



Burness, Phillip & McMullin, Kevin & Chandlee, Jane. 2021. Long-distance phonological processes as tier-based strictly local functions. *Glossa: a journal of general linguistics* 6(1): 99, pp. 1–37. DOI: <https://doi.org/10.16995/glossa.5780>



Open Library of Humanities

Long-distance phonological processes as tier-based strictly local functions

Phillip Burness, University of Ottawa, CA, pburness@uottawa.ca

Kevin McMullin, University of Ottawa, CA

Jane Chandlee, Haverford College, US

Whether we analyze phonological processes using a system of rules or constraints, the resulting map from underlying representations to surface pronunciations can be characterized as a function. Viewing processes as mathematical objects in this way allows us to study properties of phonology that hold no matter how it is implemented. Work in this vein has found that a majority of phonological processes only consider information within a finite window, placing them in the highly restrictive class of Strictly Local (SL) functions (Chandlee 2014; Chandlee et al. 2014; 2015). Long-distance phonological processes, however, lie outside the capabilities of the SL functions since they consider information that can be arbitrarily distant. The more powerful class of subsequential functions has been offered as a potential characterization of long-distance phonology (Heinz & Lai 2013; Luo 2017; Payne 2017), but we argue that an intermediate class offers a more natural model. Specifically, by incorporating an autosegmental tier (e.g., Goldsmith 1976) into the structure of an SL function, the non-local information crucial for applying long-distance processes can be rendered local. In addition to assessing the typological coverage of these Tier-based Strictly Local functions (Burness & McMullin 2019; Hao & Andersson 2019; Hao & Bowers 2019), we show that they fail to generate two pathological behaviours (minimum distance requirements and modulo counting) that can be accomplished with a subsequential function. We therefore conclude that tier-based computation is a better characterization of long-distance phonology than subsequential computation.

Glossa: a journal of general linguistics is a peer-reviewed open access journal published by the Open Library of Humanities. © 2021 The Author(s). This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC-BY 4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited. See <http://creativecommons.org/licenses/by/4.0/>.

OPEN ACCESS



1 Introduction

A typical assumption of generative phonological theory is that abstract underlying representations are mapped to surface representations by means of a phonological grammar. This grammar—whether it comprises a series of rules, requires optimization over a set of candidates and ranked constraints, or invokes some other mechanism entirely—can therefore be formally understood as a *function* that takes an underlying form as its input and produces a surface form as its output. With respect to segmental phonology, we can view this as a transformation from an input *string* (a sequence of underlying segments) to an output string (a sequence of segments produced in the surface form).

The interest in modeling phonological patterns as string-to-string transformations stems from the goal of identifying their individual computational properties as well as the overall computational complexity of the phonological grammar. A common assumption in this line of work is that local and long-distance phenomena have different computational properties (Heinz 2010): while local phonotactics and processes can be modeled with grammars that pay attention to contiguous substrings, long-distance patterns cannot. For example, the Samala language (also known as Ineseño Chumash; Applegate 1972) exhibits a regressive pattern of sibilant harmony, requiring all sibilants to agree for anteriority no matter how far apart they are in the word (with the value of $[\pm\text{ant}]$ being dictated by the rightmost sibilant). This can be seen in words such as $/\text{ha-s-xintila-wa}\text{f}/ \rightarrow [\text{ha}\text{f}\text{xintilawa}\text{f}]$ ‘his former gentile name’, and $/\text{k-su-k'ili-mekeken-f}\text{/} \rightarrow [\text{k}\text{f}\text{uk'ilimekeket}\text{f}]$ ‘I straighten myself up’ (Applegate 1972). To address the challenge of needing to look ahead (or back) across arbitrarily large distances, many theoretical frameworks use phonological ‘tiers’ (e.g., Goldsmith 1976; Clements 1980; Goldsmith 1990; Odden 1994; Clements & Hume 1995; Heinz et al. 2011; McMullin 2016), rendering long-distance dependencies local on the relevant tier, such as a tier of sibilants for Samala.

This paper demonstrates how various processes can be modelled with string-to-string functions that incorporate tier-based computation. In formal language theory, a tier is generally defined as a subset of the segment inventory, with non-members acting as if invisible to the model (Heinz et al. 2011). The use of such tiers has led to a better understanding of the computational nature of long-distance phonotactic dependencies (modelled as sets of well-formed strings; see McMullin 2016; McMullin & Hansson 2016; Aksënova & Deskmukh 2018; Lambert & Rogers 2020), and this paper explores the computational implications of extending the same idea to functions in order to characterize the types of long-distance *processes* that result in morpho-phonological alternations. Formal characteristics of these functions have already been fleshed out (Burness & McMullin 2019; Hao & Andersson 2019; Hao & Bowers 2019; Andersson et al. 2020), although to date there has not been a thorough investigation of how well tier-based functions approximate the attested typology of non-local phonological processes. Our endeavour to fill this gap reveals typological predictions that are supported cross-linguistically and offers insights into the computational nature of long-distance phonology.

The remainder of this paper is organized as follows. Section 2 summarizes relevant work on the computational complexity of local and long-distance phonotactic patterns as well as work that models local phonological processes using the class of Strictly Local functions. Section 3 discusses a previous computational model of long-distance processes, the subsequential functions, showing that this class readily generates two behaviours that we consider pathological. Section 4 shows how the class of Tier-based Strictly Local functions can model basic long-distance patterns, and how the class excludes the pathologies raised for the subsequential functions in the previous section. Section 5 demonstrates the wide-ranging capabilities of the TSL functions, including their ability to model important aspects of attested long-distance rules such as segmental transparency and segmental blocking. Section 6 then considers attested behaviours that lie outside the reach of tier-based functions, suggesting ways in which these limitations might be overcome. Finally, Section 7 concludes.

2 Background and context

2.1 Formal language theory and phonotactics

To better understand the limits of possible language patterns, early work in computational linguistics modeled patterns as sets of well-formed strings and classified patterns according to the machinery necessary for deciding whether a given string is well-formed, giving rise to what is known as the Chomsky Hierarchy of formal languages (Chomsky 1956). At the bottom of the Chomsky hierarchy are the *finite languages*, which can be modelled as finite sets of strings. Anything in the set is considered good and included in the language, whereas anything not in the set is considered bad and excluded from the language. The original Chomsky hierarchy distinguishes four classes above the finite languages. From least to most powerful, these are the regular languages, the context-free languages, the context-sensitive languages, and the recursively enumerable languages. Phonology is argued to be at most regular in terms of complexity (Johnson 1972; Kaplan & Kay 1994) but does not, in most cases, require the full capabilities of regular languages and functions. Accordingly, much work has divided the regular region into a hierarchy of *subregular* formal language classes (e.g., McNaughton & Papert 1971; Rogers & Pullum 2011; Rogers et al. 2013). Importantly, while formal languages are well-suited to static phonotactic patterns, they cannot directly model dynamic processes, and so a parallel hierarchy of functions and relations (i.e., collections of well-formed *pairs* of strings) is being explored as well. The remainder of this section provides an overview of two important classes of subregular languages: the Strictly Local (SL) languages that form the basis of the SL functions developed by Chandlee (2014) and Chandlee et al. (2014; 2015),¹ as well as the Tier-based Strictly Local (TSL) languages developed by Heinz et al. (2011) that form the basis of the TSL functions explored in this paper.

¹ Chandlee (2014) first defined SL functions in the context of phonology, but local functions themselves have precedent in the literature (see Berstel 1982; Vaysse 1986; Lind & Marcus 1995; Sakarovitch 2009).

SL languages can be defined positively or negatively, although negative definitions may be more intuitive to phonologists working in constraint-based frameworks like Optimality Theory (OT: Prince & Smolensky 2004) since they function like inviolable markedness constraints. From the negative perspective, SL languages ban particular contiguous sequences of up to a fixed length, excluding any string that contains one or more illegal contiguous sequences. When all the banned sequences are of length k or shorter, we say that the language is Strictly k -Local (SL_k). Many phonotactic patterns can be described as an SL_k language. For example, in Japanese the alveolar fricative [s] cannot occur immediately before the high front vowel [i]. Not only is the sequence *[si] absent from the language’s lexicon, but there are alternations that occur specifically to avoid it. One such alternation occurs when a verb root ending in /s/ is put into the past tense, as illustrated by the comparison of /hos-u/ → [hosu] ‘dry-NPST’ and /hos-ita/ → [hoçita] ‘dry-PST’. This aspect of Japanese phonotactics can be captured by an SL_2 language that bans the sequence *[si], preventing strings such as *[hosita] from surfacing. The TSL languages augment the SL languages with intuitions from autosegmental phonology — originally developed by Goldsmith (1976) for the analysis of tonal phenomena — in order to describe non-local phonotactic restrictions. While the specifics of autosegmental analyses can differ, the approach can be distilled down to the following. First, a phonological form consists of multiple connected levels of representation called ‘tiers’. Second, a given phonological element (e.g., a feature or segment) can be present on one tier but absent on another. Third, material that is non-local on one tier may be local on another. Finally, only material that is local *on the relevant tier* matters when enforcing a pattern. This intuition that non-local configurations can be rendered local with the right representational choices forms the foundation of the TSL languages and functions at the center of this paper.

Despite the wide-ranging interest in tiers and their utility, it is only recently that their expressive power has been explored in the context of formal language theory (Heinz et al. 2011; Jardine & Heinz 2016; McMullin 2016; McMullin & Hansson 2016; Jardine & McMullin 2017; Lambert & Rogers 2020). From this perspective, a tier is a subset of the alphabet,² and a Tier-based Strictly k -Local (TSL_k) language is a language that is SL_k once all non-tier elements have been erased. The result of erasing the non-tier elements from a string is often called the *projection* of that string onto the given tier. More sophisticated methods of projection that consider surrounding local material in addition to segment identity have been explored (Graf & Mayer 2018; Mayer & Major 2018; De Santo & Graf 2019), although we limit ourselves in this paper to the maximally simple means of projection where a segment’s identity is the only determining factor. A TSL language can then be said to ban a given string if that string’s projection onto the specified tier contains any impermissible contiguous sequences.

² In formal language theory the *alphabet* is the set of symbols that strings can be built from. In the context of phonology the alphabet is typically the segment inventory, but can also include boundary symbols, structural bracketing, etc.

Consider the pattern of liquid harmony in Bukusu, where sequences of [...r...l...] are not permitted (Odden 1994; Hansson 2010a). Data for this pattern are provided in (1), which show that the applicative suffix appears as [-ira] if the nearest leftward liquid is [r], else it appears as [-ila]. The pattern can be expressed as a TSL₂ language banning the sequence *[rl] on the liquid tier $T = \{r, l\}$. Forms like *[rumila] are excluded, since projecting them to the liquid tier yields the illegal string *[rl]. Note that such a language model does not actually describe the apparent *process* of liquid harmony, just its *result*.

- (1) Bukusu liquid harmony (Odden 1994)
- a. xam-ila ‘milk-APPL’
 - b. lim-ila ‘cultivate-APPL’
 - c. kar-ira ‘twist-APPL’
 - d. rum-ira ‘send-APPL’

A good question to ask at this point is how an autosegmental analysis of liquid harmony compares to the TSL analysis just presented. Both make a distinction between a representational level where the whole phonological string is present and a representational level where only the liquid consonants are present. Also in both approaches, the prohibition against a contiguous sequence of [-lateral][+lateral] is enforced on the less-inclusive level of representation containing just liquids. Where the approaches differ concerns the way these levels of representation are defined. An autosegmental approach might postulate, for example, that there is a tier containing all and only the instances of the feature [\pm lateral]. Assuming that only liquid consonants can be specified for the feature [\pm lateral], the entities present on this [\pm lateral] tier will then always correspond to an [l] or [r] in the full phonological string. For its part, the TSL approach arbitrarily stipulates that there is a representational level including all and only the instances of [l] and [r]. While inspired by the autosegmental approach, the TSL perspective does not follow it to the letter, allowing us to gain insight into the formal underpinnings of the structures of long-distance phonology with minimal theoretical commitments. Although the tier of a TSL language can in principle be an arbitrary set of segments, TSL analyses are by no means incompatible with approaches that define tiers according to perceptual similarity (as in the Agreement by Correspondence framework; see e.g. Rose & Walker 2004 and Hansson 2010a) or according to feature geometry (Clements & Hume 1995). Many of the tiers used in the computational literature (and in our analyses below) happen to form a natural class relative to typical feature theories, and one can of course decide to limit themselves only to similarity-driven or feature-definable tiers, but there is formally no such requirement. However one chooses to restrict the types of tiers that are possible, that choice will not change the fact that TSL languages over those tiers still model the intuition that a non-local pattern can hold locally at an appropriate representational level. Despite their name, the TSL languages do not themselves provide a theory

of possible tiers; rather, they provide a model of the effect that a tier (motivated or not) has upon the assessment of locality.

2.2 Strictly Local functions

To model phonological *processes* we need to shift our focus from *languages* to *functions* that map input strings to output strings. While it is not as well understood as the subregular hierarchy of languages, the analogous subregular hierarchy of functions has seen many developments in recent years (Heinz 2018). Particularly important to this paper are the Input Strictly Local (ISL) functions and the Output Strictly Local (OSL) functions, which adapt the notion of strict locality to input-output maps (Chandlee 2014; Chandlee & Heinz 2018; Chandlee et al. 2018). Various equivalent characterizations of these functions exist. This paper will make use of their automata-theoretic characterization, representing processes with *finite-state transducers* (FSTs) that meet particular criteria. For the full details of this characterization see Chandlee (2014), Chandlee et al. (2014), and Chandlee et al. (2015).

A (one-way) FST produces an output string incrementally by reading an input string one element at a time in a single direction. Such a machine consists of a finite set of states (which can be thought of as a primitive sort of memory) and a finite set of transitions between these states (which include the machine’s instructions for what to write at each step). The machine begins in a designated initial state, and traverses a path through the state space by following transitions in response to the input that it reads. **Figure 1** presents a visual diagram of an FST. States are represented with circles, with the initial state marked with an unlabeled incoming arrow. Transitions are represented with labelled arrows between states; a label ‘a:b’ is an instruction to take that transition when reading ‘a’ from the input and write ‘b’ to the output. The transducer in **Figure 1** operates over the input alphabet {a, b}, transforming all odd-numbered positions to ‘a’ and all even-numbered positions to ‘b’. For example it maps /bab/ to [aba] and maps /aabbabbb/ to [abababab].

ISL and OSL FSTs have particular conditions on the state set and transitions. States in an ISL_k FST correspond to strings of up to $k - 1$ input symbols, whereas states in an OSL_k FST correspond to strings of up to $k - 1$ output symbols. Transitions in an ISL_k FST always go to the




Figure 1: A simple finite-state transducer.

state matching the most recent $k - 1$ symbols that have been read from the input string, whereas the transitions in an OSL_k FST always go to the state that matches the most recent $k - 1$ symbols that have been written to the output string. Additionally, ISL and OSL FSTs are required to be *deterministic*, meaning that there is one transition per input element per state.

As a concrete example of these restrictions, consider the transducers in **Figure 2**, both of which compute the rule of post-nasal voicing in the Puyu Pungo dialect of Quechua (Orr 1962; Rice 1993). The transducers operate over simplified alphabets where ‘V’ represents a vowel, ‘N’ represents a nasal consonant, ‘T’ represents a voiceless obstruent, and ‘D’ represents a voiced obstruent. Data for the pattern are provided in (2).

- (2) Post-nasal voicing in Puyu Pungo Quechua (Orr 1962; Rice 1993)
- [sinik-pa] ‘porcupine-GEN’ [kam-ba] ‘you-GEN’
 - [sat̪a-pi] ‘jungle-LOC’ [hatum-bi] ‘big.one-LOC’
 - [wasi-ta] ‘house-OBJ’ [wakin-da] ‘others-OBJ’

In ISL and OSL transducers, the initial state corresponds to the empty string, which we write as λ . For maximum clarity, we write the labels of input-oriented FSTs between forward slashes (e.g., /X/) and the states of output-oriented FSTs between square brackets (e.g., [X]). Structurally, the crucial difference between the ISL and OSL analysis concerns the ‘T:D’ transitions that leave the ‘N’ states. These are depicted with dashed lines in the figures. This transition ends up in the /T/ state of the ISL transducer, but ends up in the [D] state of the OSL transducer. Despite the difference in structure, the FSTs compute the same function in this particular case (i.e., they are extensionally equivalent).

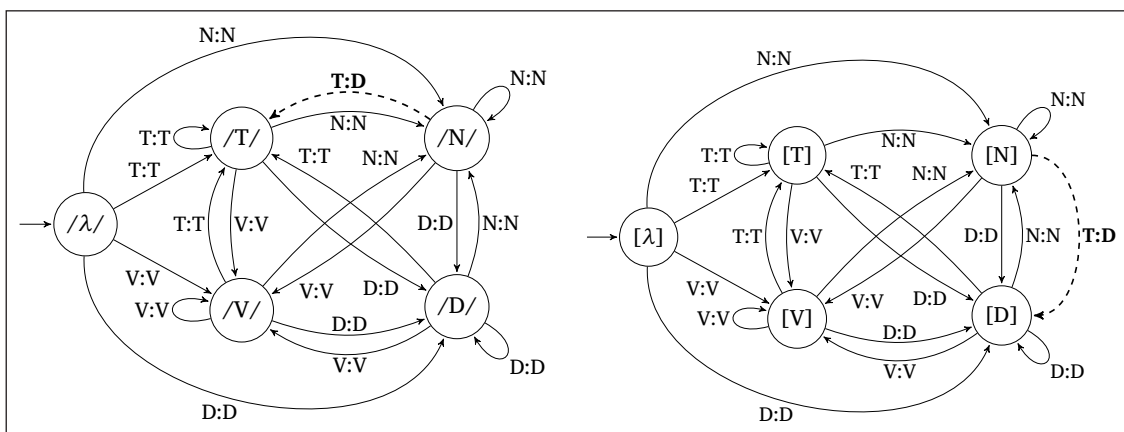


Figure 2: An ISL₂ FST (left) and an OSL₂ FST (right) computing post-nasal voicing. ISL state labels are enclosed in slashes and OSL labels are enclosed in square brackets. The transition that differs across the transducers is dashed and in bold.

2.3 Long-distance processes are not Strictly Local

As mentioned at the outset, long-distance processes cannot be modeled by tracking only recent input/output as the SL functions do. Consider, for example, the sibilant harmony found in Aari and shown in (3). The perfective suffix /-s/, which stays faithful in (3a), surfaces instead as [-ʃ] when preceded (at any distance) by a lamino-postalveolar sibilant, as shown in (3b–f). The process is ‘asymmetric’ in that underlying /ʃ/ does not analogously become [s] when preceded by an alveolar fricative (Hansson 2010a: p. 357).

- (3) Progressive long-distance sibilant harmony in Aari (Hayward 1990)
- | | | | |
|----|---------------|----------|---------------------|
| a. | /baʔ-s-e/ | baʔse | ‘bring-PERF-3SG’ |
| b. | /ʔuʃ-s-it/ | ʔuʃʃit | ‘cook-PERF-1SG’ |
| c. | /ʃʔa:q-s-it/ | ʃʔa:qʃit | ‘swear-PERF-1SG’ |
| d. | /ʒaʔ-s-it/ | ʒaʔʃit | ‘arrive-PERF-1SG’ |
| e. | /ʃed-er-s-it/ | ʃederʃit | ‘see-PASS-PERF-1SG’ |
| f. | /ʒa:g-er-s-e/ | ʒa:gerʃe | ‘sew-PASS-PERF-3SG’ |

Figure 3 demonstrates what goes wrong when we attempt to model such a process as OSL for $k = 2$. For readability we use ‘s’ to represent [+ant] sibilants, ‘ʃ’ for [-ant] sibilants, and ‘&’ as a placeholder for all non-sibilant segments. According to the generalization, an input /s/ may be mapped in two different ways. Compare (3a) /baʔ-s-e/ → [baʔse] to (3e) /ʃed-er-s-it/ → [ʃederʃit]. In both cases, the FST will be in the state [&] when the underlying /s/ is read since the last element produced was a non-sibilant ([ʔ] in 3a and [r] in 3e). The output forms show, however, that these two /s/’s must be outputted differently, meaning the transducer would need both an ‘s:s’ and an ‘s:ʃ’ transition out of this state. These two transitions are depicted in Figure 3 with

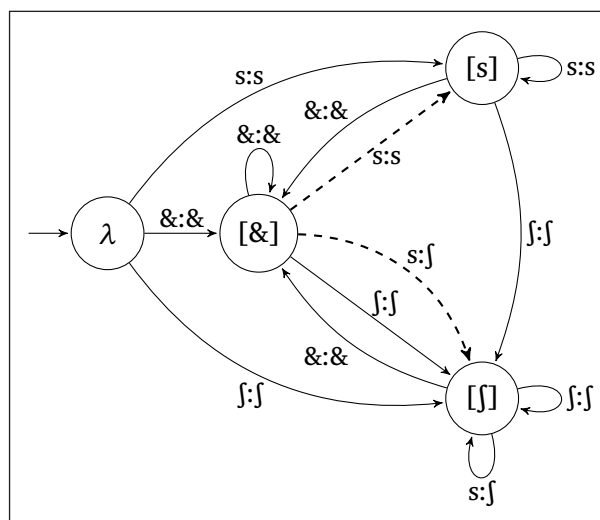


Figure 3: Failed attempt to model Aari sibilant harmony with a OSL_2 transducer.

dashed lines. The requirement of determinism prevents the FST from including both of these transitions.

The problem, of course, is due to the long-distance nature of the process. The crucial distinction between (3a) and (3e) is that the former does not contain a prior sibilant but the latter does. That distinction cannot be made by an FST that only remembers the previous output segment. And since long-distance processes apply regardless of the number of intervening segments, simply increasing the locality window will not suffice. As such, we must pursue other means for characterizing long-distance phonological processes.

3 Subsequential functions

To capture long-distance phenomena, we will require a class of functions that has greater generative capacity than the SL functions, but how high must we climb in computational power? Early work from Johnson (1972) and Kaplan & Kay (1994) established that phonological processes fit into the class of *regular relations*. These are the relational analogue to the regular languages within the Chomsky hierarchy (Chomsky 1956) that was described briefly in Section 2.1, and they can be defined as relations computable by some FST (among other converging definitions). Specifically, the above-mentioned work proved that every re-write rule of the shape ‘ $A \rightarrow B/C_D$ ’ (where A, B, C, and D are regular languages) can be computed by an FST provided it does not reapply to its own structural change. Additionally, the regular relations are closed under composition (i.e., if $f(w)$ and $g(w)$ are regular then $f(g(w))$ is regular) so any ordered series of such re-write rules also computes a regular relation.

Using regular relations as a model of long-distance phonology is attractive for a number of reasons aside from the fact that the regular region subsumes the attested range of processes. For instance, the regular region requires relatively low computational power. Moreover, it excludes a type of process known as *majority rules* harmony which is widely regarded as being pathological (Lombardi 1999; Baković 2000; Heinz & Lai 2013; Finley 2017). In a majority rules process, determining the outcome requires comparing the number of occurrences of one segment type versus another within the same underlying form. For example, in majority rules backness vowel harmony, a stem would surface containing only front vowels if front vowels are more numerous than back vowels in the underlying form, but would surface containing only back vowels if back vowels are more numerous than front vowels in the underlying form. Human phonology does not seem to track relative frequency in this way, so the fact that the regular hypothesis excludes majority rules behaviour is a welcome result.

More recently, long-distance phonological processes have been argued to instantiate *subsequential functions*. These are functions that can be computed by a *deterministic* FST, making them a proper subclass of the regular relations (every subsequential process is regular, but not every regular process is subsequential). As support for using the subsequential functions to

model long-distance phonology, several types of unidirectional processes have been shown to be subsequential, namely vowel harmony (Gainor et al. 2012; Heinz & Lai 2013), consonant harmony (Luo 2017), and consonant dissimilation (Payne 2017). A further argument in favour of the subsequential hypothesis parallels the discussion of majority rules harmony above; a pathological process which Wilson (2003; 2006) calls *sour grapes* vowel harmony—adapting a term from Padgett (1995)—is regular but not subsequential (Heinz & Lai 2013). In a *sour grapes* process, a harmony-triggering segment will spread its harmonic feature to other segments only if it can affect all possible targets; if any one of these targets would resist being changed, the spreading never gets initiated. In other words, the potential harmony trigger first looks ahead to see whether it can affect everything, and preemptively gives up if it sees any antagonistic elements down the line. Human phonology does not seem to have such unbounded lookahead capabilities, and so the fact that the subsequential hypothesis excludes *sour grapes* behaviour is a welcome result (but see Section 6.3 for some counterexamples in the form of *unbounded circumambience*). That being said, we contend that the subsequential functions do not accurately represent the functioning of long-distance processes, since two pathological behaviours can result from the freedom afforded to what states may represent in a subsequential FST.

The first pathology is what we call *minimum distance requirements*, where a trigger affects a target only if it is some set distance or further from the target. An example of such a system would be *strictly beyond-transvocalic dissimilation*, where two segments dissimilate only when they are separated by *at least* a vowel and a consonant; no dissimilation occurs across just a vowel (i.e., transvocalic contexts). McMullin (2016) and McMullin & Hansson (2019) ran a series of artificial grammar learning experiments, finding that strictly beyond-transvocalic patterns of liquid harmony and dissimilation are not reliably acquired in an experimental setting. Even when given unambiguous training data, participants tended to either infer a more typical unbounded pattern or else not infer any pattern. We accordingly propose that phonological processes should be prevented from containing minimum distance requirements like those in a beyond-transvocalic system. The subsequential hypothesis fails in this regard, since states in a subsequential FST can represent arbitrary configurations such as “*n* or more syllables away from the most recent *x*”, and so minimum distance requirements are readily enforceable by a subsequential FST.

The FST in **Figure 4** provides a concrete demonstration of how beyond-transvocalic dissimilation is subsequential. For ease of interpretability we consider only inputs with perfect consonant-vowel alternation, although this assumption does not crucially affect the discussion. The rule computed by the transducer is expressed in (4), where V represents a vowel, C represents a consonant other than [l], and $\{CV\}^+$ represents a string of one or more CV syllables. To see how this transducer computes first-last harmony, consider the input /lolokol/. Starting in state 0, we produce [l] upon reading word-initial /l/ bringing us to state 1. Next we read /o/ which produces [o] and brings us to state 2. On the following step, we read another /l/ but this is not separated




Figure 4: A subsequential FST that computes beyond-transvocalic liquid dissimilation.

from the previous one by a beyond-transvocalic distance and is produced faithfully as [l], moving us back to state 1. The next input elements are /o/, /k/ and /o/ which together produce [oko] and move us to state 4. After that comes another /l/ but this time we *are* separated from the previous /l/ by a beyond-transvocalic distance. Accordingly, we dissimilate and produce [r], so the full mapping is then /lolokol/ → [lolokor]. Once state 4 is reached, the machine can only cycle between states 4 and 5, so if any more instances of /l/ are read in this situation, all of these will dissimilate to become [r].

(4) $l \rightarrow r / IV\{CV\}^+ _$

The second pathology is what we call *modulo counting*, where sequences of segments get collected into groups of equal size. An example of such a system would be transvocalic harmony that is sensitive to the parity of potential harmony targets (see also McMullin (2016) and McMullin & Hansson (2016) who discuss how this pathology falls out of certain constraint-based models). In a transvocalic sibilant harmony system, sibilants agree in anteriority specifically if they are separated by at most a single vowel. When there is a sequence of multiple such sibilants, we expect the anteriority value of the first sibilant to be inherited by all sibilants in the chain. Essentially, the value gets passed from each sibilant to the next, skipping over one vowel on each pass. In the parity-sensitive system, however, the chain of sibilants will organize itself into non-overlapping pairs, with each first member of a pair passing its value to the second member of the pair. Furthermore, there is no interaction between sibilants in different pairs. As a result, the system permits disharmonic sequences so long as the disharmony straddles a ‘pair boundary’. For example, we would expect that /saʃasasa/ maps to [sasasasa] but it will instead map to [s_ia_si_ra_sa_sa], where pair membership is marked with subscripts. Unlike for the minimum distance pathology, there are to our knowledge no experiments that investigate the learnability

of modulo counting patterns like parity-sensitive harmony. Nonetheless, we believe that it is safe to consider modulo counting pathological for segmental phenomena.³

The FST in **Figure 5** provides a concrete demonstration of how parity-sensitive harmony is subsequential. To improve the transducer's readability, we consider a segment inventory of only vowels and sibilant fricatives and consider only inputs with perfect consonant-vowel alternation. Neither of these assumptions crucially affect the discussion. The rule computed by the transducer is expressed in (5), where V represents a vowel, S represents a sibilant and $\{ \}^{2n}$ represents an even number of instances (including 0) of some sequence. To see how this transducer computes parity-sensitive transvocalic harmony, consider the input /sofofoso/. Reading the first sibilant-vowel syllable produces [so] bringing us to state 2. The next sibilant-vowel pair syllable must harmonize since it forms a pair with the previous one, so we write [so] for /fo/. Doing so returns us to state 0, which means that the following syllable starts a new pair so /fo/ is free to surface faithfully bringing us to state 4. Finally, the fourth syllable forms a pair with the previous one, so we write [fo] for /so/. This gives us the full mapping /sofofoso/ → [sosofofo].

$$(5) \quad [+\text{sibilant}] \longrightarrow [\alpha_{\text{anterior}}] / \{SV\}^{2n} \left[\begin{array}{c} +\text{sibilant} \\ \alpha_{\text{anterior}} \end{array} \right] \text{V} \text{ —}$$

The above pathologies suggest that subsequential functions are not fully emblematic of possible phonological computations. Interestingly, the weaker class of Tier-based Strictly Local functions




Figure 5: A subsequential FST that computes parity-sensitive transvocalic harmony.

³ A reviewer points out that the parity-sensitive example looks like a pattern applying only within a prosodic foot. The example is thus not unequivocally pathological, but it remains that subsequential functions can perform modulo counting for any positive integer, and can do so without any regard for prosodic boundaries.

(defined below in Section 4) avoids these two pathologies. Since TSL functions can model most attested non-local processes (as we discuss in detail in Section 5) without generating the two pathologies just described, we will propose that the TSL functions act as a better characterization of long-distance phonology than the subsequential functions.

4 Tier-based Strictly Local functions

To address the inability of SL functions to model long-distance processes, recent formal work (Burness & McMullin 2019; Hao & Andersson 2019; Hao & Bowers 2019; Andersson et al. 2020) has aimed to instead extend the TSL languages to functions, although to date there has not been a thorough investigation of their empirical coverage. Intuitively, TSL functions are simply SL functions that operate with reference to a tier.⁴ As we did for the SL functions above, we will illustrate TSL functions using FSTs and describe their properties primarily from an informal perspective. For various formal definitions (automata-theoretic and otherwise) of TSL functions, readers are directed to the works cited above.

In FST representations of ITSL_k and OTSL_k functions,⁵ each state is a record of the most recent $k - 1$ tier segments that were read from the input string or written to the output string, respectively. We illustrate with a simplified version of long-distance sibilant harmony, similar to that of Aari shown above in (3). Specifically, an input string $/\dots f \dots s \dots /$ maps to $[\dots f \dots f \dots]$ and $/\dots s \dots f \dots /$ maps to $[\dots s \dots s \dots]$. The left-hand FST in **Figure 6** models this process as an ITSL_2

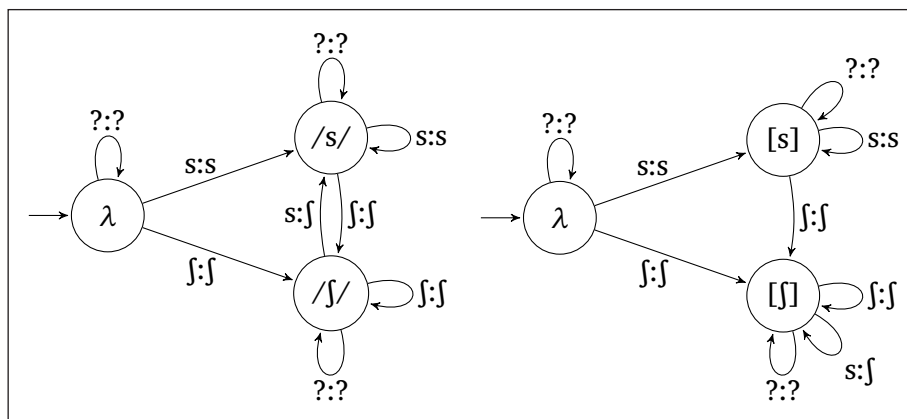


Figure 6: An ITSL_2 FST (left) and an OTSL_2 FST (right) computing patterns of sibilant harmony.

⁴ A word of caution: the image of a TSL function (i.e., the language formed by collecting all of its possible outputs) is not necessarily a TSL language, just as the image of an SL function is not necessarily an SL language (Chandlee 2014). The Karajá pattern analyzed in Section 5.4 is a TSL function that does not produce a TSL language.

⁵ These functions are not to be confused with the ITSL and OTSL *languages* of Graf & Mayer (2018), Mayer & Major (2018), and De Santo & Graf (2019). The acronyms ITSL and OTSL consistently refer to the function classes throughout the rest of this article.

function, whereas the right-hand FST is OTSL_2 . Throughout, we generally use ‘?’ to represent non-tier segments. Non-tier segments are mapped faithfully in the two functions here, so their transitions are loops and the FSTs can only change state upon reading/writing a tier segment.⁶

As with SL FSTs, the crucial difference between ITSL and OTSL FSTs is whether the transitions go to the state corresponding to their input or output side. This difference can be observed in the figures for a transition like ‘s:ʃ’, which goes to state /s/ in the ITSL_2 FST and state [ʃ] in the OTSL_2 FST. This differing behavior of the two FSTs is exemplified in (6) and (7), which show the respective paths and outputs through both FSTs for the hypothetical input /ʃasas/.

(6) Path through the ITSL FST (Figure 6 left side) for the input /ʃasas/

Input:	ʃ	a	s	a	s						
Path:	λ	\Rightarrow	/ʃ/	\Rightarrow	/ʃ/	\Rightarrow	/s/	\Rightarrow	/s/	\Rightarrow	/s/
Output:	ʃ	a	ʃ	a	s						

(7) Path through the OTSL FST (Figure 6 right side) for the input /ʃasas/

Input:	ʃ	a	s	a	s						
Path:	λ	\Rightarrow	[ʃ]	\Rightarrow	[ʃ]	\Rightarrow	[ʃ]	\Rightarrow	[ʃ]	\Rightarrow	[ʃ]
Output:	ʃ	a	ʃ	a	ʃ						

Importantly, the two pathologies described in Section 3 are excluded from both the ITSL and OTSL classes precisely because local tier-based computation makes them impossible to generate. To impose a minimum distance requirement for the triggering of a process using a TSL function, the trigger and intervening material must be on the provided tier else they could not affect the state of the transducer. For example, if we were to try computing strictly beyond transvocalic dissimilation with an ITSL function, the tier would need to be the entire input alphabet. There is, however, a finite maximum of material that can factor into the state label. Consequently, the trigger will eventually be pushed out of this window and forgotten. Continuing with the ITSL attempt at beyond transvocalic dissimilation, suppose that $k = 5$ and we are on the last steps of reading /kolo-l/, /loko-l/, and /lokoko-l/. In the first, the state label will be /kolo/ and so we do not apply dissimilation since the preceding /l/ is not far enough away. In the second, the state label will be /loko/ and so we do apply dissimilation since the preceding /l/ is far enough away. In the third, the state label will be /koko/ and so we do not apply dissimilation since we (erroneously) do not see any preceding /l/. TSL functions are subject to a *maximum distance limit* if they wish to impose a minimum distance requirement, a constraint that does not apply to subsequential functions.

⁶ This is not enforced by the formal definitions of ITSL/OTSL, which only dictate what states are present and where transitions go. Formally speaking, non-tier input segments do not necessarily need to map faithfully, nor do non-tier output segments need to be the result of a faithful mapping.

As for the other pathology, a TSL function can perform some modulo counting if it has a sufficiently high maximum length for its state labels. For example, let us try modelling parity-sensitive sibilant harmony using an ITSL₄ transducer whose tier is $T = \{s, \int\}$. Suppose that we are on the last steps of reading /soso- \int /, /sososo- \int / and /sosososo- \int /. In the first, the state label /ss/ has an even number of sibilants and so we do not harmonize with the most recent one. In the second, the state label /sss/ has an odd number of sibilants and so we harmonize with the most recent one. In the third, the state label /sss/ leads us to (falsely) believe that there is an odd number of preceding sibilants, and so we harmonize. State labels in a TSL transducer will eventually become saturated and cannot subsequently become unsaturated, so the modulo counting can only be enforced up to a particular multiple, a constraint which does not apply to subsequential functions. By virtue of being able to model attested long-distance phonological processes (see Section 5) while ruling out two arguably non-phonological behaviours, the TSL functions offer a more natural characterization of long-distance phonology than the subsequential functions.

Finally, it is worth commenting at this point on how the TSL approach to patterns like sibilant harmony compares to an autosegmental analysis, in which instances of the feature [\pm anterior] are dominated by [+sibilant] nodes, and spreading is a matter of reassociation. For their part, TSL functions are agnostic about why {s, \int } form a tier to the exclusion of other segments and why the tier segments affect certain segments the way they do. While the operations being computed by the transducer can be interpreted as the result of a spreading rule, the transducer does not directly represent the reasons behind tier membership and tier influence in any way. It simply represents the map from one string to another. It is important to remember though that the overarching goal of computationally-oriented work in phonology is to characterize the mathematical properties of the maps, and not necessarily their underlying causes. Regardless of one's preferred theoretical mechanism for accomplishing something like sibilant harmony, the fact remains that it is a TSL function. The next sections will show how most other long-distance phonological maps also fall within either the class of ITSL functions, the class of OTSL functions, or both.

5 Expressivity of Tier-based Strictly Local functions

5.1 Transparency

A well-studied aspect of long-distance phonological processes is the non-participation or *transparency* of material between the trigger and target (see, e.g., Archangeli & Pulleyblank 2007; Gafos & Dye 2011; Rose & Walker 2011; Finley 2017). Consider the Turkish pattern of backness harmony, where suffix vowels take on the backness value of the rightmost vowel in the base (Clements & Sezer 1982; Nevins 2010). Consonants do not participate and are 'skipped over' when determining appropriate suffix allomorphs.⁷ Data for the pattern are provided in (8), where

⁷ Note that some analyses argue that the consonants are not actually skipped over (e.g., Waterson 1956).

the plural suffix alternates between front [-ler] and back [-lar] and the genitive suffix alternates between front [-in] and back [-un].

- (8) Turkish backness vowel harmony (Nevins 2010: 28–29)
- [ip-ler] ‘rope-PL’ [ip-ler-in] ‘rope-PL-GEN’
 - [el-ler] ‘hand-PL’ [el-ler-in] ‘hand-PL-GEN’
 - [pul-lar] ‘stamp-PL’ [pul-lar-un] ‘stamp-PL-GEN’
 - [son-lar] ‘end-PL’ [son-lar-un] ‘end-PL-GEN’

An OTSL₂ transducer for this process is depicted in **Figure 7**. The tier for this function is all and only the output vowels, though for reasons of space, the transducer collapses all the front vowel states and the all back vowel states together. The transitions are also collapsed such that a transition labelled ‘[αbk]:[βbk]’ means that an input vowel specified as [αback] gets mapped to its [βback] equivalent (with all other features unchanged). Turkish allows disharmonic bases, and some suffixes do not alternate, instead remaining faithful and starting a new domain of harmony, so we assume that harmony can only affect segments with missing backness values, marking under-specification as [0bk].⁸ The transducer models the intuition that a vowel underspecified for [±back] will take on the [±back] value of the closest specified vowel to its left. Consonant transparency is captured in this transducer by the fact that reading and producing a consonant never causes a change of state.

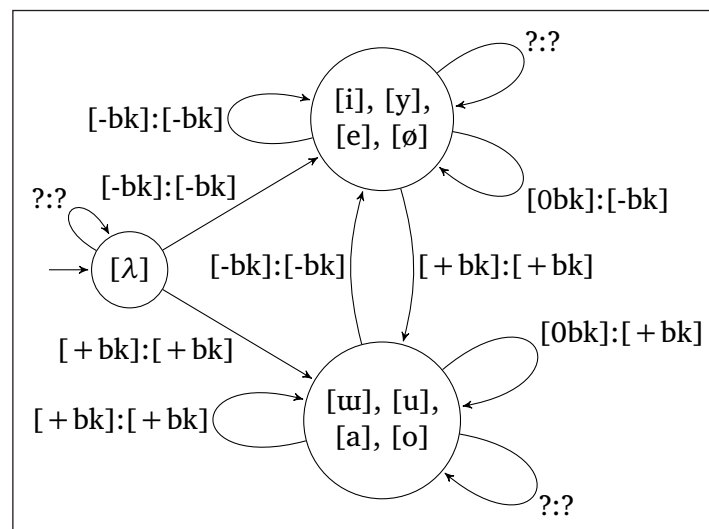


Figure 7: An OTSL₂ transducer that produces Turkish backness harmony.

⁸ We are able to account for such morpheme-specific information in this way since a function’s input and output alphabets are not required to be the same.

The transparency of consonantal segments in vowel harmony is a matter of some debate. Some researchers adopt a “local spreading” approach where the harmonic feature spreads onto and through consonants via coarticulation, (e.g., Ní Chiosáin & Padgett 2001; Jurgec 2011) rather than skipping them. Transparent segments would then be those for which the harmonizing feature does not carry any perceptual or contrastive force. Support for this idea comes from research on the reportedly transparent vowels of Hungarian. Like Turkish, Hungarian has backness harmony, although the non-low front vowels [i] and [e] are unaffected by the harmony and neither trigger nor block it (Gafos & Dye 2011). Phonetic studies reveal, however, that instances of [i] and [e] between two back vowels are produced slightly back, albeit to a sub-phonemic and likely imperceptible level (Benus et al. 2004; Gafos & Benus 2006; Benus & Gafos 2007).

Some other harmony processes, however, are not fully amenable to a spreading analysis. For example, the usual [d] of the perfective suffix /-idi/ in Yaka becomes nasal if there is a nasal obstruent in the root (Hyman 1995; Walker 2000; Rose & Walker 2004; 2011; Archangeli & Pulleyblank 2007; Jurgec 2011). The root in (9a) contains no nasal consonant, so the suffix surfaces faithfully. Compare this to the roots in (9b–c) which *do* have a nasal consonant. No matter whether this nasal consonant is in the adjacent syllable (9b) or several syllables away (9c), it forces the suffix to harmonize and surface as [-ini]. This harmony for the feature [\pm nasal] occurs across vowels and voiceless obstruents without nasalizing them to any extent (Rose & Walker 2004). We must instead allow nasality to skip over them, which we can do by excluding them from an appropriate OTSL function’s tier.

- (9) Yaka nasal consonant harmony (Hyman 1995)
- a. /tsub-idi/ [tsub-idi] ‘roam-PERF’
 - b. /tsum-idi/ [tsum-ini] ‘sew-PERF’
 - c. /nutuk-idi/ [nutuk-ini] ‘lean on-PERF’

Transparency in the Yaka pattern is no more or less difficult to describe using a TSL function than transparency in Turkish or Hungarian vowel harmony. This is because the TSL class places no limits on which collections of segments can act as a tier. As we mentioned earlier in Section 4, the TSL functions are agnostic as to *why* the tier takes a particular shape. We believe that it is good practice to try and define tiers relative to a feature theory, but also believe that the possibility of arbitrary tiers is advantageous since it is not always possible to derive a segment’s tier membership from external factors (e.g., the shape of the phoneme inventory). For example, the low vowel /a/ undergoes tongue root harmony in Kinande (Cole & Kisseberth 1994; Gick et al. 2006) while it resists and blocks tongue root harmony in Pulaar (Archangeli & Pulleyblank 1994), even though there are no relevant differences between the languages that can explain this discrepancy (Rose & Walker 2011).

5.2 Parasitic harmony

Another behaviour that has received plenty of attention in the phonological literature is so-called *parasitic* harmony, where the source and target interact on one dimension only when they agree along another. The most prominent cases of parasitic harmony involve rounding harmony depending on vowel height. Take for instance the Kachin dialect of Khakass. Suffix high vowels in this language become round if the nearest leftward vowel is also high (Korn 1969; Kaun 1995). The data in (10) show a [+high] suffix vowel harmonizing with [+high] vowels (a and b) but failing to harmonize with [−high] vowels (c and d). As the additional data in (11) show, a [−high] suffix vowel never harmonizes, even if the preceding vowel is also [−high].

(10) Kachin Khakass rounding harmony in [+high] suffixes (Korn 1969: 102–103)

- a. /kyn-ni/ [kyn-ny] ‘day-ACC’
- b. /kuʃ-tun/ [kuʃ-tun] ‘of the bird’
- c. /ød-ir/ [ød-ir] ‘to kill’
- d. /ok-tun/ [ok-tun] ‘of the arrow’

(11) Kachin Khakass lack of harmony in [−high] suffixes (Korn 1969: 102–103)

- a. /kyn-ge/ [kyn-ge] ‘to the day’
- b. /kuzuk-ta/ [kuzuk-ta] ‘in the nut’
- c. /tʃør-gen/ [tʃør-gen] ‘who went’
- d. /pol-za/ [pol-za] ‘if he is’

Let us consider how the OTSL₂ transducer in **Figure 8** models the pattern. The tier includes all vowels, so the machine is always in the state corresponding to the most recently produced vowel.

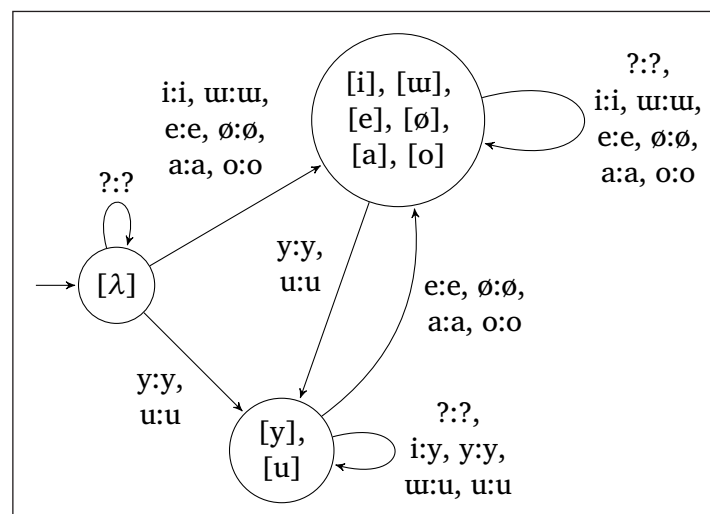


Figure 8: An OTSL₂ transducer that produces Kachin Khakass parasitic rounding harmony.

For readability, the two [+high, +round] states are collapsed together, as are the states for the six remaining vowels. While in a [+high, +round] state (i.e., the [y] or [u] state), reading /i/ will produce [y] and reading /u/ will produce [u]. The same vowels are produced faithfully while in any other state. By coordinating the work of the tier and the work of the transitions, we can ensure that (i) only high vowels are ever affected by vowel harmony and (ii) they are only so affected when the preceding vowel is also high.

5.3 Blocking

Rounding harmony in Khalkha Mongolian is very similar to rounding harmony in Kachin Khakass but also exhibits *blocking* effects. As in Kachin Khakass, rounding harmony is parasitic on height, in this case requiring the target and trigger to both be non-high (Svantesson et al. 2005; Nevins 2010; Gafos & Dye 2011). The harmony causes alternations in a variety of suffixes, such as the comitative, data for which are shown in (12).

- (12) Khalkha Mongolian rounding harmony (Nevins 2010: p. 139)
- a. [nar-tai] ‘sun-COM’
 - b. [ɔd-tɔi] ‘star-COM’

Note that there is a simultaneous process of advanced tongue root (ATR) harmony; our analysis abstracts away from this additional process. The high front vowel /i/ is transparent to rounding harmony (Nevins 2010): it does not become rounded when preceded by a round vowel of any height, but also does not prevent rounding from reaching a following [-high] vowel. This is shown in (13) using words with the accusative and reflexive suffixes.

- (13) Khalkha Mongolian transparency of /i/ (Svantesson et al. 2005: p. 50)
- a. [pɔ:r-ig-o] ‘kidney-ACC-REFL’
 - b. [xɔ:lʒ-ig-ɔ] ‘food-ACC-REFL’
 - c. [mʊ:r-ig-a] ‘cat-ACC-REFL’
 - d. [su:lʒ-ig-e] ‘tail-ACC-REFL’

The other high vowels /u/ and /ʊ/, however, *do* prevent rounding from reaching a following [-high] vowel, as can be seen from the words in (14) where the causative suffix comes between a root and the perfective suffix.

- (14) Khalkha Mongolian blocking by /u, ʊ/ (Nevins 2010: p. 137)
- a. [tor-o:d] ‘be.born-PERF’
 - b. [tor-u:l-e:d] ‘be.born-CAUS-PERF’
 - c. [ɔr-ɔ:d] ‘enter-PERF’
 - d. [ɔr-ʊ:l-a:d] ‘enter-CAUS-PERF’

Blocking is handled straightforwardly in TSL functions by including blockers on the tier. For example, the Khalkha pattern can be modelled by the OTSL₂ transducer in **Figure 9** whose tier contains all and only the rounded vowels. To keep the transducer legible we abstract away from distinctions along the ATR and length dimensions (e.g., the distinction between [o] and [ɔ:]) and ignore the parallel process of ATR harmony. The vowel [o] is included on the tier because it triggers rounding harmony and the vowel [u] is included because it blocks rounding harmony. The vowels [i] and [e] (as well as consonants) are excluded from the tier because they neither trigger nor block rounding harmony. Reading the mid vowel /e/ while in the [o] state will produce [o], no matter how many consonants or instances of [i] have been produced since the [o] that put us in the [o] state. If at any point we produce a blocker [u], we move to the [u] state, out of which all transitions are faithful. Rounding harmony can only apply again if a new instance of [o] is produced.

Segmental blocking effects are not limited to vowel harmony. Although rare, a few cases of long-distance consonantal phenomena exhibit blocking effects. These include sibilant harmony in Slovenian (Jurgec 2011), sibilant harmony in Kinyarwanda (Walker & Mpiranya 2006; Walker et al. 2008), sibilant harmony in Imdlawn Tashlhiyt (Elmedlaoui 1995; Hansson 2010b) and liquid dissimilation in Georgian (Fallon 1993; Odden 1994). Blocking can also be morphologically or lexically driven. We mentioned earlier in Section 5.1 that some suffix vowels in Turkish are invariant, like the vowel in the nominalizer suffix /-gen/ which is always [-back] (Nevins 2010: pp. 33–34). The /e/ in /-gen/ acts as a blocker in the sense that it stops the spread of [+back], but differs from the cases of blocking above by actively causing all following vowels to take on [-back]. As we showed, one way of ensuring this kind of behaviour in a TSL function is to make a distinction between input vowels unspecified for some feature and input vowels pre-specified for the same feature; the former will take on a value for the empty feature based on the most recent tier element (i.e., based on the current state), but the latter will maintain the specification that they already have.




Figure 9: An OTSL₂ transducer that computes rounding harmony with blocking.

5.4 Icy targets

The last type of segment we discuss from the TSL perspective is what Jurgec (2011) calls an *icy target*. These segments act much like a blocker in that they prevent harmony beyond themselves, but differ from blockers in that they undergo the harmony that they block. An instance of iciness comes from the Macro-Jê language Karajá, which has a regressive ATR harmony process that spreads [+ATR] (Ribeiro 2003; Rose & Walker 2011; Walker 2012). Underlyingly [-ATR] high vowels are the icy targets: an underlying [+high, +ATR] vowel will trigger harmony, while an underlying [+high, -ATR] vowel will undergo harmony but then block its spread (Ribeiro 2003). In other words, *derived* [+high, +ATR] vowels cause harmony to halt. Compare the examples in (15a–b), in which harmony spreads throughout the word, with those in (15c–d), which contain icy targets.

- (15) Karajá ATR harmony (Ribeiro 2003)
- a. /brɔrɛ-dĩ/ [brɔrɛ-ni] ‘deer-similar.to’
 - b. /bɛdɔ-dĩ/ [bedo-ni] ‘filhote-similar.to’
 - c. /krɔbi-dĩ/ [krɔbi-ni] ‘monkey-similar.to’
 - d. /kɔɖu-dĩ/ [kɔɖu-ni] ‘turtle-similar.to’
 - e. /kɔlukɔ-dĩ/ [kɔlukɔ-ni] ‘cajá (tree species)-similar.to’

Our analysis abstracts away from the local nasal displacement that derives [-ni] from /-dĩ/; it is challenging to implement multiple processes with a single TSL function, but see Section 6.1 for extensions to the TSL class designed with this issue in mind. The pattern is sensitive to the [±ATR] value that each vowel has *in the input string* and so an ITSL₂ function is required. To keep the transducer in **Figure 10** legible, we use only the following subset of the Karaja vowel inventory: /i/, /I/, /u/, /U/, /e/, /ɛ/, /o/, and /ɔ/. Limiting the inventory does not significantly affect the analysis. Because the process is regressive, the input string is read from right to left. The FST’s tier includes underlyingly [+ATR] vowels alongside /I/, /U/. Crucially, /ɛ/ and /ɔ/ must be excluded from the tier so that reading them while in the [+ATR] state does not move us to the [-ATR] state. While it might feel odd to let non-tier elements be mapped unfaithfully, doing so does not violate the definition of an ITSL function in any way.⁹

Compare this to the transitions originating in the [+ATR] state for the inputs /I/ and /U/ which *do* belong on the tier and thus *do* lead to the [-ATR] state. As a result, [+ATR] harmony can affect and pass over/through an underlyingly [-ATR] mid vowel, whereas it can affect but not pass over/through an underlyingly [-ATR] high vowel.

It is worth noting that, despite using an input-oriented function, we are generating what looks like iterative output-oriented application. Such “iterative” application is possible with an

⁹ The definition of an ITSL function requires only (i) that the tier be a subset of the input alphabet and (ii) that the corresponding transducer’s current state represent the most recent tier elements. This definition places no requirements on the output edges of transitions, and indeed, the output alphabet could in principle be entirely disjoint from the input alphabet, making all transitions “unfaithful”.




Figure 10: An ITSL_2 transducer that produces Karajá ATR harmony.

ITSL function when non-icy targets are off the input tier and do not prevent the trigger from remaining in memory (i.e., remaining in the state label). For another example, while we present Turkish vowel harmony as OTSL in Section 5.1, it can equally be computed by an ITSL function provided that affix vowels are underspecified and provided that underspecified vowels are off the tier. In fact, the ITSL transducer looks exactly the same as the OTSL transducer in Figure 7 aside from changing “[]” to “//” in the state labels.

5.5 Double, (non-)initial, and (non-)final triggers

All of the long-distance patterns discussed so far share one striking thing in common: they are all TSL_k for $k = 2$. This is not surprising when we consider the typical interpretation of long-distance processes. Long-distance harmony is usually construed as the feature value of one segment spreading in some direction to affect other eligible segments, so it makes sense that a window of length 2 is sufficient in the tier-based context, since it lets us remember the most recent “relevant” element. Long-distance dissimilation can also generally be reduced to the influence of the most recent relevant element. The prominence of $k = 2$ is also interesting with regards to learnability. Burness & McMullin (2019) showed that while any OTSL_k function can be learned efficiently from positive data when the tier is known in advance, only the OTSL_2 functions can be efficiently learned by their algorithm from positive data when the tier is *not* known in advance.¹⁰ Nonetheless, there are several behaviours that require a value of $k > 2$ for various reasons.

¹⁰ This of course does not mean that learning tiers is impossible when $k > 2$, only that no method currently exists for doing so efficiently. Whether or not such tiers can be efficiently learned is a topic of ongoing research.

A first example is the rounding harmony in one variety of Oroqen. Some sources analyze the harmony as only being triggered by a sequence of two rounded vowels (Zhang et al. 1989; Li 1996; Zhang 1996; Zhang & Drescher 1996; Walker 2001; Walker 2014). The data in (16) show the definite article suffix becoming round when preceded by two round vowels, but not when preceded by a single round vowel (even if this vowel is long). Note that the language also contains a process of ATR harmony, hence the alternation between [ɔ] and [o]. An alternative analysis of the same data comes from Drescher & Nevins (2017), who analyze the harmony as triggered specifically by non-initial vowels. They analyze the process in this way because the Russian loanword [kinɔ] ‘film’ causes rounding harmony in the definite object suffix /-wa/ giving us [kinɔ-wɔ] instead of *[kinɔ-wa] as predicted by the double-trigger analysis (Drescher & Nevins 2017). There also exists a plural suffix /-nɔr/ for kinship terms that is always round, and this suffix causes rounding harmony regardless of the rounding in preceding vowels, giving us [ətʃəxə-nɔr-wɔ-t] ‘paternal.uncle-PL-DEF-ACC’ instead of *[ətʃəxə-nɔr-wa-t], for example (Drescher & Nevins 2017). This again is expected under the non-initial trigger analysis, but not under the double trigger analysis.

- (16) Oroqen rounding harmony (Walker 2001)
- a. /ɔɔ-wa/ [ɔɔ-wɔ] ‘fish-DEF.OBJ’
 - b. /tʃoŋko-wa/ [tʃoŋko-wo] ‘window-DEF.OBJ’
 - c. /mɔ:-wa/ [mɔ:-wa] ‘tree-DEF.OBJ’

Interestingly, both the double trigger analysis and non-initial trigger analysis can be implemented as an OTSL function if we set k to be greater than or equal to 3. The states in a TSL_k transducer are labelled with sequences of *up to* length $k - 1$. Accordingly, until enough tier elements have been encountered, the transducer can be in a state labelled with fewer than $k - 1$ segments, but it will cycle through states labelled with $k - 1$ elements once that number is reached. For example, an OTSL_3 transducer with the alphabet {V, C} and the tier {V} will enter the state ‘V’ upon producing its first vowel, and will not move to the state ‘VV’ until a second vowel is produced. After that point it will remain in the state ‘VV’, never returning to the state ‘V’. This is illustrated in **Figure 11**. When implementing the double-trigger analysis as an OTSL_3 function, we would ensure that rounding harmony is only enforced out of a state whose label consists of two round




Figure 11: Example of state labels shorter than the maximum.

vowels. Similarly, when implementing the non-initial trigger analysis as an OTSL_3 function we would ensure that rounding harmony is only enforced when the state label contains two vowels and the rightmost of these is round.

Similar to how the right-hand vowel of a two-vowel state label in a progressive TSL_3 function is necessarily non-initial, a one-vowel state would necessarily reflect the first vowel when all vowels are on the tier. With $k = 3$, then, we can also model non-iterative harmony that is triggered specifically by the first or last vowel in a word. In Megisti Greek, regressive non-iterative backness and rounding harmony is triggered only by the last vowel in a word (van Oostendorp & Revithiadou 2005; McCollum & Kavitskaya 2018). Harmony preceding a final back vowel is shown in (17a) and harmony before a final front vowel is shown in (17b). The example in (17c) shows that the harmony does not extend past the penultimate vowel.

- (17) Megisti Greek non-iterative backness harmony Triggered only by final vowels
(van Oostendorp & Revithiadou 2005)
- a. /sits-a/ [sutsa] ‘fig.tree-NOM.F’
 - b. /filak-s-e/ [filekse] ‘guard-3SG.PST’
 - c. /anofli/ [anefli] ‘lintel’

Trigger status is purely positional in this pattern and has nothing to do with stress placement (McCollum & Kavitskaya 2018), so we cannot rely on a dichotomy between stressed and unstressed vowels to derive the restriction on triggers. A TSL analysis of this pattern (or its mirror image) needs a tier of all vowels and $k = 3$. A state labelled with a single vowel will then correspond to having seen only the first vowel in a left-to-right machine and correspond to having only seen the last vowel in a right-to-left machine. Restricting harmonic transitions so that they can only come from these states will limit triggers to the desired word edge.

Values of k higher than 2 can thus account for certain behaviours not captured by the TSL_2 functions, such as double triggering and some instances of positional triggering. Not all instances of positional triggering, however, can be analyzed in the way outlined above. For example, suffixes in the Eastern Meadow dialect of Mari harmonize with the initial vowel in backness across all other vowels (Vaysman 2009; Walker 2011).¹¹ This particular process cannot be modelled by inflating k since it is iterative and there is no way to exclude non-initial vowels from the tier by identity alone, so the initial vowel will eventually be pushed out of the memory window (i.e., the state label of the FST). We conjecture that positional triggering of this kind requires something like the structure-sensitive tier projection methods explored for formal languages by Graf & Mayer (2018), Mayer & Major (2018), and De Santo & Graf (2019), which consider surrounding material in addition to segment identity.

¹¹ We thank an anonymous reviewer for bringing this pattern to our attention.

6 Limitations of Tier-based Strictly Local functions

As shown in the previous sections, the TSL functions can model a range of non-local processes, including those with blocking. They are not, however, sufficient to capture all such processes that are attested in the literature. This section presents three phenomena that TSL functions cannot model—multiple simultaneous processes, bidirectional application, and two-sided contexts—along with extensions to the TSL functions that might circumvent the relevant inadequacies.

6.1 Multiple processes

The Tamashek dialect of Tuareg contains a process of long-distance regressive sibilant harmony and a process of long-distance regressive labial dissimilation (Heath 2005; McMullin 2016). The sibilant harmony can be seen in words where the causative prefix /s-/ is followed non-locally by another sibilant, whereupon it takes that other sibilant's values for anteriority, voicing, and pharyngealization as shown by the data in (18) from Heath (2005: p. 442). Note that the language has considerable vowel allophony, and we write 'V' where Heath (2005) does not provide the surface vowel quality.

- (18) Tamashek sibilant harmony (Heath 2005; McMullin 2016)
- causative /s-/
 -s-VŋŋV- 'cook'
 -s-VsVfVr- 'treat (patient)'
 -s^ʕ-Vs^ʕuhV- 'strengthen'
 -ʃ-VluʃV- 'clean sand from'
 -z-VjVzzV- 'scrutinize'

The labial dissimilation can be seen in words where a prefix /m/ (such as in the mediopassive) is followed non-locally by a labial consonant other than /w/, whereupon the prefix /m/ will dissimilate to [n]. Data for this process is shown in (19) from Heath (2005: p. 472).

- (19) Labial dissimilation: mediopassive /m-/
 -m-VrtVj- 'become mixed'
 -n-VkmVm- 'be squeezed'

That both processes can occur simultaneously is demonstrated by the word in (20) from Heath (2005: p. 462), which contains the causative and mediopassive prefixes together.

- (20) Both prefixes/processes
 α-z^ʕ-ən:-ət-əlməz^ʕ 'spitting saliva'

The combination of sibilant harmony and labial dissimilation cannot, however, be computed by a single TSL function. Consider what happens when we try with an OTSL₂ function whose tier

consists of all sibilants and labial consonants. Producing a sibilant will push the most recent labial consonant (if any) out of the $k - 1$ window, and producing a labial consonant will push the most recent sibilant (if any) out of the $k - 1$ window. At any given point, then, we can only know how to correctly map an input /s/ or only know how to correctly map an input /m/, but not both. Increasing k does not eliminate the issue, since any number of sibilants/labials can in principle occur between two labial/sibilant consonants.

A further difficulty for TSL functions posed by Tamashek Tuareg is that long-distance regressive labial dissimilation is overridden by a local process of regressive nasal place assimilation. An input /m/ fails to dissimilate specifically when it is immediately followed by an oral labial stop, in order to avoid a heterorganic cluster. This can be seen in a word like (21) from Heath (2005: p. 476) where the reciprocal prefix /Vm-/ comes immediately before a /b/.

- (21) Local blocking of dissimilation, reciprocal /Vm-/
 -æm-bæbba- ‘carried each other’

We are faced with a paradox if we attempt to model this interaction as a single TSL function. In order to know when an input labial needs to dissimilate we need to ignore everything that is not a labial consonant, but in order to know when an input labial needs to obey place assimilation we cannot ignore anything.

Whether such interactions between multiple local and/or long-distance processes are problematic for using TSL functions as a model of long-distance phonology depends on whether we are trying to study properties of individual processes/maps or properties of entire phonological systems. The latter perspective is explored by Chandlee & Heinz (2018) and Chandlee et al. (2018) in the context of formal language theory and automata theory. In an appendix to their article, Chandlee et al. (2018) prove that the ISL functions are not closed under composition (i.e., given two ISL functions $f(w)$ and $g(w)$, the single function equivalent to $f(g(w))$ is not guaranteed to be ISL),¹² and yet the main body of their article shows that certain opaque rule interactions such as counterfeeding and counterbleeding can nevertheless be modelled using a single ISL function, suggesting that overall phonological systems might generally obey strict locality. In a similar vein, preliminary work by Burness & McMullin (2020) explores how to combine multiple TSL functions into a single function based on insights from research by Aksënova & Deskmukh (2018) and McMullin et al. (2019) into phonotactic patterns that operate over multiple tiers. The particular subclass of multi-tiered functions that Burness & McMullin (2020) argue for lets each input element separately depend on its own set of tiers, provided that the set forms a strict superset-subset hierarchy. Further work is needed, however, to assess the degree to which

¹² Although Chandlee & Lindell (under review) conjecture that closure is guaranteed when neither of the functions being combined contains a “null cycle” where infinite deletion is possible.

tier-based strict locality is obeyed by phonological systems overall. On this note, a reviewer observes that the interaction between dissimilation and assimilation in Tamashek Tuareg is similar to a minimum distance requirement, suggesting that tier-based computation may avoid certain pathologies only at the level of an individual process.

6.2 Bidirectional application

One aspect of long-distance patterns is tacitly assumed above, namely that they tend to apply in a single direction. This seems to be the norm, but there are patterns that apply in both directions, typically from the root outwards. Take for instance vowel harmony in Degema, where prefixes *and* suffixes take on the $[\pm\text{ATR}]$ specification of the root (Archangeli & Pulleyblank 2007; Kari 2007). The data in (22) show both parts of a circumfix alternating to match the root's value for $[\pm\text{ATR}]$. The combination of regressive harmony to prefixes and progressive harmony to suffixes gives the impression of a harmony system that begins from the centre. TSL functions obligatorily start at either the left or right edge of the input, so even if we make a distinction between root and affix vowels by having the latter be unspecified for $[\pm\text{ATR}]$ in the input, we will only correctly treat either the prefixes or the suffixes.

- (22) Degema ATR harmony (Kari 2007)
- | | | | | |
|----|-------|-----------|--------------------------|-----------|
| a. | [ból] | 'hold' | [u-ból ¹ -ám] | 'holding' |
| b. | [gén] | 'look' | [u-gén ¹ -ám] | 'looking' |
| c. | [dúm] | 'create' | [o-dúm ¹ -ám] | 'creator' |
| d. | [hór] | 'sharpen' | [ɔ-hór ¹ -ám] | 'sharpen' |

A deceptively simple solution presents itself: why not apply the same TSL function once going from left to right, then a second time going from right to left? The first pass will fill in the missing $[\pm\text{ATR}]$ value on suffixes and the second pass will fill in the missing $[\pm\text{ATR}]$ values for prefixes. A similar idea of decomposing a single bidirectional process into an ordered pair of opposite-direction subsequential functions is how Heinz & Lai (2013) account for stem-controlled vowel harmony, though they caution against freely composing subsequential functions due to a theorem from Elgot & Mezei (1965). The theorem states that running a first subsequential transducer in one direction and then a second subsequential transducer in the reverse direction can describe any regular function. This happens because the first transducer can “annotate” the input string to give the opposite-direction transducer information about distant upcoming material. Heinz & Lai (2013) accordingly suggest that when we need to split a single process into two parts, the first transducer's output alphabet must be equal to its input alphabet (which is then the input alphabet to the second transducer), and must not be allowed to produce an output longer than the input. These restrictions prevent some of the annotation necessary for producing truly regular functions, and functions describable in this manner are known as the

weakly deterministic functions. Since we argue that the TSL functions are a more natural model of long-distance phonology than the subsequential functions, it would be interesting to see whether we can model bidirectional processes as a pair of opposite-direction TSL functions without the use of intermediate encoding.¹³ We leave a thorough investigation of this possibility for future research.

6.3 Two-sided contexts

ISL functions are able to model two-sided contexts since they can wait up to a finite number of steps before transforming an input element. This ability is not in general available to ITSL functions, so a challenge for the TSL models is that there do exist patterns that can be interpreted as being affected by non-local information in one direction as well as bounded information in the other. Consider the lowering harmony in C'Lela, where suffix high vowels become mid if they are preceded by a non-high vowel in the root (Dettweiler 2000; Pulleyblank 2002; Archangeli & Pulleyblank 2007; Michel 2009; Jurgec 2011). When there is more than one vowel following the triggering root vowel, the harmony only affects the very last one, which will always be in word-final position thanks to the syllabic structure of C'Lela (Jurgec 2011: p. 186). The data in (23) show that the class membership suffix /-i/ will lower to [e] when it is the only suffix following a non-high root vowel, but will be unaffected when it is followed by another suffix vowel.¹⁴

(23) C'Lela lowering harmony (Dettweiler 2000)

- | | | | |
|----|------------|------------|-----------------|
| a. | /zis-i/ | [zis-i] | 'long-CL' |
| b. | /rek-i/ | [rek-e] | 'small-CL' |
| c. | /zis-i-ni/ | [zis-i-ni] | 'long-CL-ADJM' |
| d. | /rek-i-ni/ | [rek-i-ne] | 'small-CL-ADJM' |

The basic facts of lowering are simple enough to account for with an ITSL₂ or OTSL₂ function: the tier consists of non-high vowels, and input high vowels will be output as non-high while in a non-high state. Accounting for the fact that targets are only lowered in word-final position, on the other hand, lies outside the powers of the TSL class. In order to confirm that an input vowel is word-final, we need to postpone its output until we read the next input element: if anything except the word boundary follows, it is not word-final. A paradox ensues: the lowering requires a tier with only non-high vowels, while determining whether the vowel is word-final requires a tier with everything. The multi-tiered functions of Burness & McMullin (2020) seem well-suited

¹³ A reviewer also makes the interesting observation that the tier projection of a TSL function can only delete a segment or reproduce it exactly, precluding the ability to add information.

¹⁴ Dettweiler (2000) was unable to determine the exact meaning of the suffix /-ni/, and so treats it as an adjective marker of uncertain status.

to resolving this paradox, since a local tier is by necessity a strict superset of any tier used to keep track of non-local information. Another possibility, as suggested by a reviewer, would be to appeal to structure-sensitive tier projection (Graf & Mayer 2018; Mayer & Major 2018; De Santo & Graf 2019)

Related to hybrid local/non-local cases like C’Lela lowering harmony is what Jardine (2016) calls *unbounded circumambience*, where the application of a process depends on potentially non-local information in both directions simultaneously. Take for instance Tutrugbu vowel harmony, as analyzed by McCollum & Essegbey (2018) and McCollum et al. (2020). Root vowels that are [+ATR] spread this value leftwards to prefix vowels, and when all prefix vowels are of the same height, the harmony will reach the left word edge (McCollum & Essegbey 2018; McCollum et al. 2020), as seen in (24a) for a string of [+high] prefixes and (24b) for a string of [–high] prefixes. Roots are underlined for clarity, and tone is not indicated to keep [ɪ] and [i] visually distinct.

(24) Conditional blocking by low vowels in Tutrugbu ATR harmony (McCollum et al. 2020)

- | | | | |
|----|-----------------------|-----------------------|---------------------|
| a. | /bu-ti- <u>e</u> / | [bu-ti- <u>e</u>] | ‘1PL-NEG-grow’ |
| b. | /ka-ba- <u>e</u> / | [ke-be- <u>e</u>] | ‘2SG-FUT-grow’ |
| c. | /ka-ti-ba- <u>e</u> / | [ke-ti-be- <u>e</u>] | ‘CL-NEG-FUT-grow’ |
| d. | /ɪ-ba-di- <u>wu</u> / | [ɪ-ba-di- <u>wu</u>] | ‘1SG-FUT-ITV-climb’ |

The behaviour of strings of prefixes with different heights depends on the height of the leftmost vowel. Harmony spreads to all prefix vowels when the leftmost vowel is [–high] as seen in (24c), but harmony is blocked by [–high] vowels when the leftmost vowel is [+high] as seen in (24d), although harmony still spreads up until that [–high] blocker (McCollum & Essegbey 2018; McCollum et al. 2020). The outcome of an input [–high] vowel can therefore depend simultaneously on non-local information to its left (the height of the leftmost vowel) and non-local information to its right (the ATR specification of the root), creating clear instances of unbounded circumambience. Jardine (2016) hypothesized that unbounded circumambient processes were not weakly deterministic according to Heinz & Lai’s (2013) definition, although O’Hara & Smith (2019), Smith & O’Hara (2019), and Lamont et al. (2019) show that loopholes in the definition can be exploited to model several unbounded circumambient patterns.¹⁵ We similarly conjecture that cases of unbounded circumambience like the one in Tutrugbu could be modeled as a pair of opposite-direction single-tiered or multi-tiered functions without the need for abstract intermediate encoding, although we leave to future research the task of determining the limits on such composition that would prevent the modelling of fully regular processes.

¹⁵ The alternative definition of weak determinism under development by Meinhardt et al. (2020) is an attempt to eliminate such loopholes.

7 Conclusion

Long-distance phonological processes are traditionally modelled using segmental tiers, and we have shown how such a tier can be incorporated into the structure of a Strictly Local (SL) function. The resulting Tier-based Strictly Local (TSL) functions are computational reflections of key insights from autosegmental phonology, readily capturing key behaviours exhibited by long-distance processes (such as the transparency of intervening material) that lay outside the modelling capabilities of the SL functions. TSL functions are computationally more powerful than SL functions, although they are less powerful than the subsequential functions, which were previously offered as a hypothesized upper bound on phonological complexity. Some pathological patterns can be characterized as a subsequential function but not as a TSL function, and so we argued that the TSL functions may be a better characterization of the computational mechanisms utilized by human phonology. That being said, the TSL functions have difficulty modelling two-sided contexts, bi-directional application, and the simultaneous application of multiple local and/or non-local dependencies. These difficulties can be alleviated through function composition and/or the addition of more tiers, although work is still ongoing to determine the limits that ought to be imposed on these powerful tools.

Abbreviations

1 = first person, 2 = second person, 3 = third person, ACC = accusative, ADJM = adjective marker, APPL = applicative, CL = classifier, DEF = definite, COM = comitative, F = feminine, FUT = future, GEN = genitive, GER = gerundive, ITV = itive, LOC = locative, M = masculine, NEG = negative, NMZR = nominalizer, NOM = nominative, NPST = non-past, OBJ = object, PERF = perfective, PL = plural, POSS = possessive, PST = past, REFL = reflexive, SG = singular

Acknowledgements

We would like to thank three anonymous reviewers for their detailed and insightful comments. This research was supported by the Social Sciences and Humanities Research Council of Canada.

Competing Interests

The authors have no competing interests to declare.

References

- Aksënova, Alëna & Deskmukh, Sanket. 2018. Formal restrictions on multiple tiers. In *Proceedings of the Society for Computation in Linguistics (SCiL) 2018*, 64–73.
- Andersson, Samuel & Dolatian, Hossep & Hao, Yiding. 2020. Computing vowel harmony: The generative capacity of search and copy. In Baek, Hyunah & Takahashi, Chikako & Yeung, Alex Hong-Lun (eds.), *Proceedings of the 2019 annual meeting on phonology*. DOI: <https://doi.org/10.3765/amp.v8i0.4752>
- Applegate, Richard B. 1972. *Ineseño Chumash grammar*. Berkeley. University of California Doctoral Dissertation.
- Archangeli, Diana & Pulleyblank, Douglas. 1994. *Grounded phonology*. Cambridge, MA: MIT Press.
- Archangeli, Diana & Pulleyblank, Douglas. 2007. Harmony. In de Lacy, Paul (ed.), *The Cambridge handbook of phonology*, 353–378. Cambridge University Press. DOI: <https://doi.org/10.1017/CBO9780511486371.016>
- Baković, Eric. 2000. *Harmony, dominance and control*. New Brunswick, NJ. Rutgers University Doctoral Dissertation.
- Benus, Stefan & Gafos, Adamantios. 2007. Articulatory characteristics of Hungarian “transparent” vowels. *Journal of Phonetics* 35. 271–300. DOI: <https://doi.org/10.1016/j.wocn.2006.11.002>
- Benus, Stefan & Gafos, Adamantios & Goldstein, Louis. 2004. Phonetics and phonology of transparent vowels in Hungarian. In *Proceedings of the 29th Annual Meeting of the Berkeley Linguistics Society*, 485–497.
- Berstel, Jean. 1982. Fonctions rationnelles et addition, actes de l'école de printemps de théorie des langages. *Laboratoire d'Informatique de Paris*, 177–183.

- Burness, Phillip & McMullin, Kevin. 2019. Efficient learning of output tier-based strictly 2-local functions. In *Proceedings of the 16th Meeting on the mathematics of language*, 78–90. Association for Computational Linguistics. DOI: <https://doi.org/10.18653/v1/W19-5707>
- Burness, Phillip & McMullin, Kevin. 2020. Multi-tiered strictly local functions. In *Proceedings of the 17th sigmorphon workshop on computational research in phonetics, phonology, and morphology*, 245–255. Association for Computational Linguistics. DOI: <https://doi.org/10.18653/v1/2020.sigmorphon-1.29>
- Chandlee, Jane. 2014. *Strictly local phonological processes*. University of Delaware Doctoral Dissertation.
- Chandlee, Jane & Eyraud, Rémi & Heinz, Jeffrey. 2014. Learning strictly local subsequential functions. *Transactions of the Association for Computational Linguistics* 2. 491–503. DOI: https://doi.org/10.1162/tacl_a_00198
- Chandlee, Jane & Eyraud, Rémi & Heinz, Jeffrey. 2015. Output strictly local functions. In *Proceedings of the 14th Meeting on the Mathematics of Language (MOL 2015)*, 112–125. DOI: <https://doi.org/10.3115/v1/W15-2310>
- Chandlee, Jane & Heinz, Jeffrey. 2018. Strict locality and phonological maps. *Linguistic Inquiry* 49. 23–60. DOI: https://doi.org/10.1162/LING_a_00265
- Chandlee, Jane & Heinz, Jeffrey & Jardine, Adam. 2018. Input strictly local opaque maps. *Phonology* 35. 171–205. DOI: <https://doi.org/10.1017/S0952675718000027>
- Chandlee, Jane & Lindell, Steven. under review. A logical characterization of Strictly Local functions. Ms., Haverford College.
- Chomsky, Noam. 1956. Three models for the description of language. *IRE Transactions on Information Theory* 2. 113–124. DOI: <https://doi.org/10.1109/TIT.1956.1056813>
- Clements, George N. 1980. *Vowel harmony in nonlinear generative phonology: an autosegmental model*. Bloomington, IN: Indiana University Linguistics Club.
- Clements, George N. & Hume, Elizabeth. 1995. The internal organization of speech sounds. In Goldsmith, John (ed.), *The handbook of phonological theory*, 245–306. Cambridge, MA and Oxford, UK: Blackwell.
- Clements, George N. & Sezer, Engin. 1982. Vowel and consonant disharmony in Turkish. In van der Hulst, Harry & Smith, Norval (eds.), *The structure of phonological representations (Part II)*, 213–255. Dordrecht: Foris.
- Cole, Jennifer & Kisseberth, Charles. 1994. An optimal domains theory of harmony. In Yoon, James H. (ed.), *Proceedings of the Formal Linguistics Society of Mid-America* 5, 101–114. University of Illinois.
- De Santo, Aniello & Graf, Thomas. 2019. Structure sensitive tier projection: Applications and formal properties. In Bernardi, Rafaella & Koble, Greg & Pogodalla, Sylvain (eds.), *Formal grammar 2019* (Lecture Notes in Computer Science, vol. 11668), 35–50. Springer. DOI: https://doi.org/10.1007/978-3-662-59648-7_3
- Dettweiler, Stephen H. 2000. Vowel harmony and neutral vowels in C'Lela. *Journal of West African Languages* 18. 3–18.

- Dresher, B. Elan & Nevins, Andrew. 2017. Conditions on iterative rounding harmony in Oroqen. *Transactions of the Philological Society* 115. 365–394. DOI: <https://doi.org/10.1111/1467-968X.12104>
- Elgot, Calvin C. & Mezei, Jorge E. 1965. On relations defined by generalized finite automata. *IBM Journal of Research and Development* 9. 47–68. DOI: <https://doi.org/10.1147/rd.91.0047>
- Elmedlaoui, Mohamed. 1995. *Aspects de représentations phonologiques dans certains langues chamito-sémitiques [Aspects of phonological representations in certain Chamito-Semitic languages]*. Rabat, Morocco. Université Mohammed V Doctoral Dissertation.
- Fallon, Paul D. 1993. Liquid dissimilation in Georgian. In Katholl, Andreas & Bernstein, Michael (eds.), *Proceedings of the 10th eastern states conference on linguistics*, 105–116. Ithaca, NY: DMLL Publications.
- Finley, Sara. 2017. Locality and harmony: Perspectives from artificial grammar learning. *Language and Linguistics Compass* 11. DOI: <https://doi.org/10.1111/lnc3.12233>
- Gafos, Adamantios & Benus, Stefan. 2006. Dynamics of phonological cognition. *Cognitive Science* 30. 1–39. DOI: https://doi.org/10.1207/s15516709cog0000_80
- Gafos, Adamantios & Dye, Amanda. 2011. Vowel harmony: Transparent and opaque vowels. In van Oostendorp, Marc (ed.), *The Blackwell companion to phonology*, vol. 4. Wiley-Blackwell. DOI: <https://doi.org/10.1002/9781444335262.wbctp0091>
- Gainor, Brian & Lai, Regine & Heinz, Jeffrey. 2012. Computational characterizations of vowel harmony patterns and pathologies. In *Proceedings of the 29th West Coast Conference on Formal Linguistics*, 63–71. Somerville, MA: Cascadilla Press.
- Gick, Bryan & Pulleyblank, Douglas & Campbell, Fiona & Mutaka, Ngessimo. 2006. Low vowels and transparency in Kinande vowel harmony. *Phonology* 23. 1–20. DOI: <https://doi.org/10.1017/S0952675706000741>
- Goldsmith, John. 1976. *Autosegmental phonology*. MIT Doctoral Dissertation.
- Goldsmith, John A. 1990. *Autosegmental and metrical phonology*. Oxford: Blackwell.
- Graf, Thomas & Mayer, Connor. 2018. Sanskrit n-retroflexion is Input-Output Tier-Based Strictly Local. In *Proceedings of SIGMORPHON 2018*, 151–160. DOI: <https://doi.org/10.18653/v1/W18-5817>
- Hansson, Gunnar Ólafur. 2010a. *Consonant harmony: Long-distance interaction in phonology* (University of California Publications in Linguistics 145). Berkeley, CA: University of California Press.
- Hansson, Gunnar Ólafur. 2010b. Long-distance voicing assimilation in Berber: Spreading and/or agreement. In *Proceedings of the 2010 annual conference of the Canadian Linguistic Association*.
- Hao, Yiding & Andersson, Samuel. 2019. Unbounded stress in subregular phonology. In *Proceedings of the 16th sigmorphon workshop on Computational Research in Phonetics, Phonology and Morphology*, 135–143. Florence, Italy: Association for Computational Linguistics. DOI: <https://doi.org/10.18653/v1/W19-4216>
- Hao, Yiding & Bowers, Dustin. 2019. Action-sensitive phonological dependencies. In *Proceedings of the 16th Workshop on Computational Research in Phonetics, Phonology and Morphology*, 218–228.

- Florence, Italy: Association for Computational Linguistics. DOI: <https://doi.org/10.18653/v1/W19-4225>
- Hayward, Richard J. 1990. Notes on the Aari language. In Hayward, Richard J. (ed.), *Omotiic language studies*, 425–493. London: School of Oriental and African Studies.
- Heath, Jeffrey. 2005. *A grammar of Tamashek (Tuareg of Mali)*. Berlin: Mouton de Gruyter. DOI: <https://doi.org/10.1515/9783110909586>
- Heinz, Jeffrey. 2010. Learning long-distance phonotactics. *Linguistic Inquiry* 41. 623–661. DOI: https://doi.org/10.1162/LING_a_00015
- Heinz, Jeffrey. 2018. The computational nature of phonological generalizations. In Hyman, Larry & Plank, Frans (eds.), *Phonological typology* (Phonetics and Phonology) chap. 5, 126–195. De Gruyter Mouton. DOI: <https://doi.org/10.1515/9783110451931-005>
- Heinz, Jeffrey & Lai, Regine. 2013. Vowel harmony and subsequentity. In Kornai, Andras & Kuhlmann, Marco (eds.), *Proceedings of the 13th Meeting on the Mathematics of Language (MOL 13)*, 52–63. Sofia, Bulgaria: Association for Computational Linguistics.
- Heinz, Jeffrey & Rawal, Chetan & Tanner, Herbert G. 2011. Tier-based strictly local constraints for phonology. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics*, 58–64. Portland, OR: Association for Computational Linguistics.
- Hyman, Larry. 1995. Nasal consonant harmony at a distance: The case of Yaka. *Studies in African Linguistics* 24. 5–30.
- Jardine, Adam. 2016. Computationally, tone is different. *Phonology* 33. 247–283. DOI: <https://doi.org/10.1017/S0952675716000129>
- Jardine, Adam & Heinz, Jeffrey. 2016. Learning Tier-based Strictly 2-Local languages. *Transactions of the Association for Computational Linguistics* 4. 87–98. DOI: https://doi.org/10.1162/tacl_a_00085
- Jardine, Adam & McMullin, Kevin. 2017. Efficient Learning of Tier-Based Strictly k-Local Languages. In *International Conference on Language and Automata Theory and Applications (LATA 2017)*, 64–76. DOI: https://doi.org/10.1007/978-3-319-53733-7_4
- Johnson, C. Douglas. 1972. *Formal aspects of phonological description*. The Hague: Mouton. DOI: <https://doi.org/10.1515/9783110876000>
- Jurjec, Peter. 2011. *Feature spreading 2.0: A unified theory of assimilation*. University of Tromsø Doctoral Dissertation.
- Kaplan, Ronald M. & Kay, Martin. 1994. Regular models of phonological rule systems. *Computational Linguistics* 20. 331–378.
- Kari, Ethelbert E. 2007. Vowel harmony in Degema, Nigeria. *African Study Monographs* 28. 87–97.
- Kaun, Abigail R. 1995. *The typology of rounding harmony: An Optimality-Theoretic approach*. University of California, Los Angeles Doctoral dissertation.
- Korn, David. 1969. Types of labial vowel harmony in the Turkic languages. *Anthropological Linguistics* 11. 98–106.
- Lambert, Dakotah & Rogers, James. 2020. Tier-Based Strictly Local Stringsets: Perspectives from Model and Automata Theory. In *Proceedings of the Society for Computation in Linguistics (SCiL) 2020*, 330–337. New Orleans, Louisiana.

- Lamont, Andrew & O'Hara, Charlie & Smith, Caitlin. 2019. Weakly deterministic transformations are subregular. In *Proceedings of the 16th workshop on computational research in phonetics, phonology and morphology*, 196–205. Association for Computational Linguistics. DOI: <https://doi.org/10.18653/v1/W19-4223>
- Li, Bing. 1996. *Tungusic vowel harmony*. The Hague: Holland Academic Graphics.
- Lind, Douglas & Marcus, Brian. 1995. *Symbolic dynamics and coding*. Cambridge UP.
- Lombardi, Linda. 1999. Positional faithfulness and voicing assimilation in Optimality Theory. *Natural Language & Linguistic Theory* 17. 267–302. DOI: <https://doi.org/10.1002/9780470756171.ch17>
- Luo, Huan. 2017. Long-distance consonant agreement and subsequentiality. *Glossa: A Journal of General Linguistics* 2. 1–25. DOI: <https://doi.org/10.5334/gjgl.42>
- Mayer, Connor & Major, Travis. 2018. A challenge for tier-based strict locality from Uyghur backness harmony. In *Formal Grammar 2018* (Lecture Notes in Computer Science 10950), 62–83. Berlin: Springer. DOI: https://doi.org/10.1007/978-3-662-57784-4_4
- McCollum, Adam G. & Baković, Eric & Mai, Anna & Meinhardt, Eric. 2020. Unbounded circumambient patterns in segmental phonology. *Phonology* 37. 215–255. DOI: <https://doi.org/10.1017/S095267572000010X>
- McCollum, Adam G. & Essegbey, James. 2018. Unbounded harmony is not always myopic: Evidence from Tutrugbu. In Bennett, William G. & Hrac, Lindsay & Storoshenko, Dennis R. (eds.), *Proceedings of the 35th West Coast Conference on Formal Linguistics*, 251–258. Somerville, MA: Cascadilla Proceedings Project.
- McCollum, Adam G. & Kavitskaya, Darya. 2018. Non-iterative vowel harmony in Crimean Tatar. In Bennett, William G. & Hrac, Lindsay & Storoshenko, Dennis Ryan (eds.), *Proceedings of the 35th West Coast Conference on Formal Linguistics*, 259–268.
- McMullin, Kevin. 2016. *Tier-Based Locality in Long-Distance Phonotactics: Learnability and Typology*. Vancouver, BC. University of British Columbia Doctoral Dissertation.
- McMullin, Kevin & Aksënova, Alëna & De Santo, Aniello. 2019. Learning phonotactic restrictions on multiple tiers. In *Proceedings of the Society for Computation in Linguistics (SCiL) 2019*, vol. 2. 377–378.
- McMullin, Kevin & Hansson, Gunnar Ólafur. 2016. Long-distance phonotactics as Tier-based Strictly 2-Local Languages. In Albright, Adam & Fullwood, Michelle A. (eds.), *Proceedings of the 2014 Annual Meeting on Phonology*. Washington, DC: Linguistic Society of America. DOI: <https://doi.org/10.3765/amp.v2i0.3750>
- McMullin, Kevin & Hansson, Gunnar Ólafur. 2019. Inductive learning of locality relations in segmental phonology. *Laboratory Phonology* 10. DOI: <https://doi.org/10.5334/labphon.150>
- McNaughton, Robert & Papert, Seymour A. 1971. *Counter-free automata*. Cambridge, MA: MIT Press.
- Meinhardt, Eric & Mai, Anna & Baković, Eric & McCollum, Adam G. 2020. On the proper treatment of weak determinism: Subsequentiality and simultaneous application in phonological maps. Unpublished manuscript.
- Michel, Daniel. 2009. Positional transparency in C'lela. In *Chicago linguistic society (cls) 45*. Chicago: Chicago Linguistic Society.
- Nevins, Andrew. 2010. *Locality in Vowel Harmony* (Linguistic Inquiry Monographs 55). MIT Press. DOI: <https://doi.org/10.7551/mitpress/9780262140973.001.0001>

- Ní Chiosáin, Máire & Padgett, Jaye. 2001. Markedness, segment realization, and locality in spreading. In *Segmental phonology in Optimality Theory: Constraints and representations*, 118–156. Cambridge: Cambridge University Press. DOI: <https://doi.org/10.1017/CBO9780511570582.005>
- Odden, David. 1994. Adjacency parameters in phonology. *Language* 70. 289–330. DOI: <https://doi.org/10.2307/415830>
- O'Hara, Charlie & Smith, Caitlin. 2019. Computational complexity and sour-grapes-like patterns. In *Supplemental proceedings of the Annual Meeting on Phonology 2018*. DOI: <https://doi.org/10.3765/amp.v7i0.4502>
- Orr, Carolyn. 1962. Ecuador quichua phonology. In Elson, Benjamin (ed.), *Studies in Ecuadorian Indian languages*, 60–77. Norman, Oklahoma: Summer Institute in Linguistics.
- Padgett, Jaye. 1995. Partial class behaviour and nasal place assimilation. In *Proceedings of the Arizona Phonology Conference: Workshop on features in Optimality Theory*, 145–183.
- Payne, Amanda. 2017. All dissimilation is computationally subsequential. *Language* 93. 353–371. DOI: <https://doi.org/10.1353/lan.2017.0076>
- Prince, Alan & Smolensky, Paul. 2004. *Optimality Theory: Constraint interaction in generative grammar*. Malden, MA: Blackwell.
- Pulleyblank, Douglas. 2002. Harmony drivers: No disagreement allowed. In Larson, Julie & Paster, Mary (eds.), *Proceedings of the 28th Annual Meeting of the Berkeley Linguistics Society*, 249–297. University of California, Berkeley. DOI: <https://doi.org/10.3765/bls.v28i1.3841>
- Ribeiro, Eduardo R. 2003. Directionality in vowel harmony: The case of Karajá (Macro-Jê). In Larson, Julie & Paster, Mary (eds.), *Proceedings of the Berkeley Linguistics Society 28*, 475–485. Berkeley, CA: Berkeley Linguistics Society. DOI: <https://doi.org/10.3765/bls.v28i1.3859>
- Rice, Keren. 1993. A reexamination of the feature [sonorant]: The status of 'sonorant obstruents'. *Language* 69. 308–344. DOI: <https://doi.org/10.2307/416536>
- Rogers, James & Heinz, Jeffrey & Fero, Margaret & Hurst, Jeremy & Lambert, Dakotah & Wibel, Sean. 2013. Cognitive and sub-regular complexity. In *Formal Grammar* (Lecture Notes in Artificial Intelligence 8036), 90–108. Springer.
- Rogers, James & Pullum, Geoffrey K. 2011. Aural pattern recognition experiments and the subregular hierarchy. *Journal of Logic, Language and Information* 20. 329–342. DOI: <https://doi.org/10.1007/s10849-011-9140-2>
- Rose, Sharon & Walker, Rachel. 2004. A typology of consonant agreement as correspondence. *Language* 80. 475–531. DOI: <https://doi.org/10.1353/lan.2004.0144>
- Rose, Sharon & Walker, Rachel. 2011. Harmony Systems. In Goldsmith, John & Riggle, Jason & Yu, Alan C. (eds.), *The Handbook of Phonological Theory*. Blackwell 2nd edn. DOI: <https://doi.org/10.1002/9781444343069.ch8>
- Sakarovitch, J. 2009. *Elements of automata theory*. Cambridge University Press.
- Smith, Caitlin & O'Hara, Charlie. 2019. Formal characterizations of true and false sour grapes. In *Proceedings of the Society for Computation in Linguistics 2019*, vol. 2. 338–341.

- Svantesson, Jan-Olof & Tsendina, Anna & Karlsson, Anastasia & Franzén, Vivian. 2005. *The phonology of Mongolian*. New York: Oxford University Press.
- van Oostendorp, Marc & Revithiadou, Anthi. 2005. Quasi-opacity and headed spans in Silly and Megisti Greek. Paper presented at the 13th Manchester Phonology Meeting.
- Vaysman, Olga. 2009. *Segmental alternations and metrical theory*. MIT Doctoral dissertation.
- Vaysse, Odile. 1986. Addition molle et fonctions p -locales. *Semigroup Forum* 34. 157–175. DOI: <https://doi.org/10.1007/BF02573160>
- Walker, Rachel. 2000. Yaka nasal harmony: Spreading or segmental correspondence. In *Annual Meeting of the Berkeley Linguistics Society (BLS 26)*, 323–332. DOI: <https://doi.org/10.3765/bls.v26i1.1164>
- Walker, Rachel. 2001. Round licensing and bisyllabic triggers in Altaic. *Natural Language & Linguistic Theory* 19. 827–878.
- Walker, Rachel. 2011. *Vowel patterns in language*. New York: Cambridge University Press. DOI: <https://doi.org/10.1017/CBO9780511973710>
- Walker, Rachel. 2012. Vowel Harmony in Optimality Theory. *Language and Linguistics Compass* 6. 575–592. DOI: <https://doi.org/10.1002/lnc3.340>
- Walker, Rachel. 2014. Nonlocal trigger-target relations. *Linguistic Inquiry* 45. 501–523. DOI: https://doi.org/10.1162/LING_a_00165
- Walker, Rachel & Byrd, Dani & Mpiranya, Fidèle. 2008. An articulatory view of Kinyarwanda coronal harmony. *Phonology* 25. 499–535. DOI: <https://doi.org/10.1017/S0952675708001619>
- Walker, Rachel & Mpiranya, Fidèle. 2006. On triggers and opacity in coronal harmony. In Cover, Rebecca T. & Kim, Yuni (eds.), *Proceedings of the 31st Annual Meeting of the Berkeley Linguistics Society*. University of California, Berkeley. DOI: <https://doi.org/10.3765/bls.v31i1.880>
- Waterson, Natalie. 1956. Some aspects of the phonology of the nominal forms of the Turkish word. *Bulletin of the School of Oriental and African Studies, University of London* 18. 578–591. DOI: <https://doi.org/10.1017/S0041977X00088066>
- Wilson, Colin. 2003. Experimental investigation of phonological naturalness. In Tsujimura, Mimura & Garding, Gina (eds.), *Proceedings of the 22nd West Coast Conference on Formal Linguistics*, 533–546. Somerville, MA: Cascadia Press.
- Wilson, Colin. 2006. Unbounded spreading is myopic. In *Phonology Fest Workshop on Current Perspectives on Phonology*. Indiana University.
- Zhang, Xi. 1996. *Vowel systems of the Manchu-Tungus languages of China*. University of Toronto Doctoral Dissertation.
- Zhang, Xi & Drescher, B. Elan. 1996. Labial harmony in written Manchu. *Saksaha: A Review of Manchu Studies* 1. 13–24.
- Zhang, Yanchang & Li, Bing & Zhang, Xi. 1989. *The Oroqen language*. Changchun: Jilin University Press.

