Consonant Harmony

Gunnar Ólafur Hansson, University of British Columbia, Department of Linguistics, Vancouver B.C. V6T 1Z4 Canada

Key Points

- Consonant harmony is long-distance assimilation between consonants, where intervening vowels and/or consonants appear unaffected.
- Consonant harmony always involves agreement in some specific phonological feature; it is not a general preference for co-occurring consonants to be "similar".
- Harmony is typically limited to consonants that are already very similar to each other (i.e., share many properties besides the harmony feature).
- While the transparency of intervening consonants has often been attributed to their lacking the harmony feature, many cases exist where such an explanation fails.
- Consonant harmony shares many characteristics with phonological speech errors, and may have its roots in motor planning and in lexical storage and retrieval.
- The cross-linguistic typology of consonant harmony patterns has contributed significantly to theories of phonological representation, computation, and learnability.

Abstract

Consonant harmony, or long-distance assimilation between consonants, is a comparatively rare but important type of sound pattern. It is attested for a wide range of phonological features, though it most commonly affects sibilants. The long-range character of consonant harmony makes it unusual and raises important questions for phonological theory. How are the interacting consonants able to "see" each other across intervening vowels and consonants, such that assimilation is triggered? How is the harmony feature transmitted from one consonant to the other, across those intervening segments? How are learners able to detect, and generalize, such long-distance dependencies within the input?

Introduction

The term *consonant harmony* is generally used to refer to any type of systematic sound pattern that involves *assimilation* between consonants that are *nonadjacent*—that is, separated from each other by one or more intervening vowels and/or consonants, which do not themselves appear to participate in the assimilation or be affected by it in any way—and may thus alternatively be referred to as *long-distance consonant agreement* (Rose & Walker, 2004). While a similar phenomenon is well known in child language, usually

 https://doi.org/10.1016/B978-0-323-95504-1.01000-0

involving place of articulation, such as when an English-learning child produces [gʌk] for /dʌk/ 'duck' (see Levelt, 2011 for an overview), this article focuses exclusively on consonant harmony as manifested in adult phonological systems.

While it is less frequent in the world's languages than *vowel harmony*, consonant harmony has played (and continues to play) an important role in the development of theoretical models of phonology as an aspect of grammar. Its study has contributed to our understanding of the factors that shape phonological systems and their cross-linguistic typology. This is especially due to its seemingly "nonlocal" character (see **Locality and Proximity** and **Transparency: Specification versus Similarity**) and the learnability challenges that such long-distance dependencies pose (see **Learning and Computational Complexity**).

Attestation

Although it is generally considered to be a relatively rare type of sound pattern, consonant harmony is attested in a wide range of language families and geographical regions that span the globe (with the possible exception of Australia). The most comprehensive cross-linguistic survey of consonant harmony to date is Hansson (2010a), but many additional cases have come to phonologists' attention since then, most notably a large number of South Asian languages with nonsibilant coronal harmony (Arsenault, 2012).[1]

Manifestations

The clearest manifestation of consonant harmony is when it gives rise to *alternations*. Nearly all of the examples in **Table 1** are of this type, such as when, in the Samala word /s-am-net-in-waʃ/ → [ʃamnetiniwaʃ], the 3rd person subject prefix /s-/ is realized as [ʃ-] due to the presence of a postalveolar sibilant later in the word (in the PAST suffix /-waʃ/).

However, consonant harmony is often confined to morpheme–internal contexts (see **Domains and Productivity**). In such cases, where the harmony effectively amounts to a *morpheme structure constraint*, its overt manifestation is mainly as a static *phonotactic* generalization over the sound shapes of (root) morphemes within the lexicon. In Basque, for example, laminal sibilants /t͡s̻, s̻/ (orthogr. <tz, z>) and apical sibilants /t͡s̺, s̺/ (orthogr. <ts, s>) do not co-occur with each other within roots, though they are free to do so in heteromorphemic contexts. Such morpheme-internal consonant harmony may in some cases be a *gradient* (i.e., probabilistic) phonotactic rather than an absolute gap (see, e.g., Arsenault & Kochetov, 2011 on Kalasha), analogous to what is often found with dissimilatory co-occurrence restrictions (e.g., Frisch et al., 2004).

Nevertheless, even morpheme-internal consonant harmony can have "active" manifestations, confirming its psychological reality and its status as part of the phonological grammar. This is the case, for instance, when diachronic processes (e.g., other sound changes, borrowing, rebracketing) lead to disharmonic forms which then get harmonized, as seen in the English loanwords into Zulu in **Table 1**. As another example, Basque /s̺inet͡s̺i/ 'believe' was earlier /s̺inet͡s̻i/ (etymologically a compound /s̺in/ 'truth' + /(h)et͡s̻i/ 'consider').

Table 1 Illustrative examples of consonant harmony involving a variety of phonological features. For more detailed descriptions and data sources, see Hansson (2010a).

Type	*Language (Family)*	*Example*	*Assimilation*
Coronal sibilant harmony	Samala, a.k.a. Ineseño (Chumashan)	/s-am-net-in-waʃ/ → [ʃamnetiniwaʃ] 'they did it to you'	s…ʃ → ʃ…ʃ
Coronal non-sibilant harmony	Mayak (Nilotic)	/tiɗ-ʌt̪/ → [tiɗ-ʌt] 'doctor'	t…t̪ → t…t
Liquid harmony	Bukusu (Bantu)	/-rum-il-a/ → [-rumira] 'send for'	r…l → r…r
Dorsal harmony	Tlachichilco Tepehua (Totonacan)	/ʔuks-laqt͡s'in/ → [ʔoqslaqt͡s'in] 'look at (across surface)'	k…q → q…q
Laryngeal harmony	Zulu (Bantu)	English *court* /kʰɔt/ → /i-kʰôtʰo/ English *beat* /bit/ → /ùm-bídi/	kʰ…t → kʰ…tʰ b…t → b…d
Stricture harmony	Yabem (Oceanic)	/se-dàgù/ → [tédàgù] 'they follow (irrealis)'	s…d → t…d
Secondary-articulation harmony	Tsilhqot'in (Dene/Athabaskan)	/sˤɛ-i-ɬ-t͡ʃʰæz/ → [siɬt͡ʃʰæz] 'I roasted it'	sˤ…z → s…z

[1]In the interest of brevity, sources describing individual languages are not cited in this article; the reader is referred to Hansson (2010a) for these. Exceptions are made for examples or important references not included in that survey.

Features

Although consonant harmony is attested for practically every phonological feature, there are clear asymmetries in frequency of attestation. The most frequent type is *coronal harmony*, where coronals (or rather, some subset of these; see Similarity) are required to agree with respect to some inherently coronal-specific distinction in place of articulation. Particularly common is *sibilant harmony*, typically involving the contrast between alveolars and either postalveolars or retroflexes (or all three, e.g., in some Omotic languages), or between laminal and apical alveolars (e.g., in Basque), though such harmony may extend to (nonsibilant) dentals as well (e.g., Tahltan; Shaw, 1991).

Usually *nonsibilant* coronals (e.g., /t/, /n/ or /l/) are neutral and inert to coronal harmonies that involve sibilant distinctions such as /s/ versus /ʃ/ (cf. the Samala example in **Table 1**). However, cases exist where alveolar plosives serve as the counterpart to postalveolar (sibilant) affricates (e.g., /t…t͡ʃ/ → [t͡ʃ…t͡ʃ] in Kera). Coronal harmony based on the retroflex versus nonretroflex (denti-alveolar) distinction in stops—that is, plosives and nasals (e.g., /t, d, n/ vs. /ʈ, ɖ, ɳ/)—is particularly widespread on the South Asian subcontinent (Arsenault, 2012). Coronal harmony is even attested for clicks, involving the dental versus post-alveolar versus lateral distinction, for instance in Ndebele. Participation of lateral clicks makes that case an exception to the general tendency for lateral consonants to be neutral to coronal harmony (cf. Samala /ha-s-xintila-waʃ/ → [haʃxintilawaʃ] 'his former non-Christian name'). Such exceptions can occasionally be found with pulmonic lateral obstruents as well (e.g., historical /s…ɬ/ > /ɬ…ɬ/ in Thao). Finally, so-called *liquid harmony*, where lateral approximants and rhotics assimilate at a distance (e.g., in Bukusu; see **Table 1**), could also be viewed as a subtype of coronal harmony.

Analogous to such coronal-specific contrasts, the velar versus uvular distinction can be the basis for *dorsal harmony* (e.g., in Totonacan languages) Alternatively, consonant harmony may take the form of *laryngeal harmony*: long-distance assimilation in voicing, aspiration (as in Zulu), or the pulmonic versus glottalic (ejective or implosive) distinction. As for nasality, it is useful to distinguish *nasal consonant harmony* (e.g., in Yaka) from the far more common phenomenon of *nasal harmony* (Botma, 2024; Walker, 2011), whereby nasality spreads in a local fashion from a nasal vowel or consonant to vowels and (some) consonants alike (e.g., Johore Malay /pəŋawasan/ → [pəŋãw̃ãsan] 'supervision'). In very rare instances, consonant harmony may operate over a manner distinction purely in *constriction degree* (e.g., in Yabem) or in *secondary articulation* (e.g., in Tsilhqot'in). Finally, Danis (2019) argues that the co-occurrence ban on velars /k, ɡ/ and labial-velars /k͡p, ɡ͡b/ in Ngbaka constitutes a unique example of *major-place harmony*, though it might be more comparable to secondary-articulation harmony (e.g., the ban on labials /p, m/ co-occurring with (labio-)velarized labials /pˠ, mˠ/ in Pohnpeian).

Characteristics

Similarity

Consonant harmony is very often sensitive to the relative similarity of the interacting segments. In essence, consonant harmony is frequently *parasitic* on other features, such that harmony (agreement) in feature [F] is only enforced between segments that also happen to share identical specifications for some set of additional features {[G], [H], …}. Alternatively, enforcement of harmony between less similar consonants may be subject to limitations that do not affect harmony between more similar consonants.

For example, root-internal coronal harmony in Kalasha (Arsenault, 2012; Arsenault & Kochetov, 2011) is enforced almost categorically for pairs of coronals that agree in *manner of articulation*: for plosive/plosive, affricate/affricate, and fricative/fricative combinations, if one of the two is retroflex, the other must also be (e.g., /ʈ͡ʂãɖ͡ʐa/ 'pinewood torch', historically harmonized from Middle Indo-Aryan */t͡ʃanɖ͡ʐa/), while for different-manner pairs, disharmony is not uncommon (e.g., /t͡ʃuʂ-ik/ 'to suck'). In closely related Indus Kohistani, retroflexion harmony is not constrained by similarity in this way (cf. harmonized /ʈ͡ʂoːʂáṽ/ 'to suck out').

Besides manner of articulation, other shared properties that may contribute to such similarity effects are *place of articulation* and *laryngeal* features. In Zulu and Ndebele, root-internal laryngeal harmony in aspiration and voicing is overridden by a general phonotactic that bans /kʰ, ɡ/ outside of root-initial position (e.g., /-kʰupʰ-/ 'remove' but /-pʰek-/ 'cook, brew' rather than */-pʰekʰ-/). However, when the two plosives are homorganic, harmony instead overrides this ban (e.g., /-kʰokʰ-/ 'pull, draw out'). The voicing harmony found in many Amazigh (Berber) languages is restricted by both place and manner, encompassing only coronal (sibilant) fricatives. Thus, in Tamajaq Tuareg, causative /s- / becomes voiced before any of the voiced coronal fricatives /z, zˤ, ʒ/ (e.g., [z-əntəz] 'cause to extract'), but not before a voiced noncoronal fricative or voiced coronal nonfricative (e.g., [s-əɣdər] 'cause to betray'). In Nkore/Kiga, sibilant harmony (in anteriority) is likewise limited to fricatives but is also constrained by laryngeal features (see Hansson, 2010a, pp. 333–335): harmony seems to apply to postalveolar…alveolar sequences only when these agree in voicing (e.g., *[ʃ…s] is avoided in favor of [s…s], but disharmonic [ʃ…z] is tolerated).

Locality and Proximity

Consonant harmony is never sensitive to the properties of any *intervening vowels* that stand between the two consonants. This is in striking contrast to vowel harmony, which may be affected by intervening consonants in a variety of ways (Hansson, 2024). As for potential interference by *intervening consonants*—that is, where certain types of (nonharmonizing) consonants can act as *blockers*, disrupting harmony from applying across them—it was long believed that this, too, was completely unattested (Gafos, 1999; Rose & Walker, 2004). However, in the last two decades, a handful of cases have come to light that unambiguously involve blocking effects in consonant harmony, most notably sibilant harmony in Kinyarwanda (Walker et al., 2008) and in western dialects of Slovenian (Jurgec, 2011), as well as (sibilant) voicing harmony in Imdlawn Tashlhiyt (Hansson, 2010b). For example, in the variety of Slovenian described by Jurgec (2011), sibilant harmony can apply (optionally) across intervening noncoronals as well as coronal sonorants (e.g., /za-klon-iʃt͡ʃe/ → [ʒakloniʃt͡ʃe] 'bomb shelter', cf. [zaklon] 'shelter') but never across an intervening coronal obstruent (e.g., /zida-ʃ/ → [zidaʃ] 'you build', not *[ʒidaʃ]).

More commonly, consonant harmony is subject to restrictions on *proximity*, that is, the distance between the two potentially interacting consonants. A very frequent limitation is for harmony to apply only when these are separated just by an intervening vowel (…C̲VC̲…) but not when the intervening material includes one or more consonants (e.g., …C̲VCVC̲…); in effect, all nonparticipating consonants act as blockers (McMullin, 2016). This has been referred to as *strictly transvocalic* locality (Hansson, 2010a), but others see *syllable adjacency* as the relevant criterion (Bennett, 2015; Odden, 1994; Rose & Walker, 2004); see Hansson (2020) for discussion of these two interpretations.

As an illustrative example of the distinction between strictly transvocalic consonant harmony and its *unbounded* counterpart, where harmony applies regardless of distance, consider two Omotic languages, Koorete (a.k.a. Koyra) and Aari. Both display sibilant anteriority harmony, by which /s/ in suffixes assimilates to a postalveolar sibilant in the stem. In Koorete, this harmony occurs only in transvocalic contexts (e.g., /tut͡ʃ'-us-/ → [tut͡ʃ'uʃ-] 'cause to shut') but not at greater distances (e.g., /t͡ʃ'aːn-us-/ → [t͡ʃ'aːnus-] 'cause to load', not *[t͡ʃ'aːnuʃ-]). In Aari, by contrast, harmony applies even when the two sibilants are further apart (e.g., /ʒa̤ːg-er-s-e/ → [ʒa̤ːgerʃe] 'it was sewn').

Directionality

Other things being equal, consonant harmony seems to display a preference for *regressive* (i.e., *anticipatory*) assimilation. For instance, in Samala (see **Table 1**), the last sibilant within the word determines the anteriority of all preceding sibilants, regardless of the morphological affiliation or phonological characteristics of the consonants involved; thus, for example, /s…t͡ʃ/ and /ʃ…t͡s/ are harmonized to [ʃ…t͡ʃ] and [s…t͡s], respectively. This tendency for anticipatory rather than perseveratory influence has been cited as one of several notable parallels between consonant harmony and phonological speech errors (Hansson, 2010a; see also Rose & Walker, 2004 for discussion of other such parallels).

Very frequently, however, the direction of assimilation is instead dictated by morphological constituent structure, such that affix consonants assimilate to consonants within the base of affixation (that is, either in the root or in another, "inner" affix); such "outward" directionality is referred to as *stem control* (Baković, 2000). Sibilant harmony in Koorete (see **Locality and Proximity**) illustrates how progressive assimilation may result from stem control: causative /-us/ harmonizes with a [−anterior] sibilant within the preceding root. Contrast this with a Samala word like /s-api-t͡ʃʰo-us/ → [sapit͡sʰolus] 'he has a stroke of good luck', where the 3rd person object suffix /-us/ instead triggers regressive [+anterior] harmony onto the preceding root /t͡ʃʰo/ 'good'.

A few cases seem to exhibit progressive consonant harmony that cannot straightforwardly be attributed to stem control (e.g., Santiago Tz'utujil; Lyskawa & Ranero, 2021). Results from artificial grammar learning experiments also suggest that purely progressive consonant harmony is no less robustly learned than its regressive counterpart (Conklin et al., 2024).

Domains and Productivity

As noted above, it is quite common for consonant harmony to be confined to *morpheme-internal* contexts, manifesting as a static co-occurrence restriction. Such morpheme-internal harmony may nevertheless be productively enforced (see **Manifestations** for examples). Even when consonant harmony does apply across morpheme boundaries, resulting in alternation, it tends to be limited to some relatively more lexicalized, word-internal domain. For instance, Misantla Totonac dorsal harmony applies to derivational affixes but not inflectional ones. Similarly, in Bantu languages like Yaka, nasal consonant harmony is confined to the derivational stem, consisting of the verb root and suffixes (most of which have valence-changing or aspectual functions), while prefixes (denoting inflectional properties like agreement, negation or tense) lie outside of the domain of harmony.

Where there is evidence that application of phonological processes is "interleaved" with morpho-syntactic structure—as assumed, for example, in Stratal Phonology (Bermúdez-Otero, 2018)—consonant harmony is often limited to earlier/inner layers (levels, strata, cycles). This can be seen when later phonological processes create apparent exceptions to harmony or obscure its application. For instance, in a Yaka word like /N-mak-idi/ → [mbakiní] 'I carved', the root-initial /m/ has triggered nasal consonant harmony in the perfective suffix /-idi/, even though it later undergoes denasalization due to the 1sg prefix /N- /, rendering the motivation for harmony opaque (cf. /N-bak-idi/ → [mbakidí] 'I caught', with no harmony).

Explanations

As was noted in **Directionality**, various resemblances exist between patterns of consonant harmony and spontaneous phonological speech errors (slips of the tongue; Hansson, 2010a; Rose & Walker, 2004). Consider, for instance, how the Samala sibilant harmony seen in **Table 1** seems analogous to the mispronunciation of /siʃoɹ/ 'seashore' as [ʃiʃoɹ] in the classic English tongue-twister "She sells seashells by the seashore." In part for this reason, consonant harmony is often considered to be rooted in *motor planning* (Tilsen, 2019); in essence, these synchronic sound patterns may be "phonologized speech errors" (Garrett & Johnson, 2013).

It is striking, however, how a large majority of consonant harmony cases (though, crucially, not all!) involve articulatory properties which could potentially be sustained through intervening vowels (and some consonants) with little or no impact on their acoustic realization (Gafos, 1999). This suggests that *phonetic coarticulation* may also play a contributing role in the emergence of consonant harmony, at least in some cases.

Another domain whose potential relevance has not been sufficiently explored is that of *memory* and *recall*, that is, lexical storage and retrieval. Patterns of consonant harmony provide an advantage in this regard (Gafos, 2021). Related arguments for the role of *perceptual contrast* in shaping nonadjacent consonant co-occurrence patterns are raised by Gallagher (2010).

That being said, some individual cases of consonant harmony appear to have come about by idiosyncratic sequences of historical changes that do not quite fit any of these descriptions (Hansson, 2007). It is thus unlikely that a one-size-fits-all explanation can be provided for consonant harmony as a phenomenon.

Theoretical Significance

The study of consonant harmony, and of the typological characteristics of such sound patterns, has had a greater impact on theoretical developments in phonology than might be expected given its relative rarity in the world's languages. The theoretical significance of consonant harmony has mainly been in relation to three areas, which I will refer to as the "transparency problem," the "mechanism problem" and the "learning problem."

Transparency: Specification versus Similarity

The *transparency problem* is how to account for the fact that segments that intervene between a harmony trigger and harmony target may be invisible to, and apparently unaffected by, the assimilatory interaction that is taking place across them. For instance, when the Yaka perfective verb stem /-mí:tuk-idi/ 'sulked' surfaces as [-mí:tuk-ini] due to harmony between the root-initial /m/ and the /d/ of the perfective suffix, without any noticeable nasalization on the intervening vowels or voiceless plosives, how are we to understand the inertness and transparency of the latter?

A highly influential approach to the transparency problem in harmony in general (not just consonant harmony) has been to treat the intervening segments in question as being effectively absent at some representational level where the assimilatory interaction takes place. One example, popular within autosegmental and feature-geometric approaches to phonological representation (see, e.g., Uffmann, 2011) is to assume that transparency of intervening segments is always due to their being *unspecified* for the harmonizing feature. As a result, the harmony trigger and target segments are effectively *adjacent* to each other, as long as adjacency is being computed on the representational *tier* defined by the harmony feature itself (or, rather, on the *plane* defined by that tier and the one that immediately dominates it; Odden, 1994). Notable examples of analyses in this vein include Ao (1991) for nasal consonant harmony in Kongo and Shaw (1991) for coronal harmony in Tahltan.

This is illustrated in Fig. 1, showing the application of sibilant harmony in the Samala word /s-am-net-in-waʃ/ → [ʃamnetiniwaʃ] 'they did it to you' from Table 1. Harmony involves assimilation in [±anterior], which is here assumed to be a dependent of the (unary, privative) place feature [coronal], which is in turn a dependent of the class node C-Place. Noncoronal consonants like /m/ *inherently* lack the feature [coronal] (and its dependent [±anterior]), and so do vowels (at least as a direct dependent of C-Place). Nonsibilant coronals like /t/ or /n/ are also assumed to lack any specification for [±anterior], due to their being *noncontrastively* (i.e., predictably, redundantly) alveolar in Samala. On some accounts, these nonsibilant alveolars would lack even a [coronal] feature, owing to coronal being the (universally) *unmarked* place of articulation (Shaw, 1991).

However, nonspecification for the harmony feature—whether attributed to irrelevance, noncontrastiveness, or unmarkedness—is inadequate as a general account of transparency in consonant harmony. For instance, in Tamajaq Tuareg and many other Amazigh (Berber) languages, sibilant voicing harmony applies across contrastively voiceless and voiced nonsibilant obstruents alike. Thus, we see the causative prefix /s- / become voiced in harmony with a voiced sibilant in the root even in cases

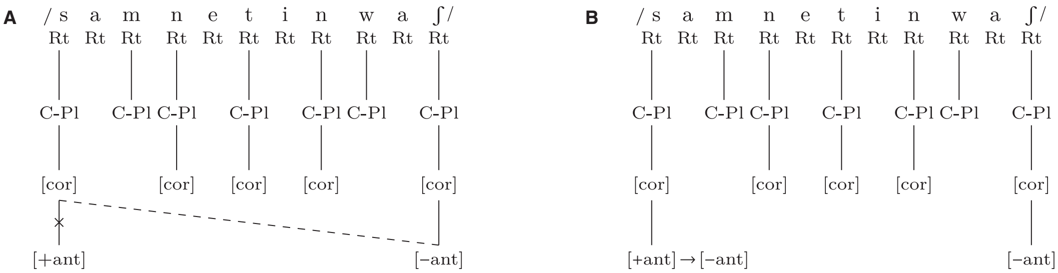

Fig. 1 Two alternative construals of the assimilatory mechanism involved in consonant harmony, illustrated using the Samala example from **Table 1**: (A) *feature spreading*, resulting in the sharing of a single featural element (autosegment); (B) *feature copying*, resulting in agreement (identical feature values). Abbreviated labels for feature-geometric nodes: Rt = Root; C-Pl = C-Place; [cor] = [coronal]; [ant] = [anterior].

like /s-əkləzˁ/ → [z-əkləzˁ] 'cause to invent' or /s-guləz/ → [z-əgːuləz] 'cause to remain', with [+voice] being transmitted from one sibilant to another across [−voice] /k/ and [+voice] /g/ alike.

A more recent approach maintains the core idea behind the nonspecification hypothesis—that some *tier* can indeed be defined, on which the harmony trigger and target segments do have a presence but the intervening transparent segments do not, such that the harmonizing segments are *adjacent* (on that tier)—but recasts the tier notion in a more contingent, nonrepresentational way. A tier is defined in terms of a *set* of segments, and a given word contains, on that tier, just those of its segments that are members of this set. For Tamajaq Tuareg, the sibilants [s, sˁ, ʃ, z, zˁ, ʒ] define what we may call the "sibilant tier," and forms like [z-əkləzˁ] and [z-əgːuləz] contain, on this tier, just the sequences [zzˁ] and [zz], respectively; their (ungrammatical) counterparts *[s-əkləzˁ] and *[s-əgːuləz] would instead contain [szˁ] and [sz]. The generalization, then, is that adjacent [−voice][+voice] sequences are not permitted on the sibilant tier and are repaired by regressive assimilation. To the extent that tier-defining sets correspond to *natural classes*, this also accounts for the strong tendency for consonant harmony to display *similarity* effects (see **Similarity**), including many that are challenging to explain for representational (e.g., feature-geometric) approaches.

Formal implementations of this latter approach include the optimality-theoretic Agreement by Correspondence model of long-distance assimilation and dissimilation (ABC; e.g., Bennett, 2015; Hansson, 2010a; Rose & Walker, 2004). Computational formalizations in terms of Tier-based Strictly Local (TSL) stringsets or functions (see **Learning and Computational Complexity** for discussion and references) also fall under the same rubric.

Mechanisms: Spreading versus Agreement

Intimately related to the transparency problem is the *mechanism problem*: by what means does one member of a pair of co-occurring (but nonadjacent) consonants assimilate to the other, and what is the most appropriate way to model this assimilation process—and the components of the phonological grammar (rules, or constraints) that give rise to it—formally?

From the perspective of how the phonological representation gets modified, there are two main options, which I will refer to as "feature sharing" and "feature copying," respectively. Consider yet again the Samala example /s-am-net-in-waʃ/ → [ʃamnetiniwaʃ] from **Table 1**, where an underlying alveolar…postalveolar sequence /s…ʃ/ surfaces as [ʃ…ʃ], with both sibilants articulated as postalveolar ([−anterior]). On a *feature sharing* analysis, this assimilation is understood as involving a single [−anterior] featural element (autosegment) coming to be simultaneously associated with both sibilants. To put it differently, the [−anterior] feature *spreads* from one sibilant to the other, as illustrated in **Fig. 1**A. By contrast, a *feature copying* analysis assumes that the internal featural composition of the one sibilant is simply altered under the influence of the other, changing the value of its [anterior] feature (from + to −), as shown in **Fig. 1**B.[2] That is, one sibilant *copies* the feature [anterior] from the other sibilant, resulting in the two having *identical* values for that feature.

On the assumption that a phonological feature (autosegment) corresponds, more or less directly, to an articulatory gesture, feature spreading should entail the temporal extension of that single gestural event (see, e.g. Gafos, 1999). This is problematic for many cases of consonant harmony, in which intervening vowels and consonants are clearly unaffected by the harmonizing feature (e.g., the Yaka and Yabem examples in **Table 1**). As a result, most current analyses of consonant harmony assume that it is driven by constraints (or rules) that merely require *agreement*—that is, identical specifications for the harmony feature—rather than the sort of feature-sharing configuration that would necessitate *spreading* (see, e.g. Hansson, 2010a; Rose & Walker, 2004). In any case, even when there is articulatory evidence that intervening segments are affected, suggesting that a single, extended gesture may be involved (e.g., sibilant retroflexion harmony in Kinyarwanda; Walker et al. 2008), this seems to be just an alternative means of achieving featural agreement between the harmony trigger and the (distal) target consonant (what Hansson, 2010b refers to as "agreement by spreading").

Learning and Computational Complexity

The long-distance character of consonant harmony interactions gives rise to the *learning problem*. By what means do learners detect such (potentially unbounded) dependencies between sounds in the input data? It is simply not feasible for a learner to track the (non-)co-occurrence, at any distance, of every conceivable segment or natural class of segments with every other conceivable segment or class. Therefore, the discovery of such nonadjacent dependency patterns must be guided by some *learning biases*: principles and limitations that narrow down the hypothesis space traversed by the learner.

For this reason, consonant harmony has figured prominently in the literature on computational phonology. The Subregular Hypothesis (Heinz, 2018) holds that all possible phonological patterns belong to small, well-defined subregions of the formal-language-theoretic class of *regular* languages (stringsets) and relations (string-to-string functions). The cross-linguistic typology of consonant harmony patterns has focused attention on the class known as Tier-based Strictly Local (TSL) stringsets (Heinz et al., 2011) and functions (Burness et al., 2021) and ongoing development of algorithms whereby these can be inferred from positive evidence (e.g., Jardine & McMullin, 2017).

[2]Although the diagram in **Fig. 1**B retains the same assumptions as **Fig. 1**A in terms of autosegmental representation (e.g., underspecification), these are not essential under a feature copying analysis.

A related area of research where consonant harmony figures frequently is in the use of *artificial grammar learning* (AGL) experiments to probe the ways in which human learners infer phonological patterns and generalize these to novel contexts (for an overview, see Finley, 2017). For example, AGL studies have found that learners robustly generalize harmony from beyond-transvocalic contexts to transvocalic ones but not vice versa (McMullin & Hansson, 2019). Learners also appear to be largely unable to learn consonant-consonant dependencies whose level of computational complexity exceeds that of the TSL class, for example, a harmony that applies exclusively at beyond-transvocalic distances (McMullin & Hansson, 2019), or a "first—last" harmony that holds only between word-initial and word-final sibilants, ignoring any intervening word-medial sibilants (Avcu & Hestvik, 2020; Lai, 2015).

Conclusions

Consonant harmony, or long-distance consonant agreement, may involve assimilation in any of a wide range of phonological features, though it most frequently involves distinctions that are specific to coronal consonants. This type of sound pattern may give rise to alternations, but it is frequently manifested merely as a static phonotactic restriction on morpheme shapes within the lexicon. Intervening vowels are always inert and transparent, and although the same is usually true of intervening consonants, several cases have come to light where certain types of consonants act as blockers. The long-distance character of consonant harmony interactions poses questions of theoretical importance, especially with regard to the role of locality and how it is computed. As a result, consonant harmony has figured prominently in the formal modelling of phonological systems, and of their acquisition, and will likely continue to do so in the future.

References

Ao, B. (1991). Kikongo nasal harmony and context-sensitive underspecification. *Linguistic Inquiry, 22*(1), 193—196.

Arsenault, P. (2012). *Retroflex consonant harmony in South Asia*. Doctoral dissertation. University of Toronto.

Arsenault, P., & Kochetov, A. (2011). Retroflex harmony in Kalasha: Agreement or spreading? In S. Lima, K. Mullin, & B. Smith (Eds.), *Proceedings of the 39th meeting of the North Eastern Linguistic Society* (pp. 55—66). Amherst, MA: GLSA Publications, University of Massachusetts.

Avcu, E., & Hestvik, A. (2020). Unlearnable phonotactics. *Glossa, 5*(1), 56.

Baković, E. (2000). *Harmony, dominance and control*. Doctoral dissertation. Rutgers University.

Bennett, W. G. (2015). *The phonology of consonants: Harmony, dissimilation and correspondence*. Cambridge: Cambridge University Press.

Bermúdez-Otero, R. (2018). Stratal phonology. In S.-J. Hannahs, & A. R. K. Bosch (Eds.), *The Routledge handbook of phonological theory* (pp. 100—134). London: Routledge.

Botma, B. (2024). Nasal harmony. In H. van der Hulst, & N. Ritter (Eds.), *The Oxford handbook of vowel harmony* (pp. 37—48). Oxford: Oxford University Press.

Burness, P., McMullin, K., & Chandlee, J. (2021). Long-distance phonological processes as tier-based strictly local functions. *Glossa, 6*(1), 99.

Conklin, J. T., Stork, E., Bastimar, R., & He, X. (2024). Impacts of direction and morphological locus of control on learnability of sibilant harmony. *Laboratory Phonology, 15*(1), 1—38.

Danis, N. (2019). Long-distance major place harmony. *Phonology, 36*(4), 573—604.

Finley, S. (2017). Locality and harmony: Perspectives from artificial grammar learning. *Language and Linguistics Compass, 11*(1), e12233.

Frisch, S. A., Pierrehumbert, J. B., & Broe, M. B. (2004). Similarity avoidance and the OCP. *Natural Language and Linguistic Theory, 22*(1), 179—228.

Gafos, A. (1999). *The articulatory basis of locality in phonology*. New York: Garland.

Gafos, A. (2021). Consonant harmony, disharmony, memory and time scales. In A. Ettinger, E. Pavlick, & B. Prickett (Eds.), *Proceedings of the Society for Computation in Linguistics (SCiL) 2021* (pp. 188—205). Association for Computational Linguistics. Online.

Gallagher, G. (2010). Perceptual distinctness and long-distance laryngeal restrictions. *Phonology, 27*(3), 435—480.

Garrett, A., & Johnson, K. (2013). Phonetic bias in sound change. In A. C. L. Yu (Ed.), *Origins of sound change: Approaches to phonologization* (pp. 51—97). Oxford: Oxford University Press.

Hansson, G.Ó. (2007). On the evolution of consonant harmony: The case of secondary-articulation agreement. *Phonology, 24*(1), 77—120.

Hansson, G.Ó. (2010a). *Consonant harmony: Long-distance interaction in phonology*. Berkeley, CA: University of California Press.

Hansson, G.Ó. (2010b). Long-distance voicing assimilation in Berber: Spreading and/or agreement? In M. Heijl (Ed.), *Actes du Congrès de l'ACL 2010 / 2010 CLA Conference Proceedings*. Canadian Linguistic Association. Online.

Hansson, G.Ó. (2020). Consonant harmony. In M. Aronoff (Ed.), *Oxford research encyclopedia of linguistics*. Oxford University Press.

Hansson, G.Ó. (2024). The role of consonants in vowel harmony. In H. van der Hulst, & N. Ritter (Eds.), *The Oxford handbook of vowel harmony* (pp. 17—36). Oxford: Oxford University Press.

Heinz, J. (2018). The computational nature of phonological generalizations. In L. M. Hyman, & F. Plank (Eds.), *Phonological typology* (pp. 126—195). Berlin: De Gruyter Mouton.

Heinz, J., Rawal, C., & Tanner, H. G. (2011). Tier-based strictly local constraints for phonology. In D. Lin, Y. Matsumoto, & R. Mihalcea (Eds.), *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics* (pp. 58—64). Portland, OR: Association for Computational Linguistics.

Jardine, A., & McMullin, K. (2017). Efficient learning of tier-based strictly *k*-Local languages. In F. Drewes, C. Martín-Vide, & B. Truthe (Eds.), *Proceedings of the 11th Language and Automata Theory and Applications conference (LATA 2017)* (pp. 64—76). Berlin: Springer.

Jurgec, P. (2011). *Feature spreading 2.0: A unified theory of assimilation*. Doctoral dissertation. University of Tromsø.

Lai, R. (2015). Learnable vs. unlearnable harmony patterns. *Linguistic Inquiry, 46*(3), 425—451.

Levelt, C. C. (2011). Consonant harmony in child language. In M. van Oostendorp, C. J. Ewen, E. Hume, & K. Rice (Eds.), *The Blackwell companion to phonology* (Vol. 3, pp. 1691—1716). Oxford: Wiley-Blackwell.

Lyskawa, P., & Ranero, R. (2021). Sibilant harmony in Santiago Tz'utujil (Mayan). *Proceedings of the Linguistic Society of America, 6*(1), 265—279.

McMullin, K. (2016). *Tier-based locality in long-distance phonotactics: Learnability and typology*. Doctoral dissertation. University of British Columbia.

McMullin, K., & Hansson, G.Ó. (2019). Inductive learning of locality relations in segmental phonology. *Laboratory Phonology, 10*(1), 1—53, 14.

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

Rose, S., & Walker, R. (2004). A typology of consonant agreement as correspondence. *Language, 80*(3), 475—531.

Shaw, P. A. (1991). Consonant harmony systems: The special status of coronal harmony. In C. Paradis, & J.-F. Prunet (Eds.), *The special status of coronals: Internal and external evidence* (pp. 125—158). San Diego, CA: Academic Press.

Tilsen, S. (2019). Motoric mechanisms for the emergence of non-local phonological patterns. *Frontiers in Psychology, 10*, 2143.

Uffmann, C. (2011). The organization of features. In M. van Oostendorp, C. J. Ewen, E. V. Hume, & K. Rice (Eds.), *The Blackwell companion to phonology* (Vol. 1, pp. 643–668). Oxford: Wiley-Blackwell.

Walker, R. (2011). Nasal harmony. In M. van Oostendorp, C. J. Ewen, E. Hume, & K. Rice (Eds.), *The Blackwell companion to phonology* (Vol. 3, pp. 1838–1865). Oxford: Wiley-Blackwell.

Walker, R., Byrd, D., & Mpiranya, F. (2008). An articulatory view of Kinyarwanda coronal harmony. *Phonology, 25*(3), 499–535.