Omnivorous Agreement, like Uyghur Backness Harmony, is a Challenge for Tier-Based Strict Locality

Allison Verbil
Department of Linguistics
University of California, Los Angeles
averbil@ucla.edu

Tim Hunter
Department of Linguistics
University of California, Los Angeles
timhunter@ucla.edu

Abstract

A well-known exception to the characterization that phonological patterns belong to the subregular class of TSL dependencies is found in Uyghur backness harmony (Mayer and Major, 2018). At the same time, a recent line of work has argued that many long-distance syntactic phenomena are subsumed by the TSL class, revealing an interesting parallel between phonology and syntax. We show that a certain omnivorous syntactic agreement pattern, namely Mundari object agreement (Murugesan et al., 2025), poses the same challenge to TSL as Uyghur backness harmony.

1 Introduction

It is generally accepted that phonological processes are at most regular stringsets and string-to-string mappings (Johnson, 1972; Kaplan and Kay, 1994), and can be computed by regular grammars and finite-state automata (Heinz, 2018). This regular class within the Chomsky Hierarchy is, however, argued to not be sufficiently restrictive in characterizing possible phonological patterns, and it has been claimed that a stronger characterization — that phonological stringsets are tier-based strictly local languages, which are subregular — holds (Heinz, 2018). This hypothesis claims that valid segmental phonological strings can be captured via constraints on substrings, except that these substrings may belong to a tier which contains only the segments relevant to a particular process (Heinz et al., 2011). Mayer and Major (2018) show, however, that a segmental process in Uyghur cannot be captured by a TSL language due to the specific behavior of potential harmony triggers in Uyghur.

At the same time, recent work has argued that something similar holds in the domain of syntax, namely that many agreement patterns fall within the class of TSL languages when construed as constraints on strings representing "spines" of a tree (Hanson, 2025). This paper will argue that a certain *omnivorous agreement* pattern (Nevins, 2011), namely object agreement in Mundari (Murugesan et al., 2025) poses a challenge for this approach to agreement that is analogous to the challenge posed by Uyghur in the phonological domain. The critical property of this omnivorous agreement pattern is that second person DPs, which rank in the middle of a three-level person scale, are sometimes, but not always, relevant for determining the form of the object marker; we argue that this is analogous to the way certain dorsal consonants in Uyghur are sometimes, but not always, relevant to the backness harmony. We will show that, because of this formal similarity, a natural extension of TSL that increases the complexity of the tier projection function to be sensitive to symbols recently projected onto the tier (Graf and Mayer, 2018; De Santo and Graf, 2019; Graf and Shafiei, 2019) captures both the Uyghur and Mundari patterns in a unified way.

The paper is structured as follows. §2 provides a background on TSL grammars in phonology and syntax, §3 summarizes Mayer and Major (2018)'s claims about Uyghur, and §4 shows Mundari omnivorous agreement behaves analogously to Uyghur. §5 presents potential objections, §6 introduces O-TSL grammars for Uyghur and Mundari, and §7 concludes.

2 Background on TSL

2.1 The TSL class

Σ represents a finite set of symbols referred to as an *alphabet*, and Σ^* denotes the set of all strings over Σ of any finite length. The empty string is ϵ. $\rtimes$ and $\ltimes$ will be edge markers, which denote the left and right edge of a word, respectively, and are not part of Σ. A string u is a substring of a string w iff $w = xuy$ for some strings $x, y \in \Sigma^*$. The k-factors of a string are its k-length substrings. The function $f_k(w)$ returns the set of k-factors of w.

Proceedings of the Society for Computation in Linguistics (SCiL) 2026, pages 129–137

Informally, a strictly k-local grammar is a finite set of forbidden substrings of length k, and the strictly local language generated by that grammar is the set of all strings that do not contain any of those forbidden substrings. More formally, we can say that a strictly k-local grammar has a finite set of k-factors $G \subseteq (\{\rtimes, \ltimes\} \cup \Sigma)^k$ which describes forbidden, or illicit, substrings. The language generated by a strictly k-local grammar G contains a string $w \in \Sigma^*$ iff $\rtimes w \ltimes$ contains no illicit substrings, i.e. $w \in L(G)$ iff $f_k(\rtimes w \ltimes) \cap G = \emptyset$. A language L is strictly local iff, for some $k \in \mathbb{N}$, there is a strictly k-local grammar that generates L.

Tier-based strictly local languages, as a class, properly contain the strictly-local languages. Informally, tier-based strictly local grammars differ from strictly-local grammars in that only symbols projected onto some tier are checked for illicit substrings. Formally, we can define a tier $\tau \subseteq \Sigma$, and a projection function $E_\tau(\sigma_1 \ldots \sigma_n) = u_1 \ldots u_n$, where $u_i = \sigma_i$ iff $\sigma_i \in \tau$ and $u_i = \epsilon$ otherwise. This removes symbols from the string that are not in τ. The k-factors for a k-TSL grammar are taken from $(\{\rtimes, \ltimes\} \cup \tau)^k$. The language generated by a k-TSL grammar G contains a string $w \in \Sigma^*$ iff the tier-projection of w, enriched with edge markers, contains no illicit substrings, i.e. $w \in L(G)$ iff $f_k(\rtimes E_\tau(w) \ltimes) \cap G = \emptyset$. A language L is TSL iff, for some $k \in \mathbb{N}$, there is a k-TSL grammar that generates L.

2.2 TSL patterns in phonology

As the k-factors in a TSL grammar are substrings on a tier, we can account for certain kinds of long-distance phonotactic dependencies like vowel and consonant harmony with blocking (Heinz et al., 2011). Consider a language like Georgian, which exhibits long-distance liquid dissimilation, where the suffix /-uri/ surfaces as [-uli] when preceded by an [r] anywhere in the stem (Fallon, 1993):

(1) a. /p'olon-uri/ → [p'o**l**on-u**r**i] 'Polish'
b. /ungr-uri/ → [ung**r**-u**l**i] 'Hungarian'
c. /gramat'ik'a-uri/ → [g**r**amat'ik'-u**l**i] 'grammatical'

However, dissimilation is blocked when /-uri/ is preceded by an intervening [l], as in [avst'**ral**-u**r**i] 'Australian' and [ka**r**t**l**-u**r**i] 'Kartvelian'. If our tier alphabet contains only the symbols $\tau = \{r,l\}$ we can forbid adjacent identical liquids by banning the 2-factors [rr] and [ll] on the tier, which will rule out any strings that do not obey dissimilation. For example, the tier projection of the illicit word *[ungr-uri] is [rr], which (non-properly) contains the banned substring [rr]. Cases where blocking occurs are straightforward to handle with this TSL-2 grammar; the tier projection of the licit string [kartl-uri] is [rlr], and therefore does not have two adjacent [r] segments, correctly predicting that [kartl-uri] is licit. The vast majority of observed phonological patterns that have been analyzed with TSL grammars have, like this example, used $k = 2$ (Jardine and Heinz, 2016).

2.3 TSL patterns in syntax

More recently there have been suggestions that TSL stringsets, also with $k = 2$, might play an interesting role in characterizing syntactic patterns (Graf and Shafiei, 2019; Hanson, 2025). The relevant syntactic objects are trees, as usual, but the idea in this work is that syntactic constraints might be expressible as conditions on strings that correspond to "spines" of the tree. These are known as c-strings.

We assume the relevant trees are headed binary trees[1] with lexical items at the leaves, and we consider *only pre-movement positions*. Any such tree is the maximal projection of some lexical item. The c-string of maximal projection XP is the string of lexical items comprised of: (i) the c-string of XP's complement (if it has one), followed by (ii) the heads of each of XP's specifiers (if it has any), in deepest-first order, followed by (iii) the head of XP. This is illustrated in Figure 1.

Following Hanson (2025), the simple English agreement pattern in (2) can be expressed in terms of the c-strings extracted from trees of the form shown in Figure 2.

(2) a. This dog chase<u>s</u> these cats
b. * This dog chase<u>∅</u> these cats
c. * These cats chase<u>s</u> this dog
d. These cats chase<u>∅</u> this dog

Taking into account (only) the base position of the subject, and abstracting away from whatever post-syntactic process leads to the T head surfacing in a position attached to the verb, we arrive at the four c-strings shown in (3).

(3) a. cats these chase this v $T_{\text{-s}}$

[1] By this we mean a binary tree where each internal node is designated as a projection of exactly one of its daughters. This is sufficient for the notions of maximal projection, head, complement and specifier to be determined in standard ways.

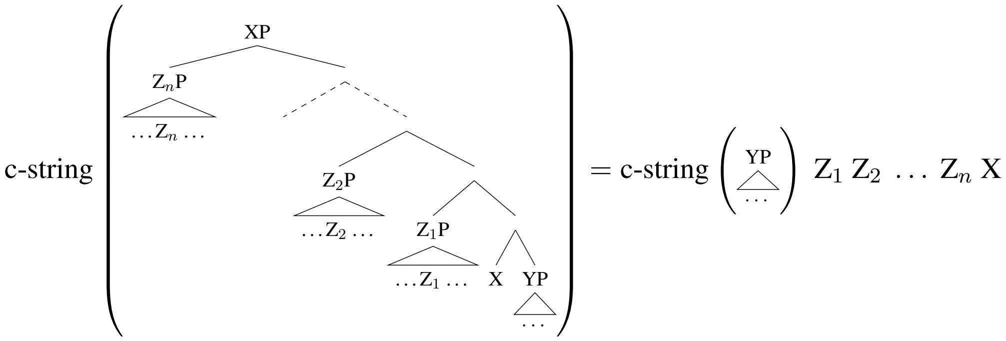

Figure 1: Schematic illustration of the definition of a tree's c-string. Notice that the c-string of XP includes the full c-string of its complement YP (which will end with the head Y), but only the heads Z_i of its specifiers.

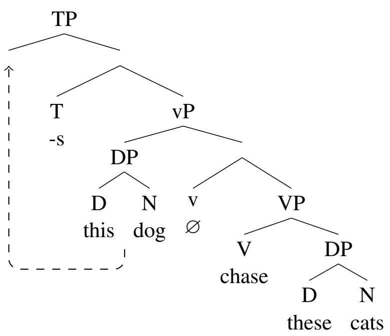

Figure 2: Tree structure for sentence (2a)

b. cats these chase this v $T_{\varnothing}$

c. dog this chase these v $T_{\text{-s}}$

d. dog this chase these v $T_{\varnothing}$

It will be useful to write these strings in a form that more clearly indicates the relevant linguistic properties. For this we adopt the strategy, discussed by Aksënova et al. (2016), of "relabeling" surface forms, in order to deal with strings of symbols such as D_{sg} and D_{pl} to collapse irrelevant distinctions (e.g. between 'these' and 'those'). Then instead of (3) we have (4).

(4) a. N D_{pl} V D_{sg} v $T_{\text{-s}}$

b. N D_{pl} V D_{sg} v $T_{\varnothing}$

c. N D_{sg} V D_{pl} v $T_{\text{-s}}$

d. N D_{sg} V D_{pl} v $T_{\varnothing}$

Picking out the c-strings corresponding to the acceptable sentences can be achieved by a TSL-2 grammar: specifically, if we project only D and finite T heads (i.e. $\tau = \{D_{\text{sg}}, D_{\text{pl}}, T_{\text{-s}}, T_{\varnothing}\}$), then the relevant requirement can be stated as a ban on the tier bigrams 'D_{sg} $T_{\varnothing}$' and 'D_{pl} $T_{\text{-s}}$'.

(5) a. N D_{pl} V D_{sg} v $T_{\text{-s}}$ $\xrightarrow{E_\tau}$ D_{pl} D_{sg} $T_{\text{-s}}$ ✓

b. N D_{pl} V D_{sg} v $T_{\varnothing}$ $\xrightarrow{E_\tau}$ D_{pl} D_{sg} $T_{\varnothing}$ ✗

c. N D_{sg} V D_{pl} v $T_{\text{-s}}$ $\xrightarrow{E_\tau}$ D_{sg} D_{pl} $T_{\text{-s}}$ ✗

d. N D_{sg} V D_{pl} v $T_{\varnothing}$ $\xrightarrow{E_\tau}$ D_{sg} D_{pl} $T_{\varnothing}$ ✓

The use of bigrams here matches the length of k-factors that have been used in phonology.

Although in these examples the T head is at a fixed distance from the D head it depends on (only v intervenes), the use of the tier allows for the same constraint to generalize to agreement in existentials; c-strings exclude the expletive 'there' for the same reason that they ignore post-movement positions.

(6) a. There seems to be a cat here
$\xrightarrow{E_\tau}$ D_{sg} $T_{\text{-s}}$ ✓

b. * There seem∅ to be a cat here
$\xrightarrow{E_\tau}$ D_{sg} $T_{\varnothing}$ ✗

c. * There seems to be two cats here
$\xrightarrow{E_\tau}$ D_{pl} $T_{\text{-s}}$ ✗

d. There seem∅ to be two cats here
$\xrightarrow{E_\tau}$ D_{pl} $T_{\varnothing}$ ✓

By also projecting C heads to the tier, and banning bigrams where a C head is followed by $\varnothing_{\text{T}}$, this TSL grammar over c-strings can capture the appropriate pattern for finite complements of verbs like 'seem' and for sentential subjects.

(7) It seems that these cats sit∅
w = V D_{pl} v $T_{\varnothing}$ C V v $T_{\text{-s}}$
$E_\tau(w)$ = D_{pl} $T_{\varnothing}$ C $T_{\text{-s}}$ ✓

(8) That these cats sit∅ surprises us
w_1 = D_{pl} V C v $T_{\text{-s}}$
w_2 = V D_{pl} v $T_{\varnothing}$
$E_\tau(w_1)$ = D_{pl} C $T_{\text{-s}}$ ✓
$E_\tau(w_2)$ = D_{pl} $T_{\varnothing}$ ✓

Note that for (8) there are two non-trivial c-strings: w_1 for the main clause's spine, on whose tier projection the C head separates the plural object D_{pl} ('us') from the T head, and w_2 for the sentential subject itself, where the standard agreement with 'these cats' is enforced as usual.

In sum, the TSL-2 grammar that covers the basic English agreement facts discussed here can be formalized as follows:

(9) $$\tau = \{D_{sg}, D_{pl}, C, T_{\text{-s}}, T_{\varnothing}\}$$
$$G = \{D_{sg}\ T_{\varnothing},\ D_{pl}\ T_{\text{-s}},\ C\ T_{\varnothing}\}$$

This TSL-2 grammar follows, in all essential respects, the proposal in Hanson (2025), although we have focused on characterizing the distributions of, and dependencies among, elements that are distinguished in relatively pre-theoretic terms (D_{sg}, $T_{\text{-s}}$, etc.), whereas Hanson presents grammars that regulate the distribution of diacritics ($[+\phi]$ and $[-\phi]$) that are assumed to drive a feature-copying agreement operation; see Hanson (2025, pp.62-63), for discussion of this choice point. We have opted for the former approach here because it more straightforwardly parallels the way vowel harmony patterns are typically described.

One more inconsequential difference between our presentation here and Hanson's is that Hanson writes c-strings in the reverse order, with heads at the beginning and embedded complement c-strings at the end, whereas we have followed the ordering introduced originally in Graf and Shafiei (2019); again, the reason for our choice here is to bring out the parallels with vowel harmony more clearly.

While our focus in this section has been on c-strings as a method of enforcing agreement patterns via TSL constraints, following Hanson (2025), Graf and Shafiei (2019) apply the c-string approach to other syntactic phenomena, and make use of a slightly more powerful formalism: IO-TSL grammars. We return to the relationship between our work and that of Graf and Shafiei (2019) in §6.2.

3 The challenge from Uyghur

One of the most well-known challenges to the generalization that harmony patterns in phonology are TSL comes from vowel harmony in Uyghur. Uyghur exhibits a backness harmony pattern which requires suffixes to agree in backness with certain vowels and consonants in the stem (Mayer and Major, 2018). The relevant segments that participate in harmony are shown in Table 1.

	Harmonizing vowels	**Harmonizing consonants**
Front	y, ø, æ	k, g
Back	u, o, a	q, ʁ

Table 1: Segments which participate in Uyghur backness harmony. The vowels [i, e] are transparent non-participatory vowels.

	Form	**Gloss**	**Harmony type**
a.	aʁinæ-**dæ**	friend-LOC	closest front vowel
b.	qoichi-**da**	shepherd-LOC	closest back vowel
c.	gezit-**tæ**	newspaper-LOC	closest front dorsal
d.	qirʁiz-**da**	Kyrgyz-LOC	closest back dorsal
e.	rak-**ta**	shrimp-LOC	closest back vowel, across front dorsal
f.	mæʃq-**tæ**	exercise-LOC	closest front vowel, across back dorsal
g.	it-**ta**	dog-LOC	no trigger, arbitrarily back
h.	biz-**dæ**	we-LOC	no trigger, arbitrarily front

Table 2: Examples of Uyghur backness harmony. Alternations of /-DA/ are bolded, and the harmony triggers are highlighted in blue. From Mayer and Major (2018).

The process for determining the backness value of a particular suffix can be described as follows:

(10) a. If there is a harmonizing vowel in the stem, the suffix should match the backness of the final harmonizing vowel: [-da] if back, [-dæ] if front.

b. If there's no harmonizing vowel in the stem, the suffix should match the backness of the final harmonizing consonant: [-da] if uvular, [-dæ] if velar.

c. If there's no harmonizing vowel or harmonizing dorsal in the stem, the suffix can surface as either [-da] or [-dæ].

Representative examples of the above process can be seen in Table 2, from Mayer and Major (2018). The suffix depicted is /-DA/, the locative case marker. Note that there is a voicing alternation in the suffix that is orthogonal to the present pattern; the description in (10) abstracts away from this alternation. In case (10c), it is lexically-determined, for any given stem, which of the two forms appears, but we treat this as distinct from the harmony requirement itself (i.e. while [it-**tæ**] is not actually a valid surface form, we do not take it to be in violation of the harmony pattern that we are modeling).

The essential takeaway is that harmonizing vowels will always take precedence over dorsal consonants as harmony triggers (rows e–f), and harmonizing consonants only serve as triggers when there is insufficient backness information available from the vowels (rows c–d). Mayer and Major's claim is that it is this exact nature — where both vowels and consonants can be harmony triggers, but consonants can only be considered in the absence of vowels — that is problematic for TSL.

To describe this pattern formally, we'll relabel all harmonizing front and back vowels in Table 2 as V_f and V_b respectively, and similarly relabel harmonizing consonants as C_f and C_b. Forms of the suffix with a front vowel are relabeled S_f, and those with a back vowel S_b. The pattern is at most regular, since it can be characterized by the regular expression $r_1 \mid r_2 \mid r_3 \mid r_4 \mid r_5$, where $V = V_f \mid V_b$, $C = C_f \mid C_b$, and:

(11)
$$\begin{aligned} r_1 &= (V \mid C \mid X)^* \, V_f \, (V_f \mid C \mid X)^* \, S_f \\ r_2 &= (V \mid C \mid X)^* \, V_b \, (V_b \mid C \mid X)^* \, S_b \\ r_3 &= (C \mid X)^* \, C_f \, (C_f \mid X)^* \, S_f \\ r_4 &= (C \mid X)^* \, C_b \, (C_b \mid X)^* \, S_b \\ r_5 &= X^* \, (S_f \mid S_b) \\ X &= \overline{(V \mid C \mid S_f \mid S_b)} \end{aligned}$$

Mayer and Major note that if we set aside cases where the suffix harmonizes with a consonant (i.e. consider only $r_1 \mid r_2 \mid r_5$), a TSL-2 grammar where we project all and only suffixes and harmonizing vowels would suffice.

(12)
$$\begin{aligned} \tau &= \{V_f, V_b, S_f, S_b\} \\ G &= \{V_f S_b,\ V_b S_f\} \end{aligned}$$

This implements a ban on vowel-suffix sequences on the tier that mismatch in backness, preventing forms like *[aʁinæ-da] where the suffix does not harmonize with the stem-final vowel.

The challenge comes from how to incorporate the role of the harmonizing consonants. If these harmonizing consonants are omitted from the tier, then they will never have any non-local effect at all on the predicted suffix forms, which clearly cannot work. If they are included in the tier, then this will allow correctly for cases like [gezit-**tæ**] where the harmonizing consonant is the only harmonizing element in the stem (rows c–d); but crucially, this will lead to incorrect results in cases where a harmonizing consonant intervenes between a harmonizing vowel and a suffix (rows e–f). And since there is no bound on the number of such intervening consonants, any k-TSL grammar will be unable to make the necessary distinctions between cases that lead to the tier projections in (13), because of the formal property that Mayer and Major refer to known as suffix substitution closure; intuitively, the critical vowel is too far away from the suffix for the k-factors to enforce the desired result.

(13)
$$\begin{aligned} &V_b \, C_f^{k-1} \, S_f \qquad && V_f \, C_f^{k-1} \, S_f \\ &V_b \, C_b^{k-1} \, S_b \qquad && V_f \, C_b^{k-1} \, S_b \end{aligned}$$

Put differently, there's no one-size-fits-all answer to the question of whether C_f and C_b should be projected onto the tier. For this pattern to be enforced by strictly-local constraints on a tier, the projection function would need to project C_f and C_b only in the absence of V_f or V_b.

4 The analogous challenge from Mundari

As discussed in §1, a recent line of work has argued that many long-distance syntactic dependencies, including agreement, are subsumed by the TSL-2 class (Hanson, 2025, among others). Here, we argue that *omnivorous agreement* presents a challenge to TSL that is analogous to the above problem presented by Uyghur. *Omnivorous agreement* (Nevins, 2011) refers to a pattern where an agreement marker tracks the highest-ranking argument along some given hierarchy, regardless of the relative positions of the arguments.

Mundari (Austroasiatic, Murugesan et al., 2025) exhibits omnivorous agreement in ditransitive constructions, where both the direct and indirect object compete for a single agreement slot[2]. The language shows complex interactions between person, number, and animacy hierarchies, but we restrict attention here to cases where both objects are singular and animate; in these scenarios, object agreement follows a 1>2>3 person hierarchy. For example, consider the interaction between first and second person DPs in (14):

(14) a. hon-ko aiɲ ke am ke-ko
children-PL 1SG EMP 2SG EMP-3PL.SM
ɛm-a-iɲ-ta-n-a
give-APPL-1SG.OM-PROG-ITR-IND
"Children are giving me to you."
1SG DO & 2SG IO

b. hon-ko am ke aiɲ ke-ko
children-PL 2SG EMP 1SG EMP-3PL.SM
ɛm-a-iɲ-ta-n-a
give-APPL-1SG.OM-PROG-ITR-IND

[2] Murugesan et al. (2025) note that this pattern is clitic doubling and not "true" inflectional agreement, but this is orthogonal to our argument.

"Children are giving you to me."
2SG DO & 1SG IO

Here, first person DPs are tracked instead of second person, regardless of whether the first person DP is the direct object or the indirect object. When no first person DP is present, second person DPs are tracked over third person DPs regardless of object status, as in (15):

(15) a. hon-ko am ke Ravi ke-ko
children-PL 2SG EMP Ravi EMP-3PL.SM
ɛm-a-m-ta-n-a
give-APPL-2SG.OM-PROG-ITR-IND
"Children are giving you to Ravi."
2SG DO & 3SG IO

b. hon-ko Ravi ke am ke-ko
children-PL Ravi EMP 2SG EMP-3PL.SM
ɛm-a-m-ta-n-a
give-APPL-2SG.OM-PROG-ITR-IND
"Children are giving Ravi to you."
3SG DO & 2SG IO

A combination of first and third person will yield the same agreement marker '-iɲ' as in (14); in a case with two third person (singular, animate) objects (i.e. no first or second person DPs are present), the distinct marker '-i' surfaces. In sum, the pattern can be described as follows:

(16) a. If there is a first person DP in the agreement domain, the object marker is '-iɲ'.
b. If there's no first person DP in the agreement domain, but there is a second person DP, the object marker is '-m'.
c. If there's no first person or second person DP in the agreement domain, the object marker is '-i'.

We take the underlying (pre-movement) structure of Mundari's ditransitives to be as shown in Figure 3, with the Appl head surfacing as the object

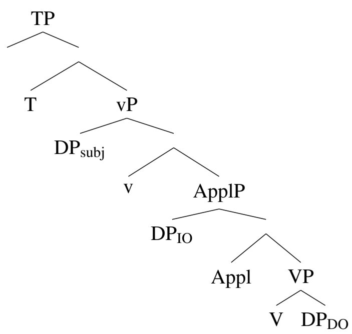

Figure 3

marker (Murugesan et al., 2025). As before, we work with a relabeling of syntactic elements that focuses attention on the critical distinctions: we use $\mathrm{Ap_A}$, $\mathrm{Ap_B}$ and $\mathrm{Ap_C}$ for applicative heads that surface as the '-iɲ', '-m' and '-i' object markers respectively, and $\mathrm{D_1}$, $\mathrm{D_2}$ and $\mathrm{D_3}$ for first, second and third person D heads. Given the tree structure in Figure 3, where the indirect object c-commands the direct object, the relevant c-strings for (14) and (15) will be as follows. (Note that the subject is third person throughout these examples; we focus only on the effect of varying the objects.)

(17) a. (14a), 1SG DO & 2SG IO
c-string: $\mathrm{D_1\ V\ D_2\ Ap_A\ D_3\ v\ T}$
b. (14b), 2SG DO & 1SG IO
c-string: $\mathrm{D_2\ V\ D_1\ Ap_A\ D_3\ v\ T}$

(18) a. (15a), 2SG DO & 3SG IO
c-string: $\mathrm{D_2\ V\ D_3\ Ap_B\ D_3\ v\ T}$
b. (15b), 3SG DO & 2SG IO
c-string: $\mathrm{D_3\ V\ D_2\ Ap_B\ D_3\ v\ T}$

(19) a. 3SG DO & 3SG IO
c-string: $\mathrm{D_3\ V\ D_3\ Ap_C\ D_3\ v\ T}$
b. 3SG DO & 3SG IO
c-string: $\mathrm{D_3\ V\ D_3\ Ap_C\ D_3\ v\ T}$

The pattern described in (16), construed as a constraint on c-strings, can therefore be characterized by the regular expression $r_1 \mid r_2 \mid r_3$, where:

(20)

$$r_1 = (\mathrm{D_1} \mid \mathrm{D_2} \mid \mathrm{X})^* \ \mathrm{D_1}\ (\mathrm{D_1} \mid \mathrm{D_2} \mid \mathrm{X})^* \ \mathrm{Ap_A}\ \mathrm{X}^*$$
$$r_2 = (\mathrm{D_2} \mid \mathrm{X})^* \ \mathrm{D_2}\ (\mathrm{D_2} \mid \mathrm{X})^* \ \mathrm{Ap_B}\ \mathrm{X}^*$$
$$r_3 = \mathrm{X}^* \ \mathrm{Ap_C}\ \mathrm{X}^*$$
$$\mathrm{X} = \overline{(\mathrm{D_1} \mid \mathrm{D_2} \mid \mathrm{Ap_A} \mid \mathrm{Ap_B} \mid \mathrm{Ap_C})}$$

Notice the parallels with (11), with $\mathrm{D_1}$ playing a role analogous to vowels, and $\mathrm{D_2}$ to consonants.

Like Uyghur, we will show that the problematic nature of this paradigm results from second person DPs only being considered as agreement goals in the absence of first person DPs; this is analogous to the way dorsal consonants in Uyghur are harmony triggers only in the absence of harmonizing vowels.

If we were not concerned with agreement triggered by second person DPs (i.e. if we consider only $r_1 \mid r_3$), then an appropriate TSL grammar could be constructed along the same lines as (12): project only first person DPs and object markers, and ensure that the choice between $\mathrm{Ap_A}$ and $\mathrm{Ap_C}$

tracks the presence or absence of D_1:

(21) $$\tau = \{D_1, Ap_A, Ap_C\}$$
$$G = \{D_1 Ap_C,\ \rtimes Ap_A\}$$

The challenge is in incorporating sensitivity to the presence of second person DPs. If we omit second person DPs from the tier, then it will be impossible to regulate the proper choice between Ap_B and Ap_C in (18) and (19). If we include second person DPs on the tier, then we can correctly account for (18) and (19), but the second person indirect object will intervene between the first person direct object and the applicative head in (17a)/(14a). As in the Uyghur case then, no choice of tier symbols will allow a TSL grammar to enforce the observed pattern.

More generally, any omnivorous agreement system which makes use of a three-level scale to determine the surface form of an agreement marker cannot be captured by a TSL grammar. The heart of the issue is that the middle level of such a scale is sometimes, but not always, relevant; neither always projecting it to the tier, nor always omitting it from the tier, will suffice. According to the typology reported by Murugesan et al. (2025), such patterns are relatively rare even among omnivorous agreement systems, although the same 1>2>3 person scale as Mundari has also been observed in Alutor (Chukotkan). Omnivorous person agreement more commonly uses a two-level scale, such as in Kichean, Chuckchi (Chukotkan), and Eastern Armenian, where first and second person are incomparable ({1,2}>3); these patterns will therefore pose no problem for TSL grammars. In the domain of omnivorous number agreement, Murugesan et al. (2025) note that two-level scales are again more common (e.g. PL > SG in Georgian, Kichean, Mara (Tibeto-Burman), or SG > PL in Ketama Berber), but TSL-incompatible three-level scales have also been observed, for example in Hayu (Sino-Tibetan) and Onondaga (Northern Iroquoian). In fact, Mundari itself also uses a three-level SG > PL > DU number scale in object agreement when animacy and person values are held constant.

5 Differences between Uyghur backness harmony and omnivorous agreement

5.1 Multiple "domains" in a single c-string

In the formal language that Mayer and Major (2018) argued is not TSL, as reviewed in §3 above, each string corresponds to one inflected Uyghur word, consisting of a stem plus an appropriate suffix. We have suggested that the way in which that suffix's surface form depends on its phonological context is analogous to the way in which the object agreement marker in Mundari depends on its syntactic context. While the analysis in §4 restricted attention to monoclausal sentences, which contain at most one object agreement marker, c-strings representing full Mundari sentences may contain multiple clauses; and a complete formal grammar for these c-strings will therefore allow strings that contain multiple object agreement markers, each with a surface form that is dependent on the DPs in the corresponding clause. So one difference between a grammar for Uyghur word forms and a grammar for Mundari c-strings will be that each string generated by the former consists of only one "harmony domain", whereas a string generated by the latter might consist of multiple distinct "agreement domains" concatenated. This difference is orthogonal, however, to our key point which is that the patterns observed inside each domain are analogous, and beyond the capacity of TSL grammars.

5.2 Unboundedness of DPs in an agreement domain

A more significant potential objection to our main claim concerns the question of how many DPs can occur in a single agreement domain (as opposed to the question of how many agreement domains can occur in a single string). In (20) we characterized the pattern in a way that allows for unboundedly many D heads in the domain that an object marker's surface form depends on, and one might question this given that all the example sentences in §4 included just two DPs in the relevant domain (a direct object and an indirect object). If the number of DPs in a domain were limited to two, then the pattern would be TSL: it would suffice to project *all* D heads onto the agreement tier, since there would be a limit on how far away the highest-ranking D head could be from the object marker. This would be crucially different from the situation in Uyghur, where it is natural to assume that unboundedly many dorsal consonants might intervene between a critical harmonizing vowel and the suffix (recall (13)).

The formulation we gave in (20) is based on an assumption that, while it may indeed be impossible to construct an acceptable Mundari sentence with more than two DPs in a single object-agreement domain, this reflects a property of some part of the language's grammar that is distinct from its agreement mechanisms. We conjecture that the explanation of this boundedness of DPs will come from theories of thematic structure or argument licensing, and that it would therefore be redundant to build into our theories of agreement a limit on how many DPs might need to be considered. This separation of concerns is consistent with the way the mechanisms responsible for agreement patterns are described in the syntactic literature (e.g. Chomsky, 2001; Béjar, 2003): in such theories, a "probe" typically searches its c-command domain for an appropriate "goal" in a manner that does not involve keeping track of how much structure has been searched so far.

Having said that, even if one rejects our assumed separation of the number of DPs from the mechanisms of agreement, there is still something significant about the Mundari object agreement pattern. The overwhelming majority of TSL accounts of linguistic patterns (in both phonology and syntax) fall within the particularly restrictive TSL-2 class, i.e. they are stated in terms of length-two substrings on the relevant tier (see e.g. Hanson, 2025, pp.87–88; Jardine and Heinz, 2016, p.88). If the number of relevant DPs for Mundari object agreement were limited to two, then a TSL grammar could achieve the desired result by projecting all D heads onto the tier, but this would be a TSL-3 grammar, since the k-factors would include both objects and the agreement marker. The strong non-TSL claim would hold only for Uyghur vowel harmony, but the weaker (and still linguistically significant) non-TSL-2 claim would still hold in a unified way for Uyghur vowel harmony and Mundari object agreement.

6 Omnivorous agreement is O-TSL

6.1 O-TSL

Mayer and Major (2018, p.73) point out that Uyghur backness harmony can be formulated as a strictly local constraint on a tier if the projection function, rather than simply erasing certain symbols, is an output-sensitive strictly local map (Chandlee, 2014). This puts the harmony pattern itself in the class of O-TSL languages. To sharpen the parallels between Uyghur backness harmony and omnivorous agreement, this section will show how O-TSL grammars can capture both patterns.

Formally, a (j,k)-O-TSL grammar is like a k-TSL grammar, except that the tier projection function can take into account the most recent $j-1$ symbols already projected to the tier when deciding whether to project a particular surface symbol.[3] So to define the tier projection function, we need to specify not simply a subset τ of Σ, but rather a set of (σ,t) pairs, called j-contexts, where $\sigma \in \Sigma$ and $t \in (\{\rtimes\} \cup \Sigma)^{j-1}$. The intended interpretation of a j-context (σ,t) is that an occurrence of σ in the surface string is to be projected onto the tier iff the tier constructed so far ends in t.

As Mayer and Major point out, the Uyghur harmony pattern becomes strictly local on the relevant tier representations if we specify that harmonizing consonants are projected only if the most recent symbol on the tier is not a harmonizing vowel. This ensures that the suffix does not end up adjacent to an irrelevant intervening dorsal consonant in cases like [rak-**ta**] and [mæʃq-**tæ**]. The (2,2)-O-TSL grammar that achieves this is as follows.

$$\begin{aligned} C &= \{(\sigma,t) \mid \sigma \in \{\mathrm{V_f},\mathrm{V_b},\mathrm{S_f},\mathrm{S_b}\}, t \in \Sigma\} \\ &\quad \cup \{(\sigma,t) \mid \sigma \in \{\mathrm{C_f},\mathrm{C_b}\}, t \in \Sigma - \{\mathrm{V_f},\mathrm{V_b}\}\} \\ G &= \{\mathrm{V_fS_b}, \mathrm{V_bS_f}, \mathrm{C_fS_b}, \mathrm{C_bS_f}\} \end{aligned}$$

An exactly analogous strategy works for the omnivorous person agreement pattern that we described above: we project $\mathrm{D_2}$ only if the most recent symbol on the tier is not $\mathrm{D_1}$. This prevents a $\mathrm{D_2}$ from ending up adjacent on the tier to the agreement marker in a scenario where a $\mathrm{D_1}$ is present.

$$\begin{aligned} C &= \{(\sigma,t) \mid \sigma \in \{\mathrm{D_1},\mathrm{Ap_A},\mathrm{Ap_B},\mathrm{Ap_C}\}, t \in \Sigma\} \\ &\quad \cup \{(\mathrm{D_2},t) \mid t \in \Sigma - \{\mathrm{D_1}\}\} \\ G &= \{\mathrm{D_1Ap_B},\ \mathrm{D_1Ap_C},\ \mathrm{D_2Ap_A},\ \mathrm{D_2Ap_C}\} \end{aligned}$$

6.2 TSL, O-TSL, and IO-TSL

This analysis of omnivorous agreement is significant given the conjecture in Hanson (2025) that syntactic agreement will fall within TSL. However, as mentioned above in §2.3, Graf and Shafiei (2019) used tier-based constraints on c-strings to model syntactic phenomena other than agreement, namely reflexive- and NPI-licensing, and implement this with IO-TSL grammars (Graf and Mayer, 2018): these extend the TSL projection function not only with sensitivity to symbols already output to the tier, but also with sensitivity to local symbols in the input. The syntactic phenomena that Graf and Shafiei address, however, seem to also be within the bounds of the simpler TSL system. As far as

[3] A standard k-TSL grammar is therefore a special case, corresponding to a $(1,k)$-O-TSL grammar.

we are aware then, the omnivorous agreement pattern discussed above presents the most significant challenge to the idea that syntactic constraints are TSL on c-strings, as well as a clear challenge to the more specific idea that this holds for agreement.

7 Conclusion

We are not the first to point out that phonology and syntax– or even vowel harmony and syntactic agreement– seem to make use of similar underlying computational machinery (Nevins, 2010, 2011). Indeed, following the recent surge of work within subregular phonology and syntax, Graf (2022) proposes the *cognitive parallelism hypothesis* that distinct language modules– such as phonology and syntax– have the same formal complexity. What he leaves as an open question, though, is what exactly that level of complexity is.

We've shown here that phonology and syntax each have exceptions, in Uyghur and Mundari respectively, to the claim that each respective module is TSL-2, and that these patterns are exceptional in the same way. In terms of formal complexity, this distinctive exceptionality seems to put each of these patterns just beyond TSL, into O-TSL. But regardless of what exact level of complexity phonology and syntax may share, if there is indeed a similarity between the mechanisms deployed in syntax and phonology, then the formal similarities observed here are the exact kinds we expect to arise.

References

Alëna Aksënova, Thomas Graf, and Sedigheh Moradi. 2016. Morphotactics as tier-based strictly local dependencies. In *Proceedings of the 14th sigmorphon workshop on computational research in phonetics, phonology, and morphology*, pages 121–130.

Susana Béjar. 2003. *Phi-syntax: A theory of agreement*. University of Toronto Toronto, ON.

Jane Chandlee. 2014. *Strictly local phonological processes*. University of Delaware.

Noam Chomsky. 2001. Derivation by phase. In Michael J. Kenstowicz, editor, *Ken Hale: A Life in language*. MIT Press, Cambridge, MA.

Aniello De Santo and Thomas Graf. 2019. Structure sensitive tier projection: Applications and formal properties. In *International conference on formal grammar*, pages 35–50. Springer.

Paul D Fallon. 1993. Liquid dissimilation in georgian. In *Eastern States Conference on Linguistics*, volume 10, pages 105–116.

Thomas Graf. 2022. Subregular linguistics: bridging theoretical linguistics and formal grammar. *Theoretical Linguistics*, 48(3-4):145–184.

Thomas Graf and Connor Mayer. 2018. Sanskrit n-Retroflexion is Input-Output Tier-Based Strictly Local. In *Proceedings of the Fifteenth Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 151–160, Brussels, Belgium. Association for Computational Linguistics.

Thomas Graf and Nazila Shafiei. 2019. C-Command Dependencies as TSL String Constraints. In *Proceedings of the Society for Computation in Linguistics (SCiL) 2019*, pages 205–215.

Kenneth Hanson. 2025. Tier-based strict locality and the typology of agreement. *Journal of Language Modelling*, 13(1):43–97.

Jeffrey Heinz. 2018. The computational nature of phonological generalizations. In Larry M. Hyman and Frans Plank, editors, *Phonological Typology*, pages 126–195. De Gruyter.

Jeffrey Heinz, Chetan Rawal, and Herbert G Tanner. 2011. Tier-based strictly local constraints for phonology. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human language technologies*, pages 58–64.

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

Douglas C. Johnson. 1972. *Formal Aspects of Phonological Description*. Monographs in Linguistic Analysis. Mouton, The Hague.

Ronald M Kaplan and Martin Kay. 1994. Regular models of phonological rule systems. *Computational linguistics*, 20(3):331–378.

Connor Mayer and Travis Major. 2018. A challenge for tier-based strict locality from uyghur backness harmony. In *International conference on formal grammar*, pages 62–83. Springer.

Gurujegan Murugesan, Imke Driemel, and Andrew Murphy. 2025. Omnivorous person, number and gender in Mundari: A cyclic Agree analysis. *Natural Language & Linguistic Theory*.

Andrew Nevins. 2010. *Locality in vowel harmony*, volume 55. Mit Press.

Andrew Nevins. 2011. Multiple agree with clitics: Person complementarity vs. omnivorous number. *Natural language & linguistic theory*, 29(4):939–971.