written in ∀-quantified FO statements, this cannot be the entire story: the interaction of violable constraints in OT can lead to more complex relations on the surface. Thus, a study of the relationship between the logic of individual OT constraints and the logic of the surface forms resulting from their interaction could be a potentially fruitful area for future work.

5.2 Comparison to Formal Language Theoretic accounts of LDCA

LDCA patterns have also been integral to work on formal-language theoretic characterizations of phonotactics. Three sub-classes of the SF class have been identified as theories of LDCA: the *Strictly Piecewise (SP)* class (Heinz, 2010; Fu et al., 2014), the *Tier-based Strictly Local (TSL)* class (Heinz et al., 2011; McMullin and Hansson, 2016), and the *Interval-based Strictly Piecewise Class (IBSP)* class (Graf, 2017), all of which are based in some way on banned substructure constraints.

First, the SP class is defined on banned substructure constraints that operate directly over $<$ string models. For example, LDCA in Navajo, which holds that all sibilants must have the same value for anteriority, would be captured in SP by the banned substructure constraints in (42a).

(42) a. 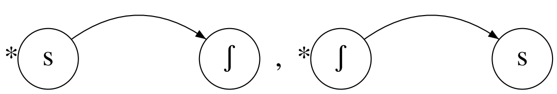

b. 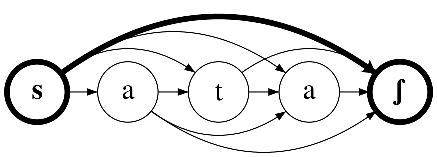

The interpretation of (42a) is simple: [s] cannot precede [ʃ], and [ʃ] cannot precede [s]. Any strings in which the sibilants do not agree will contain one of the banned substructures in (42a), as highlighted in bold in (42a). Thus, like correspondence theory, SP explains LDCA through a direct relation between the agreeing segments—the difference is that in SP, this relation is the precedence relation already present in the string model. Heinz (2010) shows that the SP class captures the fact that LDCA is not mediated through intervening segments, and also is efficiently learnable.

However, there are two issues with SP. First, because the precedence relation connects all segments in a string, it is possible to write constraints banning arbitrary combinations of symbols. For example, possible to write a banned substructure constraint that forbids an [s] followed by an [e]. Correspondence theory avoids such constraints by drawing on a finite set of conditions that restrict the correspondence relation to similar segments. One way of adjusting the SP theory of LDCA is to restrict the set of banned substructure constraints to only refer to similar segments. However, this is functionally equivalent to positing that constraints operate over models in which a precedence relation only exists between similar segments—in such a model, the precedence relation is essentially an asymmetric correspondence relation.

A second issue with the SP class is that it is essentially blind to intervening material, it cannot capture blocking patterns (Heinz, 2010; Heinz et al., 2011). The TSL class (Heinz et al., 2011) was proposed to capture such patterns. The TSL class is defined as local constraints operating over a *successor* relation defined over a *tier*, defined as a subset of the alphabet. This successor relation is a kind of intransitive precedence relation—it only relates an element and its immediately following like element. The tier is a parameter, meaning that in phonological terms it is determined on a language-specific basis. Thus, for example, to capture the LDCA pattern in Slovenian, in which an intervening [t] blocks agreement, we define a tier predicate $T(x)$ as in (43a) consisting of [s, t, ʃ].

(43) a. $T(x) \overset{\text{def}}{=} P_{\text{s}}(x) \vee P_{\text{ʃ}}(x) \vee P_{\text{t}}(x)$
b. $x \triangleleft_T y \overset{\text{def}}{=} T(x) \wedge T(y) \wedge (\forall z)[(x < z \wedge z < y) \rightarrow \neg T(z)]$

A successor relation $x \triangleleft_T y$ is then defined over this tier as in (43b), which states that $x \triangleleft_T y$ is true when both x and y are on the tier and no z on the tier intervenes. Viewed in terms of structure building as in §4, adding this 'tier successor' relation results in structures as in Fig. 8.

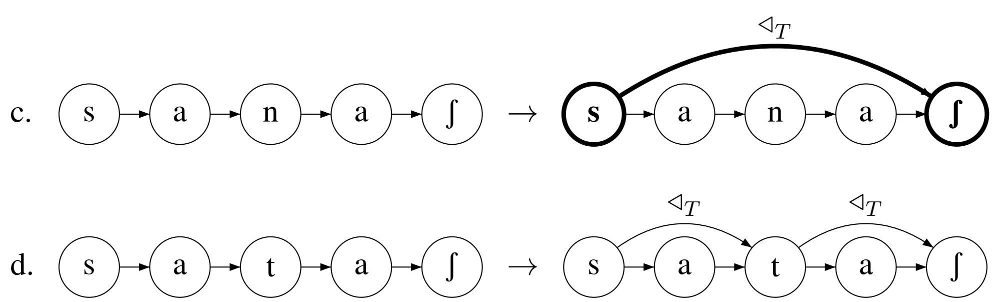

Figure 8: Adding tier successor relation (labeled with $\triangleleft_T$) as defined in (43).

In TSL, the arrows banned substructure constraints as in in (42a) are interpreted as referring only to the $\triangleleft_T$ relation. By doing so, (42a) captures the Slovenian pattern: [s] cannot precede [ʃ], as highlighted in Fig. 8c, unless there is an intervening [t], as in Fig. 8d. As exemplified in Fig. 8d, such a structure contains neither substructure in (42a); this draws from the intransitivity of $\triangleleft_T$.

TSL is similar to the proposal advanced in this paper in that it builds an enriched representation from a string and then constraints operate over that representation. Additionally, note that, indirectly, the $\triangleleft_T$ relation defined in (43) mirrors the behavior of the correspondence relation defined for Slovenian in §4.1, (29). The latter definition creates a correspondence relation between two sibilants if and only if there is no intervening non-strident coronal; (43) creates a $\triangleleft_T$ relation between the two in the same condition (otherwise, it draws intermediate $\triangleleft_T$ relations). Note that this falls out of the logical definition of $\triangleleft_T$, which refers to the intervening material in a way similar to the condition $\mathtt{transS}(x, y)$ for correspondence in (29).

The TSL class covers a wide range of patterns (McMullin, 2016; McMullin and Hansson, 2016), and is also efficiently learnable (Jardine and Heinz, 2016; Jardine and McMullin, 2017). It also excludes patterns like agreement-by-proxy (McMullin and Hansson, 2016). However, the TSL class cannot capture the syllable role condition in Kikongo LDCA. In TSL membership on the tier is based on the label of each segment. As shown above, syllable role is based not only on the label of a segment, but on local information around that segment. This is capturable if syllable information is included in the string representation, but strictly speaking, this increases the expressive power of TSL. In the other direction, as seen in (43), the tier in a TSL set is always FO-definable, meaning that any pattern describable in TSL can be mimicked with a FO-definable correspondence relation. Thus, for example, while this paper has not discussed long-distance *dissimilation*, dissimilation appears to be TSL (with the exception of Sundanese; Heinz 2010; Heinz et al. 2011; McMullin 2016). This strongly suggests that dissimilation correspondence conditions are also FO-definable.

A third subclass that bears mentioning is the IBSP class, a generalization of SP and TSL which Graf (2017) proposes to capture domain information, as well as capture some patterns that have been shown to be outside the SP and TSL classes. Essentially, IBSP uses FO definitions to parameterize SP constraints by what intervening material can come between them (like TSL, and the correspondence definitions in this paper) and the domains over which they operate. Because it properly includes SP and TSL, IBSP can capture all of the above patterns, as well as some patterns they cannot capture. However, IBSP cannot mix local and non-local constraints (Graf, 2017, p. 400), so it cannot capture Kikongo as well. Finally, Graf proposes no constraints on quantification, so it is likely that IBSP can capture patterns that the correspondence definitions proposed in §4 cannot. One further interesting difference between IBSP and the current proposal is that it defines domains and intervening material directly in its constraints, whereas the proposal here is to build these definitions into structure. It remains to be seen in detail how these two views can be integrated, but as they both use FO definitions, it is not hard to see how to accomplish this.

6 Discussion

The preceding sections have established a FO-definable theory of correspondence and shown how it is comparable to OT and formal-language theoretic explanations of LDCA. However, several questions remain.

First, are other kinds of phonological structure FO-definable? Jardine (2017b) shows that two-tier autosegmental representations are FO-definable from strings, but for other kinds of structure, such as syllables, feet, and higher-level prosodic structure, this remains an open question. This question may seem especially pressing for metrical structure, as the result in §4.2 that dividing sibilants into adjacent pairs is not FO-definable appears to bear negatively on the FO-definability of binary feet. However, virtually all stress patterns are SF (Rogers et al., 2013), and thus describable with FO sentences. This essentially guarantees that foot structure in each language is FO-definable. To see how this is possible in a simple case, consider the stress pattern of Pintupi, with initial stress and alternating secondary stress on every other non-final syllable thereafter (Hansen and Hansen, 1969).

(44) $\acute{\sigma}\sigma$, $\acute{\sigma}\sigma\sigma$, $\acute{\sigma}\sigma\grave{\sigma}\sigma$, $\acute{\sigma}\sigma\grave{\sigma}\sigma\sigma$, $\acute{\sigma}\sigma\grave{\sigma}\sigma\grave{\sigma}\sigma$, $\acute{\sigma}\sigma\grave{\sigma}\sigma\grave{\sigma}\sigma\sigma$,...

The predicate $\texttt{foot}(x, y)$ in (45a) builds a relation representing trochaic feet; the bracketing in (45b) of forms from (44) indicates pairs of syllables for which $\texttt{foot}(x, y)$ is true. (Feet are often given their own elements in a structure—see, e.g., Bird 1995—but for simplicity we can abstract away from that here. To see how structure-building predicates can add elements to a structure, see Jardine 2017b.)

(45) a. $\texttt{foot}(x, y) \overset{\text{def}}{=} \texttt{stressed}(x) \wedge \texttt{unstressed}(y) \wedge (\forall z)\,[\neg(x < z \wedge z < y)]$

b. $(\acute{\sigma}\sigma)$, $(\acute{\sigma}\sigma)\sigma$, $(\acute{\sigma}\sigma)(\grave{\sigma}\sigma)$, $(\acute{\sigma}\sigma)(\grave{\sigma}\sigma)\sigma$, $(\acute{\sigma}\sigma)(\grave{\sigma}\sigma)(\grave{\sigma}\sigma)$, $(\acute{\sigma}\sigma)(\grave{\sigma}\sigma)(\grave{\sigma}\sigma)\sigma$, ...

The predicate $\texttt{foot}(x, y)$ is true when x is a stressed syllable and y is an unstressed syllable immediately following x. Because most other stress patterns are definable with FO constraints, it is likely that these can be reformulated using FO definitions of metrical structure, although a full investigation of this will be left to future work.[8]

With respect to syllables and higher levels of prosodic structure, their FO-definability remains an open question. However, there is reason to believe that future work will find them to be FO-definable. Syllabification is based on identification of nuclei and the local neighborhood around those nuclei; this is likely FO-definable; Strother-Garcia (to appear) demonstrates this for Imdlawn Tashlhiyt Berber. In contrast, some theories of higher-level prosodic structure posit recursive domains (Ito and Mester, 2007, 2009; Selkirk, 2011), which would not in

[8] The two known exceptions to the generalization that stress is SF are Cairene Arabic and Creek (Graf, 2010a,b), in which the placement of primary stress is based on an alternating secondary stress that, in at least some descriptions, is claimed to be 'latent' or not perceptible. If it is indeed not perceptible, then foot structure is not FO-definable in either of these languages. However, whether or not secondary stress is perceptible is controversial in at least Cairene Arabic (Halle and Vergnaud, 1987), and in either case foot structure could be FO-definable if we count the natural trochaic rhythm that humans perceive (Hayes, 1995) as 'perceptible' (i.e., we include it in the signature).

and of itself be FO-definable. However, this structure is usually defined by making reference to syntactic information, made explicit in, e.g., Selkirk's MATCH theory. It is likely the case that this prosodic structure is FO-definable if syntactic information is already present in the signature (either in string or tree form). This would mean that the recursion can be attributed entirely to the syntactic, not phonological, module. Regardless, the FO-definability of prosodic structure is an important topic for future research.

Another important open question is the learnability of FO-definable structure. As the FO definitions of correspondence in §4.1 were language-specific, a full theory of FO-definable representation should have some procedure for learning these language-specific definitions. More concretely, given a finite input of strings from a LDCA pattern, it should be possible to discover both a correspondence relation and the constraints responsible for the pattern. This is a challenging problem, as it requires a learner to acquire both hidden structure and a grammar from positive data. However, this exact problem has been solved for the TSL sets (Jardine and Heinz, 2016; Jardine and McMullin, 2017). Future work on learning language-specific correspondence relations and constraints can draw on these techniques.

Finally, this paper has focused on surface well-formedness, so an important remaining question is how logical definitions of structure can be related to theories of phonological *transformations* from an underlying form to surface form. In the same way we can study the computational properties of well-formedness in terms of sets, we can study phonological transformations in terms of maps from strings to strings (Johnson, 1972; Koskenniemi, 1983; Kaplan and Kay, 1994; Chandlee and Heinz, 2012; Heinz and Lai, 2013; Chandlee, 2014; Tesar, 2014; Jardine, 2016). Much of this work has focused on finite-state characterizations of phonology, instead of logical ones. However, the theory of logical transformations used to study the computational properties of maps between different kinds of structure can also be applied to study maps from phonological underlying representations to surface representations (Chandlee and Lindell, forthcoming; Heinz, forthcoming). While such characterizations for LDCA maps are still the subject of ongoing research (Chandlee, forthcoming), pursuing such characterizations in tandem with the study of maps between different kinds of structures promises to be a fruitful avenue for even more restrictive theories of phonological representation.

7 Conclusion

This paper has introduced logical characterizations of phonological structure, and demonstrated that FO characterizations of surface correspondence form a part of a theory of phonology that is both computationally restrictive and can express substantive constraints on phonological patterns. As discussed in the immediately preceding section, this raises a rich, yet approachable, set of research questions pursuing the relationship of logic to structure, grammars, and learning in phonology.

References

Ao, B. (1991). Kikongo nasal harmony and context-sensitive underspecification. *Linguistic Inquiry*, 22(2):193–196.

Bennett, W. (2013). *Dissimilation, Consonant Harmony, and Surface Correspondence*. PhD thesis, Rutgers, the State University of New Jersey.

Bird, S. (1995). *Computational phonology: A constraint-based approach*. Studies in Natural Language Processing. Cambridge University Press.

Büchi, J. R. (1960). Weak second-order arithmetic and finite automata. *Zeitschrift für Mathematische Logik und Grundlagen der Mathmatik*, 6:66—92.

Chandlee, J. (2014). *Strictly Local Phonological Processes*. PhD thesis, University of Delaware.

Chandlee, J. (forthcoming). Logical characterizations of local vs. long-distance phonology. In Heinz, J., editor, *Doing Computational Phonology*. OUP.

Chandlee, J. and Heinz, J. (2012). Bounded copying is subsequential: Implications for metathesis and reduplication. In *Proceedings of the 12th Meeting of the ACL Special Interest Group on Computational Morphology and Phonology*, pages 42–51, Montreal, Canada. Association for Computational Linguistics.

Chandlee, J. and Lindell, S. (forthcoming). A logical characterization of strictly local functions. In Heinz, J., editor, *Doing Computational Phonology*. OUP.

Courcelle, B. (1994). Monadic second-order definable graph transductions: a survey. *Theoretical Computer Science*, 126:53–75.

Courcelle, B., Engelfriet, J., and Nivat, M. (2012). *Graph structure and monadic second-order logic: A language-theoretic approach*. Cambridge University Press.

Elmedlaoui, M. (1995). *Aspects de representations phonologiques dans certains langues chamito-semitiques*. Rabat: Faculté des Lettres et des Sciences Humaines.

Enderton, H. (1972). *A mathematical introduction to logic*. Academic Press.

Engelfriet, J. and Hoogeboom, H. J. (2001). MSO definable string transductions and two-way finite-state transducers. *ACM Transations on Computational Logic*, 2:216–254.

Fu, J., Heinz, J., Jardine, A., and Tanner, H. G. (2014). Perception-based grammatical inference for adaptive systems. Talk at ICGI 2014, Kyoto, Japan.

Graf, T. (2010a). Comparing incomparable frameworks: A model theoretic approach to phonology. *University of Pennsylvania Working Papers in Linguistics*, 16(2):Article 10.

Graf, T. (2010b). Logics of phonological reasoning. Master's thesis, UCLA.

Graf, T. (2017). The power of locality domains in phonology. *Phonology*, 34:385–405.

Halle, M. and Vergnaud, J.-R. (1987). *An Essay on Stress*. Cambridge, MA: The MIT Press.

Hansen, K. and Hansen, L. (1969). Pintupi phonology. *Oceanic Linguistics*, 8:153–170.

Hansson, G. O. (2001). *Theoretical and Typological Issues in Consonant Harmony*. PhD thesis, University of California, Berkeley.

Hansson, G. O. (2010). *Consonant Harmony: Long-distance Interaction in Phonology*. Berkeley: University of California Press.

Hayes, B. (1995). *Metrical stress theory*. Chicago: The University of Chicago Press.

Hayward, R. J. (1982). Notes on the Koyra language. *Afrika und Ūbersee*, 65:211–268.

Heinz, J. (2009). On the role of locality in learning stress patterns. *Phonology*, 26:303–351.

Heinz, J. (2010). Learning long-distance phonotactics. *LI*, 41:623–661.

Heinz, J., editor (forthcoming). *Doing Computational Phonology*. Oxford: Oxford University Press.

Heinz, J. and Lai, R. (2013). Vowel harmony and subsequentiality. In Kornai, A. and Kuhlmann, M., editors, *Proceedings of the 13th Meeting on Mathematics of Language*, Sofia, Bulgaria.

Heinz, J., Rawal, C., and Tanner, H. G. (2011). Tier-based strictly local constraints for phonology. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics*, pages 58–64, Portland, Oregon, USA. Association for Computational Linguistics.

Hodges, W. (1997). *A Shorter Model Theory*. Cambridge: Cambridge University Press.

Hopcroft, J., Motwani, R., and Ullman, J. (2006). *Introduction to Automata Theory, Languages, and Computation*. Addison-Wesley, third edition.

Immerman, N. (1980). *First-order expressibility as a New Complexity Measure*. PhD thesis, Cornell University.

Ito, J. and Mester, A. (2007). Categories and projections in prosodic structure. In *Old World Conference in Phonology 4*.

Ito, J. and Mester, A. (2009). Trimming the prosodic hierarchy. In Borowsky, T., Kawahara, S., Shinya, T., and Sugahara, M., editors, *Prosody matters: Essays in honor of Elisabeth Selkirk*. London: Equinox Publishers.

Jardine, A. (2016). Computationally, tone is different. *Phonology*, 33:247–283.

Jardine, A. (2017a). The local nature of tone-association patterns. *Phonology*, 34:385–405.

Jardine, A. (2017b). On the logical complexity of autosegmental representations. In Kanazawa, M., de Groote, P., and Sadrzadeh, M., editors, *Proceedings of the 15th Meeting on the Mathematics of Language*, pages 22–35, London, UK. Association for Computational Linguistics.

Jardine, A. and Heinz, J. (2016). Learning tier-based strictly 2-local languages. *Transactions of the Association for Computational Linguistics*, 4:87–98.

Jardine, A. and Heinz, J. (in press). Markedness constraints are negative: an autosegmental constraint definition language. In *Proceedings of CLS 51*.

Jardine, A. and McMullin, K. (2017). Efficient learning of tier-based strictly k-local languages. In Drewes, F., Martín-Vide, C., and Truthe, B., editors, *Language and Automata Theory and Applications, 11th International Conference*, Lecture Notes in Computer Science, pages 64–76. Springer.

Johnson, C. D. (1972). *Formal aspects of phonological description*. Mouton.

Jurgec, P. (2011). *Feature Spreading 2.0: A Unified Theory of Assimilation*. PhD thesis, University of Tromsø. Available at lingbuzz/001281.

Kaplan, R. and Kay, M. (1994). Regular models of phonological rule systems. *Computational Linguistics*, 20:331–78.

Koskenniemi, K. (1983). *Two-level morphology: a general computational model for word-form recognition and production*, volume 11. Helsinki: University of Helsinki Department of General Linguistics.

Lai, R. (2015). Learnable versus unlearnable harmony patterns. *LI*, 46:425–451.

Libkin, L. (2004). *Elements of Finite Model Theory*. Berlin: Springer-Verlag.

McCarthy, J. and Prince, A. (1995). Faithfulness and reduplicative identity. In Beckman, J., Dickey, L. W., and Urbanczyk, S., editors, *Papers in Optimality Theory*, number 18 in University of Massuchusetts Occasional Papers in Linguistics, pages 249—384. University of Massachusetts.

McCarthy, J. and Prince, A. (1999). Faithfulness and identity in prosodic morphology. In Kager, R., van der Hulst, H., and Zonneveld, W., editors, *The Prosody Morphology Interface*, pages 218–309. Cambridge: Cambridge University Press.

McMullin, K. (2016). *Tier-based locality in long-distance phonotactics: learnability and typology*. PhD thesis, University of British Columbia.

McMullin, K. and Hansson, G. O. (2015). Locality in long-distance phonotactics: evidence for modular learning. In *NELS 44, Amherst, MA*. in press.

McMullin, K. and Hansson, G. O. (2016). Long-distance phonotactics as Tier-Based Strictly 2-Local languages. In *Proceedings of AMP 2015*.

McNaughton, R. and Papert, S. (1971). *Counter-Free Automata*. MIT Press.

Odden, D. (1994). Adjacency parameters in phonology. *Language*, 70(2):289–330.

Place, T. and Zeitoun, M. (2015). The tale of the quantifier alternation hierarchy of first-order logic over words. *ACM SIGLOG News*, 2(3):4–17.

Prince, A. and Smolensky, P. (1993). Optimality Theory: Constraint interaction in generative grammar. *Rutgers University Center for Cognitive Science Technical Report*, 2.

Prince, A. and Smolensky, P. (2004). *Optimality Theory: Constraint Interaction in Generative Grammar*. Blackwell Publishing.

Rogers, J. (1997). Strict lt2 : Regular :: Local : Recognizable. In Retoré, C., editor, *Logical Aspects of Computational Linguistics: First International Conference, LACL '96 Nancy, France, September 23–25, 1996 Selected Papers*, pages 366–385. Springer Berlin Heidelberg, Berlin, Heidelberg.

Rogers, J., Heinz, J., Fero, M., Hurst, J., Lambert, D., and Wibel, S. (2013). Cognitive and sub-regular complexity. In *Formal Grammar*, volume 8036 of *Lecture Notes in Computer Science*, pages 90–108. Springer.

Rogers, J. and Pullum, G. (2011). Aural pattern recognition experiments and the subregular hierarchy. *Journal of Logic, Language and Information*, 20:329–342.

Rose, S. and Walker, R. (2004). A typology of consonant agreement as correspondence. *Language*, 80:475–531.

Sapir, E. and Hoijer, H. (1967). *The phonology and morphology of the Navajo language*. Berkeley: University of California Press.

Scobbie, J. M., Coleman, J. S., and Bird, S. (1996). Key aspects of declarative phonology. In Durand, J. and Laks, B., editors, *Current Trends in Phonology: Models and Methods*, pages 685–709. Manchester, UK: European Studies Research Institute (ESRI), University of Salford.

Selkirk, E. (2011). The syntax-phonology interface. In Goldsmith, J., Riggle, J., and Yu, A., editors, *The Handbook of Phonological Theory*, pages 435–484. Wiley-Blackwell.

Strother-Garcia, K. (to appear). Imdlawn Tashlhiyt Berber syllabification is quantifier-free. In *Proceedings of the first annual meeting of the Society for Computation in Linguistics*.

Tesar, B. (2014). *Output-Driven Phonology: Theory and Learning*. Cambridge Studies in Linguistics.

Thomas, W. (1982). Classifying regular events in symbolic logic. *Journal of Computer and Systems Sciences*, 25:360—376.

Walker, R. (2000). Long-distance consonantal identity effects. In Billery, R. and Lillehaugen, B. D., editors, *WCCFL 19: Proceedings of the 19th West Coast Conference on Formal Linguistics*, pages 532–545. Somerville, MA: Cascadilla Press.

Walker, R. (2001). Consonantal correspondence. In Kirchner, R., Pater, J., and Wikeley, W., editors, *Proceedings of the Workshop on the Lexicon in Phonetics and Phonology. Papers in Experimental and Theoretical Linguistics*, volume 6, pages 73–84. Edmonton: Department of Linguistics, University of Alberta.

Walker, R., Byrd, D., and Mpiranya, F. (2008). An articulatory view of kinyarwanda coronal harmony. *Phonology*, 25:499–535.

Walker, R. and Mpiranya, F. (2006). On triggers and opacity in coronal harmony. In *Berkeley Linguistics Society (BLS)*, volume 31, pages 383–394.