Gong, Shuxiao & Zhang, Jie. 2025. The Obligatory Contour Principle as a substantive bias in phonological learning. *Glossa: a journal of general linguistics* 10(1). pp. 1–34. DOI: https://doi.org/10.16995/glossa.18122

Open Library of Humanities

The Obligatory Contour Principle as a substantive bias in phonological learning

Shuxiao Gong, Beijing Language and Culture University, gong@blcu.edu.cn

Jie Zhang, The University of Kansas, zhang@ku.edu

Understanding how native speakers acquire the phonological patterns in their language is a key task for the field of phonology. Numerous studies have suggested that phonological learning is a biased process: certain phonological patterns are more easily accessed and learned by the speakers and thus more likely to appear in languages, while others show acquisition difficulties and may occur less frequently. Therefore, an important aspect of understanding phonological learning and typology is to understand the nature of these learning biases. Obligatory Contour Principle (OCP), i.e., the avoidance of adjacent similar units in the lexicon, is one of the typologically well-attested phenomena that may originate from phonological learning biases. Using an artificial grammar learning experiment testing the learnability of several phonotactic patterns, we present evidence that the OCP can directly modulate phonological learning, in that similarity avoidance is easier to learn compared to other phonotactic patterns. Specifically, an OCP-based phonotactic pattern was better learned than a complexity-matching consonant major place harmony phonotactic pattern as well as an arbitrary control pattern. Based on the AGL experiment results and the phonetic foundation of similarity avoidance, we argue that the OCP can serve as a substantive bias that influences phonological learning and, eventually, linguistic typology.

1 Introduction

Typological surveys of phonological patterns in the world's languages (e.g., Gordon 2016) help us understand what patterns keep recurring and what patterns are rare cross-linguistically. A critical task for phonologists is to account for this observed typology: what are the mechanisms causing certain patterns to occur more frequently than others? One direction to tackle this issue is to focus on phonological acquisition. It has been argued that phonological typology is influenced by various types of cognitive predispositions and/or phonetic principles, known as learning biases (Wilson 2006; Zuraw 2007; Hayes et al. 2009; Hayes & White 2013). Patterns that are favored by these biases receive an advantage in phonological acquisition and are thus more likely to occur in language and survive generations of language change (Martin 2007; Moreton 2008).

The investigation of learning biases and phonological acquisition often relies on computational modeling and/or psycholinguistic experimentation. Multiple computational algorithms have been developed to simulate generations of first language acquisition based on iterative learning, and the simulation outcome often reveals the effects of learning biases, since certain patterns are particularly difficult for the learning algorithm to pass on to the next generation (Stanton 2016; Hughto 2018; McMullin & Hansson 2019). Experimental work on learning biases has an even longer tradition. For example, in Turkish, around 54% of the nouns of the language ending with a voiceless stop involve a voicing alternation for the stop when the noun is followed by a vowel-initial inflectional suffix, while the other 46% of such nouns do not involve this alternation. Becker et al. (2011) conducted a wug-test on Turkish speakers, asking them to inflect non-existing noun stems. The results demonstrated that, even though consonant place, word length, and vowel quality can all condition the voicing alternation of stem-final consonants in the Turkish lexicon (e.g., voicing alternation is less likely when the consonant is coronal, when the root size is monosyllabic, and when the preceding vowel is [-high] and [-back]), in the wug test, only the first two factors influenced speakers' alternation behaviors, while the vowel quality factor had little to no effect. Participants' selective learning outcome on statistically equally salient lexical distributions has been attributed to phonetic naturalness and/or pattern complexity. The conditioning of voicing alternation by consonant place and word length is supported by phonetic grounding and reported in other languages as well, while the interaction between voicing and vowel quality has no comparable phonetic motivation or typological manifestation. In sum, the Turkish wug-test results indicate that phonological patterns supported by phonetic motivations or conforming to certain complexity-based restrictions and cognitive predispositions may be easier to learn than other ones. Consequently, the biased learning among different patterns leads to differences in their typological frequency. Similarly, nonword acceptability judgment experiments found that not all phonotactic patterns in the lexicon are equally learned and noticed by the native speakers. Some work has shown that patterns conforming to certain phonetic grounding or structural properties are easier to learn and have stronger impact on speakers' nonword ratings (Kager & Pater 2012; Hayes & White 2013; Gong & Zhang 2021; Jin & Wang & Lu 2023). A more direct approach to test the learnability of phonological patterns and the learning biases involved is artificial grammar learning (AGL) experiments (Gomez & Gerken 1999). Through the full control of input materials, this approach allows us to design and manipulate different types of phonological patterns we are interested in, and then directly monitor and compare the learning outcome of these patterns in real time. Natural language materials usually do not offer such flexibility, since it is rare that all patterns that we want to examine occur simultaneously in a single natural language. Previous AGL studies have made significant contributions in determining different types of learning biases by comparing the learnability of various phonological patterns. For comprehensive reviews on AGL studies in phonology, see Moreton & Pater (2012a; 2012b), Finley (2017), and Zheng & Do (2025).

The AGL experiment of this study provides further evidence for how learning biases can potentially shape phonological typology. Previous studies have addressed this question by testing the learnability of various typologically well-attested sound patterns such as vowel harmony (Pycha et al. 2003; Finley & Badecker 2009; Martin & Peperkamp 2020), consonant harmony (Finley 2011; Lai 2015; McMullin & Hansson 2019), positionally conditioned feature contrasts (Myers & Padgett 2014; Glewwe 2022), and palatalization (Wilson 2006). In this paper, we focus on the phenomenon called the Obligatory Contour Principle (OCP) effect. This principle was originally formulated to describe the tonal disharmony patterns observed in West-African languages (Leben 1973). Later, it was extended to account for the co-occurrence restrictions on segmental features such as the avoidance of homorganic consonants within a root (McCarthy 1986). Indeed, at the segmental level, OCP is predominantly manifested as the dispreference for the co-occurrence of consonants with the same major place of articulation, known as the OCP-Place effect (Boll-Avetisyan & Kager 2014; 2016). For example, in Arabic, consonants with the same place occur much less frequently within a root than chance (McCarthy 1986). Such OCP-Place effects in the lexicon are attested in a wide range of genetically unrelated languages, including Dutch (Boll-Avetisyan & Kager 2014; 2016), English (Berkley 1994), Hebrew (Berent & Shimron 1997; Berent & Shimron & Vaknin 2001), Korean (Ito 2007), Javanese (Mester 1988), Muna (Coetzee & Pater 2008), Quechua (Gallagher 2010; 2016; 2019), among many others. In addition to acting as a morpheme structure constraint that regulates the phonotactics of lexical items, the OCP-Place also plays an active role in dynamic phonological processes, such as triggering consonant place alternations in Cantonese language games (Yip 1982; 1988) and blocking the application of vowel syncope in Afar (McCarthy 1986).

1.1 The phonetic substance of OCP-Place

Since the OCP-Place effects are found in the phonotactics and phonological processes of a wide range of languages, many functional explanations have been proposed to account for them. Ohala's (1981; 1993) hypercorrection theory states that dissimilation sound change will take place when listeners assume that phonologically intended repeated features are coarticulatory and hence cancel each other out in perception. Boersma (2000) argued that the OCP is the result of a tendency to merge two acoustically similar sounds as a single percept during the perception of surface auditory forms. Similarly, Frisch (2004) argued that the encoding of repeated similar segments requires sequential activation and inhibition of the same units, which is undesirable and may lead to distortion in perception. The perceptual distinctiveness of two words also partially depends on the similarity of their constituent segments: two words both containing ejectives ([k'ap'i] ∼ [k'api]) were less reliably distinguished than word pairs with only one ejective ([k'api] ∼ [kapi]) in an AX discrimination task (Gallagher 2010). Such sequential encoding is at play in production planning as well. Dell et al. (1997) argued that there is a turn-off function that deactivates each gesture after it is articulated. When two similar segments are placed in close proximity, they will show competition in articulatory planning because the deactivation mechanism slows down the planning of the latter segment. Garrett & Johnson (2013) discussed three types of common speech errors: interchange (**sn**ow **fl**urries → **fl**ow **sn**urries), anticipation (**r**eading **l**ist → **l**eading **l**ist), and preservation (**w**aking **r**abbits → **w**aking **w**abbits). They further argued that these errors are all triggered by the confusion of articulatory plans due to adjacent phonetically similar targets. Complete segmental identity may also elicit this difficulty and lead to speech errors and repairs in production (Walter 2007). In other words, repetition of gestures is articulatorily difficult, because it involves sustained activity without a resting period for the involved articulators and requires rapid reversal of an articulator's trajectory. This production-based account for similar-consonant co-occurrence restrictions is called the *biomechanical repetition avoidance hypothesis*.

It is important to note that articulatory-based explanations discussed above are more applicable to the OCP on consonant place and manner features and may not extend to other types of OCP effects observed in the world's languages. For tonal OCP, for example, since the production of different tones involves the same articulator — the vocal folds, the articulatory motivation of avoiding the same articulator in a brief time frame does not apply[1]. Consequently, from the perspective of phonetic substance, different subtypes of the OCP (e.g., OCP-Place vs. OCP on tonal or laryngeal features) may be grounded in distinct functional mechanisms. The current study focuses on the role of OCP-Place in phonological learning and its influence on phonological typology.

[1] Studies on the learnability of tonal OCP phenomena can be found in Chen (2022), for instance.

Another characteristic observed in OCP patterns across the world's languages is that fully identical consonants are sometimes exempt from OCP-Place co-occurrence restrictions. For instance, in the Semitic languages discussed earlier, identical consonants are permitted to co-occur in stem-final positions, making stems like *samam* 'poison' relatively common in the Arabic lexicon. Similarly, in Muna, while roots with homorganic consonants are underrepresented — indicating an OCP-Place effect — identical consonants are nonetheless allowed to co-occur freely within a root (Coetzee & Pater 2008). At first glance, such cases appear to violate the OCP. From the articulatory perspective, adjacent identical consonants would require significant effort, hence we expect such sequences to be ruled out by the OCP. However, from a more theoretical perspective, phonological analyses interpret these repeated consonants as being derived from a single underlying consonant that surfaces as two identical segments through node spreading (McCarthy 1986; Berent et al. 2001). Under this analysis, only one consonantal target exists at the underlying level, thereby avoiding a true violation of the OCP.

While some languages tolerate surface identity, others adhere strictly to the OCP and treat total identity as an extreme case of similarity. A case in point is Japanese *rendaku*, a process in which the initial obstruent of the second morpheme in a compound word becomes voiced. This voicing process is more likely to occur when the morpheme boundary features identical onset obstruents (Tanaka 2023). Unlike Arabic and Muna, Japanese does not permit identical consonants to bypass OCP restrictions and actively prevents their co-occurrence through phonological alternations. Thus, whether total identity is tolerated or avoided appears to be a language-specific lexical property (Graff 2012). Furthermore, perception experiments suggest that identical and similar consonants create comparable levels of perceptual distinctiveness (Gallagher 2010; Woods et al. 2010). Therefore, despite the distinct behavior of totally identical segments and the unique learning biases they may elicit, we continue to treat them as a special case of the OCP, rather than as an exception to it.

In sum, functional pressures make adjacent similar (or identical) segments vulnerable to misperception and misproduction. These types of phonetic substance may affect synchronic learning biases to modulate language acquisition. These biases may facilitate the learning of those OCP-abiding processes, and discourage the learning of the OCP-violating ones. This difference in learning efficiency, then, leads to the typological frequency difference between patterns that abide by the OCP and patterns that violate it, as easier-to-learn patterns without OCP violations are more prone to be retained during language transmission. We focus on the learning bias stemming from OCP-Place in this paper.

1.2 The nature of learning biases

There is an ongoing debate on what types of biases are operative during the learning process. As we have seen, typologically common phonological patterns, like the OCP, are often phonetically natural, i.e., motivated by perceptually or articulatorily grounded principles (Archangeli & Pulleyblank 1994; Beguš 2020). A substantive bias explanation has been proposed for these typological patterns, which posits that there are learning biases that facilitate phonetically natural patterns and hinder unnatural ones (Wilson 2006; Hayes & White 2013). Many AGL studies returned positive results on the effect of substantive bias, because participants in the experiments did learn phonetically natural patterns better than their unnatural counterparts (Wilson 2006; White 2017; Martin & Peperkamp 2020). But there are also many studies that yielded null results when comparing the learning of natural and unnatural patterns (Skoruppa & Peperkamp 2011; Glewwe 2022).

Attempts have been made to address the uncertainty of substantive bias in AGL experiments, particularly by examining the conditions under which such biases are most likely to arise (see Zheng & Do (2025) for an extensive review). One proposed factor is the existence of a cue-robustness threshold, which determines whether a substantive bias can effectively influence learning in experimental settings. For example, Glewwe (2019) found that participants were more likely to generalize word-final voicing contrasts to word-initial positions in artificial phonotactic learning tasks, while the reverse direction showed less generalization. This asymmetry aligns with the substantive bias hypothesis, because in typology, the presence of a voicing contrast word-finally implies its presence word-initially. However, in a parallel study, Glewwe (2022) tested whether learners would generalize stop consonant major place contrasts from word-final to word-initial position. In this case, both natural and unnatural patterns were learned equally well, providing no evidence of substantive bias. The author proposes that the magnitude of the phonetic precursor may influence whether a substantive bias emerges: the perceptual distinction between voiced and voiceless stops is more salient than that between different places of articulation, making the former more likely to produce a detectable learning bias.

Another important factor in this investigation is the design of the experimental tasks. Research has shown that the effects of substantive bias are more likely to emerge when the test phase involves a production task compared to a perception task. Martin (2017), for instance, conducted several AGL experiments to compare the learnability of harmonic versus disharmonic vowel co-occurrence patterns, employing both perception- and production-based testing. Overall, participants demonstrated lower accuracy in the forced-choice perception tasks, and the learnability asymmetry between harmonic and disharmonic patterns was also less pronounced compared to the production task. The author argues that forced-choice tests, which present both correct and incorrect forms, potentially introduced confusion due to decision uncertainty or misperception. In contrast, production tasks allowed participants to express their internalized generalizations more freely, making phonetic naturalness effects more likely to surface (Glewwe 2022). Do & Havenhill (2021) argue that production-based training sessions engage the articulatory system, highlighting the ease or effort tied to natural vs. unnatural patterns. Thus, substantive bias is more likely to emerge in production-based tasks. Furthermore, the phonetic precursor of vowel harmony is argued to stem from vowel-to-vowel coarticulation: producing similar vowels in sequence facilitates smoother and more efficient transitions between sounds (Ohala 1994; Finley 2017). If the source of vowel harmony lies in articulatory ease, it follows that its associated substantive bias would be more pronounced in production-based tasks.

Another type of bias revealed by AGL experiments is the advantage of learning structurally simpler patterns over more complex ones (Pycha et al. 2003; Kuo 2009). For example, if a phonotactic pattern involves a larger number of features to be adequately described, it is found to be harder to learn compared to patterns involving fewer features (Kuo 2009). Unlike the conflicting results on substantive bias, structural bias effects are reliably found in virtually all AGL studies that investigate them (Moreton & Pater 2012a). Presumably, structural bias reflects speakers' domain-general cognitive ability, hence it is applicable to a wider range of pattern detection and learning tasks (Moreton & Pertsova 2014). Substantive and structural biases may sometimes be difficult to differentiate. For instance, perceptually more distinct sound pairs are less likely to be involved in alternation (Hayes & White 2015) — a potentially substance-based generalization. But this could also be due to the greater number of phonological features that are involved in the alternation — a structurally based generalization. Since the effect of substantive bias is weaker than structural bias, some researchers have expressed doubts on the validity of the former and suggested that much of the substantive effect can be reanalyzed and accounted for by structural properties. To instantiate a substantive bias, then, one needs to control for structural complexity or include both complexity and naturalness in the experimental data modeling (e.g., Prickett 2018).

The goal of the current study is to demonstrate that the typological frequency of OCP-Place is due to a synchronic substantive bias. In the previous section, we discussed how OCP-Place is rooted in both perception and production. We argue that the phonetic groundings of OCP-Place take effects as domain-specific cognitive predispositions and can facilitate the learning of similarity avoidance patterns during language transmission. As a consequence, these patterns are widely attested in the world's languages. AGL is an ideal experimental paradigm to test how speakers' learning procedure is biased. To stack the deck against our hypothesis, we adopted a forced-choice perception task in the testing phase, thereby strengthening the credibility of any positive result. The results of our AGL experiment suggest that OCP-based phonotactics does receive an advantage in acquisition compared to other structurally comparable and arbitrarily conditioned patterns. Such learning benefits can only be fully explained by a synchronic substantive bias account.

1.3 Previous studies on the learnability of OCP-Place

The OCP biasing effect has been found in some previous AGL studies. For example, Boll-Avetisyan & Kager (2014) created artificial languages composed of concatenated CV syllables without pauses. The shape of such speech stream followed the pattern of …$P_1P_2TP_1P_2T$…, where P stands for a CV syllable with a labial onset and T for a CV syllable with a coronal onset. In other words, a coronal-initial syllable was always followed by two labial-initial syllables. These languages were presented to Dutch native speakers to test their ability to learn word segmentation patterns. Participants' two-alternative-forced-choice nonword acceptability judgments indicated that PTP sequences were preferred over PPT or TPP sequences as well-formed words in the artificial languages. In a follow-up study, Boll-Avetisyan & Kager (2016) further found more fine-grained effects in this type of task: the strength of Labial-OCP in Dutch lexicon varies among different consonants (e.g., the co-occurrence of two adjacent [p]s is actually over-attested, while [b] pairs are highly under-attested), and participants' responses were influenced by these detailed distributional patterns in the lexicon. They interpreted these results as an OCP bias effect: since the Dutch lexicon shows an OCP-labial effect (words that contain two labial consonants are, overall, underrepresented in Dutch), speakers generalize this knowledge onto this task and are thus more likely to put a word boundary in between two labial-initial syllables.

Brohan (2014) designed a series of experiments to test whether OCP-based consonant co-occurrence restrictions can be learned in an AGL setting. Native English speakers were recruited and trained with CVCV items (C = {p b f v t d s z k g} and V = {i e a o u}) and later were asked to offer acceptability ratings on a novel set of items to see if they sounded like words that could belong to the language they heard in the exposure. In Experiment 1, there were two experimental conditions: for the OCP condition, none of the training items contained consonant pairs with homorganic places (e.g., *pVbV, *tVsV, *kVgV); and for the arbitrary condition, none of the training items allowed labials and coronals to co-occur (e.g., *pVtV, *sVfV). The arbitrary phonotactic generalization was selected so that the number of illegal consonant pairs was similar to the number for the OCP condition. The test items contained familiar words that already occurred in the exposure and novel words that either conformed to the phonotactics reflected in the exposure or not. Statistical analysis showed that the acceptability difference between legal and illegal test items was significantly larger in the OCP condition than in the arbitrary condition, suggesting that participants in an AGL task could better learn an OCP-based phonotactic restriction as compared to learning an arbitrary phonotactic restriction. In Experiment 2, Brohan further tested whether the training of an OCP pattern realized by two places of articulation could generalize to a third place, but the results were only marginally significant and inconclusive. Altogether, this study provides some preliminary evidence that an OCP-based phonotactic pattern can be better learned than an arbitrary pattern. However, one concern about this study is the source of this learning advantage: since OCP-Place is already statistically significant in the English lexicon (Berkley 1994), the preference for an OCP-conforming pattern could have originated from L1 influence. The current study circumvents this issue by recruiting native Mandarin speakers as the participants in the AGL experiment. Mandarin does exhibit certain OCP effects in its syllabary, such as co-occurrence restrictions between consonants and vowels, and between vocoids, as noted by Yi & Duanmu (2015) and Gong & Zhang (2021). Nevertheless, it differs from many other languages, including English, in that consonantal OCP-Place effects do not extend beyond the monosyllabic level in its lexicon. In a database containing 266,642 word tokens (Li & Xu 2001), CVC sequences (including both CVC.V and CV.CV) where the two consonants are homorganic are neither over- nor under-represented (Boll-Avetisyan 2012). Therefore, Mandarin native speakers serve as ideal participants for AGL experiments exploring the consonantal OCP effects across syllables, as participants' sensitivity to the target patterns cannot have originated from their L1 experience.

In addition to AGL studies, the OCP-Place effect was also attested in a wide range of other experimental paradigms. For example, Hayes & White (2013) found that for the phonotactic constraints assigned with equal weights by the lexical statistics, the phonetically natural ones, for example, ones that conform to the Sonority Sequencing Principle or homorganicity/heterogenicity requirements, had more effect on speakers' nonword judgments, while the unnatural ones had little to no effect. Some of the natural constraints they investigated, e.g., no two labials in English onset clusters, are instances of OCP-Place. Other nonword judgment tasks showed that native Arabic speakers' acceptability ratings on CVCVC nonce verb roots were sensitive to the similarity of the consonants in the roots; the more similar the consonants were, the less acceptable the nonword sounded (Frisch & Zawaydeh 2001). This finding indicates that the phonotactics of Arabic features an OCP-based constraint, and it can guide Arabic speakers to offer gradient judgments on novel forms. Similar results were found in lexical decision tasks of Arabic (Gwilliams & Marantz 2015), Dutch (Kager & Shatzman 2007), Hebrew (Berent et al. 2001), and Mandarin (Gong & Zhang & Fiorentino 2023). All of these studies compared how fast participants rejected two types of nonwords: OCP-violating vs. non-violating, and indeed the items violating the OCP showed faster reaction time. These psycholinguistic studies demonstrated that the OCP is also active in speakers' online speech processing and word recognition.

2 Experiment: The learnability of OCP-based phonotactics

Natural language experiments like nonword judgment and lexical decision only provide evidence for the role of OCP as a phonetically based principle in speakers' phonological grammar. To answer the question why OCP restrictions frequently exist in the world's languages, we need to directly examine how these patterns are acquired. The current study addresses the learning issue using the artificial grammar learning paradigm.

Previous AGL experiments have shown that both children and adults can acquire many different types of linguistic structures through statistical learning from the input (Newport 2016). The current experiment was conducted to compare the learning of an OCP-based pattern with other types of phonotactic patterns. Based on Clements & Hume's (1995) feature geometry on the internal featural organization of sounds, the place of articulation of a consonant can be categorized into three levels: labial, coronal, and dorsal, known as major consonant places. We focused on these major consonant places in CVCV nonce words. In our baseline OCP phonotactic pattern, the two consonants in the nonce words necessarily disagree in their major places of articulation.

The first type of phonotactic pattern to be compared with the OCP is an arbitrary pattern that is formally more complex than the OCP-based pattern. Phonological grammar is able to employ alpha notation or algebraic representations to capture the recurrence of identical feature specifications (Berent et al. 2012; Moreton 2012). The OCP-Place effect can thus be formalized as a co-occurrence constraint *[place_i][place_i], which assigns a violation for each recurring major place of articulation within a domain. While the OCP pattern in our design can be described by only one alpha notation, the arbitrary pattern necessarily requires more than one of such variables to be represented. As reviewed by Moreton & Pater (2012a), most of the AGL studies report that structurally simpler patterns are better learned than patterns referring to more features. The arbitrary pattern in this experiment serves as the learnability baseline. In any case, structurally simpler patterns should be more learnable than this baseline.

The second pattern to be compared with the OCP is consonant major place harmony, specifically, the agreement of place features within a root. Structurally, OCP-Place and major place harmony are comparable in complexity, as both involve constraints on a single feature. The formalization of these two patterns can be stated as *[place_i][place_i] vs. [place_i][place_i]. However, from a typological perspective, major place harmony is exceedingly rare in the world's languages, with only a few isolated cases reported, such as Ngbaka Minagende (Danis 2019). In contrast to the widespread avoidance of adjacent identical place features (as in OCP-Place), the opposite pattern — enforcement of identical place features — is virtually absent in natural language phonologies. While consonant harmony phenomena are attested in various languages, (e.g., sibilant harmony in Samala, nasal harmony in Yaka), harmony specifically targeting major place features is almost entirely lacking (Hansson 2010; Rose & Walker 2011).

The near absence of consonant major place harmony poses a challenge for phonological theory. From a functional perspective, major place distinctions (e.g., labial vs. coronal) are acoustically salient and play a crucial role in lexical contrast. A harmony pattern that neutralizes these distinctions would risk increasing lexical ambiguity, thereby being disfavored functionally. Moreover, since major place features rely on distinct articulators, there are no natural phonetic precursors — such as coarticulation — that would promote assimilation in place of articulation. Consequently, the preservation of distinct place features is generally preferred (Rose & Walker

2004). Regarding the exceptional case of Ngbaka, Hansson (2020) observes that the reported harmony may instead reflect coronal-specific agreement or similarity-based restrictions, driven by high observed/expected (O/E) ratios of specific labials and labial-velars, rather than genuine major place harmony.

Previous research on learning biases often compares vowel harmony and disharmony, highlighting that these patterns are structurally equivalent but differ in phonetic grounding. A parallel situation emerges in consonant co-occurrence patterns: OCP-Place and major place harmony both involve constraints on a single feature, yet it is the harmony pattern that lacks phonetic motivation and typological support. Therefore, comparing the learnability of these two patterns provides another ideal test case for investigating the effects of substantive bias in phonological learning. While OCP-Place is phonetically grounded and typologically robust, major place harmony is not supported by articulatory or functional considerations and represents a near-typological gap.

The two patterns in the current experiment, OCP-Place and major place harmony, are matched in terms of structural complexity, both only referring to one feature (that is, place of articulation) but differing in directions (a dissimilation pattern *[$place_i$][$place_i$] vs. an assimilation pattern [$place_i$][$place_i$]). The matching complexity ensures that if any learning asymmetry emerges from these two patterns, it would serve as evidence for a substantive bias, which in turn suggests that the learning advantage of the OCP patterns reflects a domain-specific linguistic effect.

In sum, the current experiment compares the acquisition of an OCP-based phonotactic pattern with a consonant major place harmony pattern and an arbitrary pattern by examining participants' behavior in the learning of the three artificial languages. In the OCP language, the consonants in the words all had disagreeing major places of articulation; in the harmony language, all consonants agreed in terms of major place of articulation; and in the arbitrary language, the distribution of the consonants followed an arbitrarily determined complex pattern. Participants were randomly assigned to learn either of the three languages. After hearing all words in the exposure phase, participants offered acceptability judgments for novel words that are distinct from the exposure words.

2.1 Participants

Participants were 90 native Mandarin speakers (32 men, 58 women, mean age = 19.89, age SD = 1.66) without any speech or hearing problems. They were randomly assigned to the three conditions of this experiment: OCP, Harmony, and Arbitrary. A small monetary incentive was provided for participating in this experiment. Participants were recruited via online social media (WeChat groups, etc.). Most of the participants were college students studying in mainland China.

2.2 Materials

All stimuli in the exposure and test phases were CVCV nonwords of Mandarin, where C was drawn from the set of [p^h p f k^h k x] and V from [ei a ou u]. Stimuli were created from all possible combinations of these consonants and vowels excluding identical pairs of vowels. The positional frequencies of vowels were balanced so that no specific vowel appeared especially frequently in any specific contexts, while the distributions of consonants were the critical variables for manipulation. All CV combinations except [k^{h}ei] and [pou] were existing syllables in Mandarin, although [k^{h}ei] appears in some colloquial forms. All vowels in the stimuli carried the falling tone, because two consecutive falling tones are the most frequent tonal sequence in Mandarin disyllabic words[2] (Lin 2017). This means that some CV combinations were tonal gaps of Mandarin (e.g., [fou3] with the low-dipping tone is a word 'to deny' in Mandarin, but [fou4] with the falling tone is missing). Moreover, resulting CVCV combinations that formed existing Mandarin words (e.g., [fu4fei4], 'to pay the fee') were excluded from the experimental stimuli.

The consonants in the artificial languages did not contain coronal sounds, so that in all 36 possible consonant pairs, half of them agreed in the place of articulation and half did not agree. In the exposure phase of the OCP condition, the stimuli consisted of the 18 disagreeing consonant pairs only (e.g., p^h…k, x…p, etc.), and in the Harmony condition, the stimuli consisted of the 18 agreeing pairs only (e.g., p^h…p, x…k, etc.). For the Arbitrary condition, 9 agreeing (p^h…f, p…p^h, p…p, f…p^h, f…p, k^h…x, k…k^h, k…k, k…x) and 9 disagreeing (p^h…k^h, p^h…x, p…k^h, f…k^h, k^h…p^h, k^h…p, k…p, x…p^h, x…f) consonant pairs were arbitrarily selected to appear in the exposure stimuli. Therefore, consonant pairs in the OCP and Harmony conditions abided by the generalizations that they do or do not agree in place, but no clear phonotactic generalization on consonants can be made in the Arbitrary condition.

In the exposure phase, each consonant pair was assigned five vowel pairs, giving 90 exposure stimuli for each condition. The 90 stimuli in each condition were presented to the participants randomly, with each participant receiving a different randomized order. The test phase included 36 words that were different from any of the stimuli in the exposure (i.e., novel vowel pairs) with all 36 possible consonant pairs. Half of these test words' consonant pairs were attested in the exposure and the other half were not, while exactly which pairs were attested depended on the experimental condition. Participants were asked to offer binary acceptability judgments on these test words. They were presented using the same randomization scheme as the exposure stimuli. The test words were the same across the three experimental conditions. The complete lists of all exposure and test stimuli are provided in Appendix 1. All stimuli are transcribed in Pinyin in the appendices.

[2] Adjacent repeated tones seem to be a violation of the OCP. However, if we analyze the falling tone, a contour tone, as a sequence of level tones (high plus low) (Duanmu 1994), then two adjacent falling tones do not violate the OCP. Chen (2022) named the first case as OCP-Unit, and the second case as OCP-Terminal.

Words in the artificial languages were recorded by a male native Mandarin speaker. The recording session was conducted in an office with minimal background noise using a condenser microphone. Other than applying amplitude normalization to 70 dB using Praat (Boersma & Weenink 2017), the recorded stimuli were not modified in any other way.

2.3 Procedure

All experiments in the current study were conducted using Gorilla, a web-based platform for running online experiments (Anwyl-Irvine et al. 2020). Before the experiment, participants were presented with an Information Statement and consent was requested to participate in this study. At the beginning of the experiment, participants were told that they will be learning a novel language in the exposure phase. Then, in the test phase, participants would decide whether some new words sounded like they could belong to this novel language that the participants had been listening to earlier during the exposure phase. For each trial in the exposure phase, participants heard the stimulus and were encouraged to say it out loud. Then they clicked on a button to move on to the next trial. The next-trial button would not appear until the stimulus finished playing. In the test phase, trials were presented in a similar way, except that after hearing the stimulus in each trial, participants were asked to make a binary acceptability judgment by clicking either the 'yes' or 'no' button on the screen. No orthographic forms were provided throughout this experiment. Participants' binary decisions were recorded for analysis.

After finishing the main experiment, participants were redirected to fill out a questionnaire about their basic demographic information and language background. The whole experiment lasted for around 15 minutes.

2.4 Predictions

First, we expect a main effect of consonant pair attestedness on the acceptability rates of the test stimuli, since it represents the effect of training: participants in all three conditions should be more likely to accept novel words with consonant pairs they have heard in the exposure. More importantly, we predict that there should be a significant interaction between condition and attestedness, because this would indicate that the training effect (the acceptability difference between attested and unattested test words) is different across the three experimental conditions. We specifically predict that a substantive bias will make the OCP pattern more learnable than the Harmony pattern. This is because, unlike consonant major place dissimilation, place harmony is nearly absent in typology and lacks clear phonetic grounding. Therefore, any observed learnability advantage can be attributed to the influence of substantive bias, rather than being confounded by structural bias, as both patterns are matched in structural complexity. Structural bias effects, as reported in most AGL experiments, can be assessed through the performance in the Arbitrary condition. The training effect should be weaker in the Arbitrary condition compared to the other two structurally simpler phonotactic patterns.

2.5 Results

Among the 90 participants, 23 of them (9 from the OCP condition, 7 from the Harmony condition, and 7 from the Arbitrary condition) gave a ‘yes’ response to all test words. Their results were excluded from the data analysis because they either did not learn any pattern or simply did not pay attention to the experiment. This is verified by the fact that participants who gave ‘yes’ responses to all trials responded faster in the test phase (mean log-transformed reaction time of all trials = 6.535 ms, sd = 0.865) than other remaining participants (mean = 6.824 ms, sd = 0.958). A two sample t-test suggests that this difference in mean reaction time is significant ($t = -5.101$, $p < 0.0001$). Since the experiment was conducted outside of a laboratory setting, the high attrition rate might be due to the lack of attention to the tasks in a completely self-paced environment.

The mean endorsement rates (percentage of ‘yes’ responses for the test words) of participants for different conditions and different consonant pair attestedness are shown in **Figure 1**.

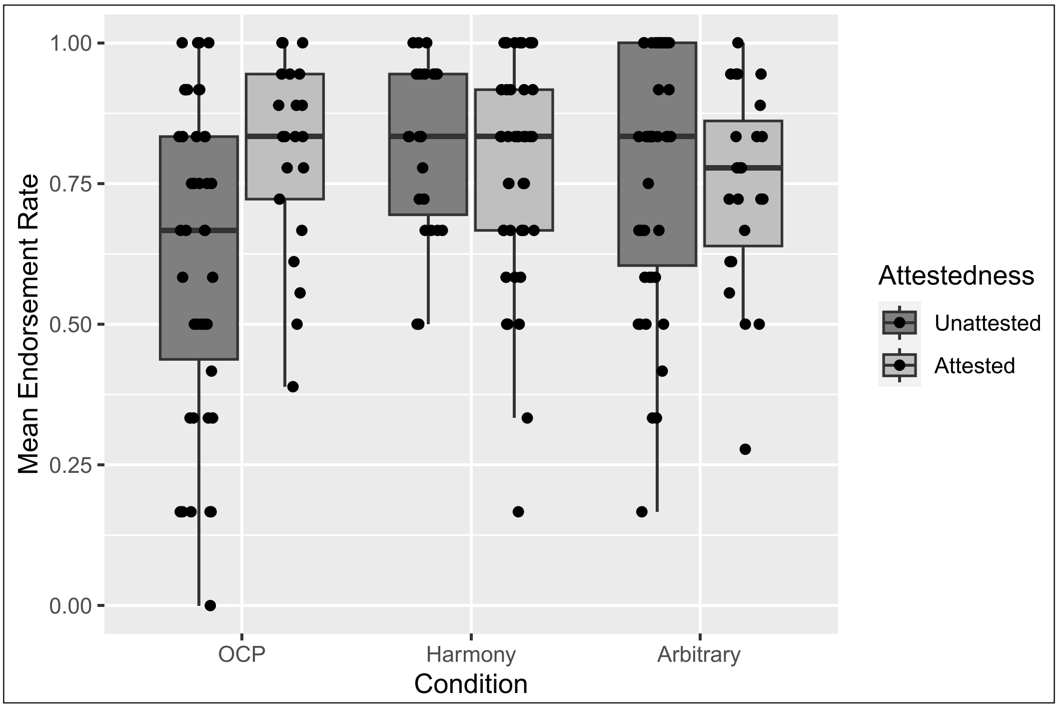

Figure 1: Proportion of ‘yes’ responses to the test stimuli grouped by condition and consonant pair attestedness. Boxes indicate the range between the first and third quartiles. Whiskers delimit the minimum and maximum data points, excluding any outliers.

Logistic mixed-effects models were fitted to participants’ yes/no acceptability responses (yes = 1, no = 0) on the test words using the *lme4* package (Bates et al. 2015) in the R software (R Core Team 2018). Fixed effects included a three-level factor experimental condition (OCP, Harmony, and Arbitrary), a binary factor consonant pair attestedness (whether the consonant pair occurs in the exposed language), and their interactions. All categorical predictors were dummy-coded. The baseline for attestedness was Unattested and the baseline for the experimental condition was Arbitrary. For the random effects to be included, we started with the most complex structure, including intercepts for participants and test words, slopes on experiment condition, test word attestedness and their interactions, but the model failed to converge. We then subtracted random slopes from the model and the maximal model that converged was the one with random intercepts for participant and test word, and random slopes for attestedness by test item. Random slopes were not attempted for participant because participants in this design were nested within conditions. The maximal model's parameter estimates are shown in **Table 1**.

	Estimate	**Std Error**	**z value**	**Pr(>|z|)**
(Intercept)	1.4960	0.2518	5.9408	0.0000***
Attested	–0.2955	0.2713	–1.0891	0.2761
OCP	–0.5871	0.2982	–1.9689	0.0490*
Harmony	0.2748	0.3582	0.7670	0.4431
Attested:OCP	1.1089	0.3220	3.4439	0.0006***
Attested:Harmony	0.0005	0.3881	0.0012	0.9990

Table 1: The maximal model for endorsement rates.

The significant effect of OCP indicates that for the unattested items, participants in the OCP condition gave significantly fewer 'yes' responses than participants in the Arbitrary condition. More importantly, there is a significant interaction between item attestedness and experimental condition. The acceptability difference between the two types of test items was significantly higher in the OCP condition than in the Arbitrary condition, and the difference was not significantly different among participants in the Harmony condition and the Arbitrary condition. To make the comparison between OCP and Harmony conditions more explicit, we switched the baseline of experimental condition to OCP in the model, and the interaction term *Attested:Harmony* was still significant ($z = -2.823$, $p = 0.0048$). We therefore conclude that the OCP pattern was better learned than both the Harmony and the Arbitrary patterns.

3 Discussion

Using the AGL paradigm, we designed phonotactic patterns with varying restrictions on consonant co-occurrences and examined how speakers acquired them in a short training phase. The different learning efficiencies across the three artificial phonotactic patterns added evidence supporting that phonological learning is substantively biased.

The generalizations for the phonotactic patterns in the OCP and Harmony conditions both only refer to one feature: *[place_i][place_i] and [place_i][place_i], respectively; these two patterns are thus matched in terms of complexity. By comparing the outcomes of these two patterns, we found a learning asymmetry such that the OCP pattern is better learned than the Harmony pattern. This constitutes strong evidence for a substantive bias effect in phonotactic learning, because the difference between these two patterns only lies in phonetic substance and typological attestedness: OCP-Place has solid perception and production motivations and is widely attested in the world's languages, while consonant major place harmony lacks such motivation and is a typological gap.

Additionally, both the Harmony and Arbitrary patterns were poorly learned, indicating a lack of structural bias effects. The major place harmony pattern can be characterized as an enforcement constraint of the form [place_i][place_i], which involves only a single feature. In contrast, the pattern in the Arbitrary condition requires reference to multiple features to be formally specified.

3.1 Identity effect

As previously discussed, identical consonants receive special treatment in certain languages. Prior AGL studies have also shown that learners can quickly and accurately detect and internalize patterns based on identity relations (Marcus et al. 1999; Gallagher & Coon 2009; Gallagher 2013; Linzen & Gallagher 2017). Since the training phase of the OCP condition did not include any items containing two identical consonants, participants may have inferred such a ban on identical consonants during the training process, in addition to a more general OCP-Place pattern. In the test phase, they may be highly inclined to reject any test items featuring two identical consonants. As a result, the learnability advantage found in the OCP condition might have originated from the dispreference for identity, not for the OCP *per se*.

To address this concern, we singled out the test items with two identical consonants in the OCP condition (6 out of 36) and compared their acceptability ratings with other types of trials (specifically, 12 unattested items with two same-place, but non-identical consonants, and 18 attested items with two different-place consonants) (See **Figure 2**). We now call these three types of trials identical, similar, and dissimilar items, respectively. We did find a difference in acceptability ratings between identical items ($\bar{x}$ = 63%, σ = 0.35) and similar items ($\bar{x}$ = 83%, σ = 0.18). Results of a two sample t-test suggested that participants' average ratings on identical items were significantly lower than similar items ($t = -2.815$, $p = 0.0074$). Therefore, it seems that the learning asymmetry may indeed be driven by the dispreference for two identical consonants. But a closer look at the data indicates that the pattern is much more nuanced.

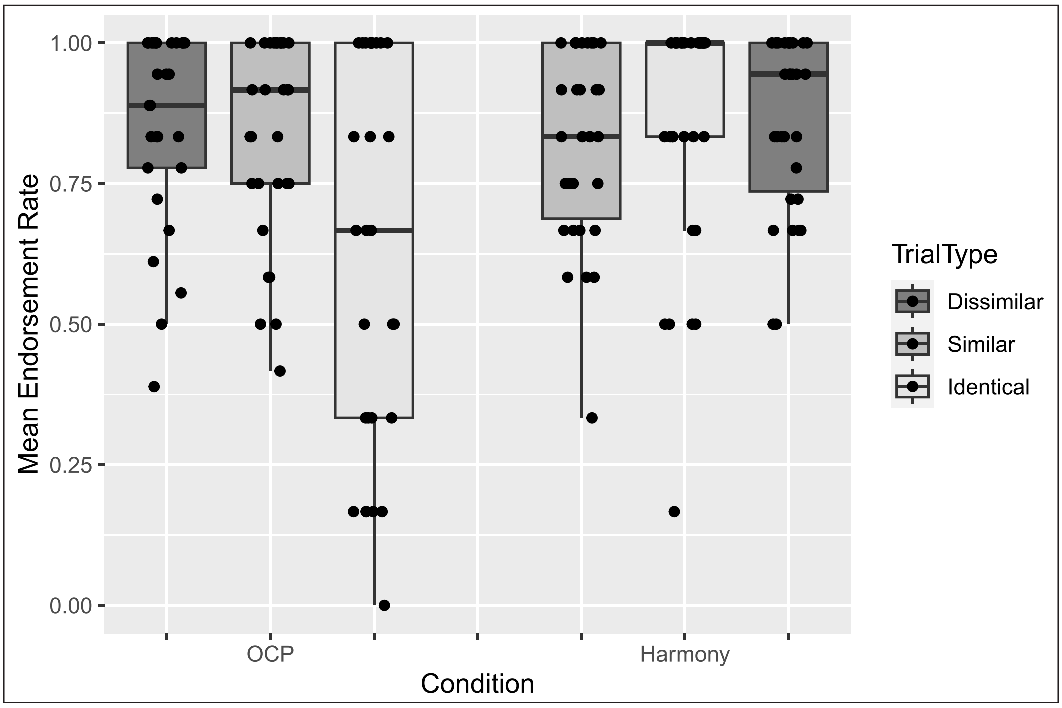

Figure 2: Proportion of 'yes' responses to the test stimuli in the OCP and Harmony conditions grouped by the type of consonant pair. Stimuli attestedness is shown in the top labels.

Because the key argument for a substantive bias comes from the different learnability between the OCP and the Harmony conditions, we did the same break-down analysis for the Harmony condition to see how the identity effect was learned there (also see **Figure 2**). Notice that in this case, the attestedness of the test items was flipped: the 18 dissimilar items are now unattested, and the 12 similar and 6 identical items are attested. Meanwhile, an unexpected learning effect emerged in the Harmony condition: unattested dissimilar items received higher ratings than the attested similar items.

We can ask the question whether the dispreference of identity in the OCP condition and the preference of identity in the Harmony condition are comparable in size. Specifically, in the OCP condition, the learning effect triggered by identity was represented by the rating difference between the attested dissimilar items and the unattested identical items (the first and third columns of the left panel in **Figure 2**). Accordingly, in the Harmony condition, identity effect was represented by the rating difference between the attested identical items and the unattested dissimilar items (the second and third columns of the right panel in **Figure 2**). This comparison reveals whether participants' identity-related learning behaviors differed between the two conditions.

For this purpose, we prepared a trimmed dataset with the ratings of the 18 dissimilar and the 6 identical items from the OCP and Harmony conditions only, and built a mixed-effects logistic model. This model now has two independent binary categorical variables: *Attestedness* and *Condition*, their interaction, and random intercepts for participants and items. The summary of this model is shown in **Table 2**. The results suggest that the rating differences between identical and dissimilar items in the OCP condition were significantly larger than those in the Harmony condition, as indicated by the significant interaction term ($z = 3.512$, $p = 0.0004$). We therefore conclude that the learnability concerning identity is higher in the OCP condition than in the Harmony condition. When the input lacked identical consonant pairs, learners efficiently detected this pattern and strongly penalized such items during the test phase. In contrast, when the input included identical pairs (along with other pairs exhibiting place harmony), learners did not differentiate them as clearly from dissimilar pairs in the test phase.

	Estimate	**Std Error**	**z value**	**Pr(>|z|)**
(Intercept)	2.4740	0.3389	7.301	<0.0001***
Attested	–0.1126	0.3112	–0.362	0.7175
OCP	–1.6266	0.4964	–3.277	0.0011**
Attested:OCP	1.7214	0.4902	3.512	0.0004***

Table 2: The model for endorsement rates on dissimilar and identical items only.

Let's now turn to participants' responses to similar items. Recall that in both conditions, the ratings for these items were somewhat unexpected. In the OCP condition, similar items were not attested during training, yet their average ratings ($\overline{x} = 83\%$, $\sigma = 0.18$) were only slightly lower than dissimilar ones ($\overline{x} = 86\%$, $\sigma = 0.17$) in the test phase (the first and second columns of the left panel in **Figure 2**). The boxplots in the figure showed that the median of similar item ratings was even higher than the median of the dissimilar group. In the Harmony condition, similar items were attested in training, but received lower average ratings in test (the first and third columns of the right panel in **Figure 2**), exhibiting a reversive learning effect. This raises the following question: whether participants ignored the absence of similarity (among the absence of identity) in the OCP condition less than they ignored the presence of similarity (among the presence of identity) in the Harmony condition. This contrast may point to another potential source of the difference in learnability between the two artificial phonotactic patterns.

To investigate the learning effects driven by similar-but-non-identical items, we selected the data of the 16 dissimilar and 12 similar items from the OCP and Harmony conditions only, and built a mixed-effects logistic model based on this trimmed dataset (**Table 3**). The structure of this model resembles the previous one. The results return a marginal effect of the *Attestedness:Condition* interaction term ($z = 1.920$, $p = 0.0548$), indicating that the two structurally comparable phonotactic patterns still exhibit a learnability difference after identical items are removed; in particular, the participants were less prone to ignore the absence of similarity in the OCP condition than the presence of similarity in the Harmony condition. These different strategies in treating similar-but-non-identical consonant pairs still constitute evidence for an OCP-based learning bias.

	Estimate	**Std Error**	**z value**	**Pr(>|z|)**
(Intercept)	2.3699	0.3133	7.564	<0.0001***
Attested	–0.3073	0.1961	–1.567	0.1171
OCP	–0.1558	0.4467	–0.349	0.7273
Attested:OCP	0.5432	0.2829	1.920	0.0548

Table 3: The model for endorsement rates on dissimilar and similar items only.

In the analyses above, we divided the learnability difference between the OCP and Harmony conditions into two components: i) a dispreference for identity, and ii) a dispreference for similarity independent of identity. Both contributed to the learning advantage of the OCP-Place pattern over the consonant major place harmony pattern. The identity effect appeared as, when identity was absent from the input (OCP condition), learners penalized identical pairs strongly at test; when identity was present in the input (Harmony condition), identical pairs received only a modest advantage over dissimilar pairs. Thus, identity had a larger impact in OCP condition than in Harmony condition. The similarity effect, by contrast, appeared as learners being less willing to ignore the absence of similarity when identity was absent (OCP condition), compared to how readily they ignored the presence of similarity when identity was present (Harmony condition). This suggests that participants discounted the presence of similarity in Harmony condition more than they penalized its absence in OCP condition.

The learning asymmetry based on similarity notwithstanding, a large portion of this advantage was indeed induced by the identity items. However, there is phonological evidence to support that identity avoidance is still a special case of the broader OCP-Place constraint. Numerous phonological processes, including Japanese *rendaku* blocking (Kawahara 2006) and Latin liquid dissimilation (Cser 2010), among others, specifically target the co-occurrence of identical consonants; a pattern that is canonically analyzed as stemming from the OCP. In other AGL experiments, identity and similarity were not necessarily treated differently by learners. For example, in Boll-Avetisyan & Kager's (2014) word segmentation experiment, they did not find any segmentation preference distinction between similar (e.g., pVm) and identical (e.g., pVp) sequences. Therefore, we believe that, even though the learnability difference between our OCP and Harmony conditions was partially driven by the dispreference of identity, the dispreference for identity itself found in the experiment also makes a contribution to our understanding of the OCP-Place effects.

3.2 Other nuisance factors

We also considered a number of other nuisance factors that may have influenced participants' acceptability judgments in the test phase. First, participants might judge the two Mandarin segmental gaps [k^hei4] and [pou4] as less acceptable than other existing syllables. We added a binary factor that distinguishes the test stimuli based on whether they contain one of these two segmental gaps into the model shown in **Table 1**, and maximum likelihood ratio test suggested that this factor did not significantly improve the model ($\chi^2 = 0.0113$, $p = 0.9152$). Second, although the vowels in the stimuli were randomly assigned, all participants were trained on the same set of stimuli in each condition. That means the vowel distributions for each participant were also fixed, so that they may have accidentally picked up on vowel co-occurrence patterns as well. If certain vowel pairs occurred less frequently in the training, it might have affected participants' acceptability judgment in the test phase. To rule out this possibility, we looked at the vowel combinations in the training to see if there were discernable patterns, and if so, whether these patterns had an effect on participants' judgments. For example, there were altogether 12 possible vowel pairs (16 pairs based on factorial combinations among [ei, a, ou, u] minus 4 identical pairs) for the CVCV test words, and there were 90 items in the exposure phase. On average, each vowel pair should occur 90 ÷ 12 = 7.5 times, yet the vowel combination […ou…ei] only occurred three times in the exposure items of the OCP condition. If vowel statistics influenced acceptability, we may expect that test items carrying this vowel combination (e.g., [koupei]) would receive lower ratings than other items in the OCP condition. Adding a factor that distinguished test items with the […ou…ei] vowel pair from others did not significantly improve the statistical model for acceptability in the OCP condition ($\chi^2 = 0$, $p = 1$). Third, participants may have accidentally noticed the similarity between the two vowels in the exposure items and make generalizations. We used the feature matrix from Hayes (2009) to define the five monophthong vowels used in the experiment (see **Table 4** below) and counted how many features were different among all the vowel pairs. When the comparison was between a monophthong and a diphthong, the monophthong was repeated twice.

	high	low	front	back	round
i	+	–	+	–	–
u	+	–	–	+	+
e	–	–	+	–	–
o	–	–	–	+	+
a	–	+	–	–	–

Table 4: Feature definition of the monophthong vowels used in the experimental stimuli

For example, for the two vowels [a] and [u], they differ in [high], [low], [back], and [round], so their distance was counted as 4. For the vowel [a] and the diphthong [ei], [a] and [e] differ in [low] and [front], [a] and [i] differ in [high], [low], and [front], so the distance between [a] and [ei] was counted as $(2 + 3)/2 = 2.5$. Using this method, we calculated the vowel distance for all exposure items, and crucially, the mean vowel distances across the three artificial languages in the experiment are the same. Adding this vowel distance factor did not significantly improve the model for acceptability ($\chi^2 = 0.5821, p = 0.446$).

3.3 Failure to learn place harmony

The lack of learning effect for the Harmony pattern is particularly interesting, because long-distance consonant harmony patterns like sibilant harmony can be successfully learned in some AGL studies (Finley 2011; 2012). Meanwhile, arbitrary and typologically unattested sound patterns are also shown to be learnable (Seidl & Buckley 2005; Koo & Callahan 2012a). For example, vowel disharmony is rarely attested in world's languages but it is learnable according to Pycha et al.'s (2003) AGL experiments. For the studies that tested arbitrary phonotactic patterns, e.g., Koo & Callahan (2012b), they often target consonant co-occurrence of a single pair (e.g., [m] always occurs after [s]), so the learning might be quite specific. The remaining question is that, given that both vowel disharmony and consonant major place harmony are typologically rare, and neither pattern has clear phonetic motivation, why did the current experiment show that place harmony is not learned, while other studies find vowel disharmony to be learnable (Pycha et al. 2003; Skoruppa & Peperkamp 2011)? Our hypothesis is that the vowel disharmony learning effects found in earlier studies may have other interpretations. Martin & Peperkamp (2020) used French accented stimuli to test the learnability of vowel harmony and disharmony patterns on English native speakers, and the results showed that vowel harmony was more easily learned than disharmony. They argued that this finding is due to the more phonetic processing of the materials when participants were presented with sounds from a foreign language. Therefore, it is possible that the proper learning of vowel disharmony patterns in earlier studies (Pycha et al. 2003; Skoruppa & Peperkamp 2011) is due to alternative strategies such as memorizing specific vowel dependencies (e.g., suffix [-u] always occurs after a preceding [e]), but in fact, vowel disharmony still receives a learning disadvantage. Finally, it should be noted that the bias against these typologically rare patterns is a 'soft' bias, i.e., it is not the case that consonant place harmony or vowel disharmony is completely unlearnable. They are disadvantaged compared to other typologically well-attested and phonetically natural patterns, but given enough exposure, these patterns may still show learning effects, as many studies have reported the successful learning of arbitrary phonological patterns (e.g., Pycha et al. 2003; Seidl & Buckley 2005; Peperkamp & Dupoux 2007; Skoruppa & Peperkamp 2011).

3.4 Substantive bias

The main findings in our experiment are that speakers learned an OCP-based phonotactic pattern better than an arbitrary consonant co-occurrence pattern referring to more than one features and a consonant harmony pattern based on one single feature as well. While most AGL experiments have shown structural bias effects, substantive bias does not always emerge in these studies (Moreton & Pater 2012b; Glewwe 2019). Our experiment hence contributes an argument for the existence of substantive bias in that the phonetically motivated OCP pattern is better learned than its structural counterpart: consonant major place harmony, which is generally missing in real languages.

The situation described here closely resembles what is attested in the typology of vowel harmony and the related learnability studies, but in an opposite direction. Vowel harmony and disharmony also involve the dis/agreement of one feature only (Harmony: *[αF][-αF]; Disharmony: *[αF][αF]); hence they are comparable in terms of structural complexity. However, in typology, vowel harmony is widely attested and has clear phonetic motivation, while vowel disharmony is almost missing and lacks a good phonetic motivation (Ohala 1994; Rose & Walker 2011). As mentioned in the literature review earlier, AGL studies comparing the learnability of these two patterns revealed a substantive bias effect as well: vowel harmony can be efficiently learned while vowel disharmony exhibits learning difficulties (Martin & Peperkamp 2020; Martin & White 2021). The results of the current study mirror this learning bias, except that in our case, it is consonant major place *disharmony* that receives an advantage in learning compared to major place harmony. Together, the opposite learnability biasing effects between vowel harmony and consonant disharmony (OCP-Place) serve as even stronger evidence for this bias to be substantive rather than structural. Both patterns are essentially dis/agreement on one single feature, yet the direction of the bias runs directly in contrast when the target sounds switch from consonants (pro-disharmony) to vowels (pro-harmony). This demonstrates that the direction of learning bias on a particular agreement structure counts on the phonetic substance of the elements involved.

3.5 The origins of OCP-Place

In discussions of the origins of phonological typology and the innateness of learning biases, two major perspectives dominate. One view emphasizes the role of channel bias, which refers to systematic errors in perception and production during language transmission. According to *Evolutionary Phonology* (e.g., Blevins 2004), channel bias alone accounts for cross-linguistic phonological typology: these phonetic predispositions lead to recurrent sound changes across unrelated languages, and speakers are not necessarily aware of these biases. In contrast, *Phonetically Based Phonology* (e.g., Hayes et al. 2004) argues that, in addition to channel bias, analytic bias also plays a crucial role. This view holds that speakers possess innate detailed phonetic and structural knowledge, which shapes the way they acquire phonological patterns synchronically.

Given that channel bias and analytic bias make similar predictions on phonological typology, even when AGL experiments show that phonetically natural patterns can be better learned, such as in our study, we still cannot unequivocally claim that the typological asymmetries stemmed from innate analytic bias. This is what Beguš (2022) termed the "duplication problem." Moreover, most of the results in the current study can be interpreted as reflecting the effects of channel bias during the experiment. Specifically, the learning advantage of the OCP-Place pattern could potentially stem from the greater perceptual salience of disagreeing consonant pairs compared to agreeing pairs, consistent with findings from Woods et al. (2010) on CVC identification. The experiment was also not designed to control for the precursor strength across the different phonotactic patterns tested. As a result, it remains unclear whether the observed OCP learning advantage is solely a product of perceptual difficulty or whether it also reflects an active component of speakers' synchronic grammatical knowledge. We interpret the findings as suggesting that both channel bias and analytic bias may contribute to the OCP advantage, potentially accounting for its robust typological attestation. This echoes other theoretical claims in the literature that do not necessarily privilege channel or analytic bias and acknowledge both of their effects (e.g., Moreton 2008).

3.6 Future directions

Finally, we would like to address some of the limitations of our study to set goals for future work. First, while our stimuli in the test phase all carried novel vowel combinations and thus never occurred in the training, we did not include unseen consonants. For the same major place, we only included obstruents in our test stimuli design and did not consider other manners such as nasals. In other words, our participants were both trained and tested on obstruents only, and this design cannot tell us if they were able to generalize the OCP pattern in the training to other unseen consonants. For example, the labial nasal [m] and velar nasal [ŋ] also exist in Mandarin sound inventory. It would be interesting to see if Mandarin participants, by being exposed to co-occurrence constraints among [p p^h f], could generalize this Labial-OCP pattern to *[p…m] or *[m…f], etc. This would serve as even stronger evidence for the learnability of OCP and argue against the hesitance that our results are merely superficial statistical learning on the co-occurrence rates of the consonants appeared in the experiments. While testing unseen consonants has its own merits, there are also practical reasons against adopting this design. For example, the test stimuli in our experiment involve all 36 possible consonant pairs among [p p^h f k k^h x]. To maintain a balanced representation of labials and dorsals, we should add both [m] and [ŋ] to this list and pair them up with other obstruents, if we want to test potential generalization to unseen sounds. This is not viable, as velar nasal [ŋ] cannot occur as the onset of a syllable in Mandarin.

Second, the current experiment was based on phonotactic learning only. The complete story of phonological learning is not limited to familiarizing with the target language's phonotactics; speakers must also learn its phonological alternations. As mentioned in the introduction section, OCP not only regulates segment co-occurrence, but also triggers or blocks phonological alternations (Yip 1988; Suzuki 1998). Therefore, we are also interested whether OCP-triggered alternations are easier to learn than other arbitrarily conditioned alternation patterns. Furthermore, static phonotactics and active alternations can often be described by the same set of constraints (known as the Duplication Problem, Kenstowicz & Kisseberth 1977), it is therefore hypothesized that the acquisition of static phonotactics and alternations is subject to the same synchronic grammatical mechanism (Tesar & Prince 2003; 2007; Hayes 2004). Results from AGL experiments also indicate that vowel harmony in the stem can aid the learning of a vowel-harmony-based alternation (Chong 2021). We believe that future studies can benefit from testing the learnability of OCP-based alternation, and if such alternation learning can be bolstered by OCP in stem phonotactics.

Third, upon examining the typology of OCP patterns, we do see that the co-occurrence of two identical segments may behave differently. For example, in Muna, roots with homorganic consonants are underrepresented in the lexicon, representing an OCP-Place effect, yet identical consonants can co-occur freely within a root (Coetzee & Pater 2008). This characteristic of identity exemption is not uncommon in the world's languages (Graff 2012). It is not uncommon for languages that exhibit OCP-Place effects in their lexicons to exempt identical consonant pairs from this constraint (e.g., Arabic, English, Muna, among others). However, Graff (2012) identified a correlation between the underattestedness of similarity avoidance and overattestedness of identical pairs in languages with OCP-Place effects, suggesting that identity exemption may be a language-specific lexical pattern. At the same time, the perceptual basis for identity exemption remains largely unclear (Woods et al. 2010). The nature of, and the relation between general OCP-Place and identity avoidance, are a promising avenue for future research. For example, would the learning advantage for OCP persist if identical items were excluded from the training input in the Harmony and Arbitrary conditions?

4 Concluding remarks

This paper has demonstrated the role of the Obligatory Contour Principle in phonological learning and helped us gain a better understanding of the nature of various learning biases speakers implement during acquisition. Moreover, despite being partly driven by the dispreference for identity, the overall learnability difference across the OCP and Harmony conditions suggests that this bias is a domain-specific linguistic substantive bias, not just a domain-general cognitive predisposition. Therefore, the current study adds another piece of evidence for the role of substantive bias in phonological learning.

Upon examining the typology, we find that some phonological patterns occur over and over again, some other patterns occur less frequently, and some logically possible patterns are completely unattested in languages. In the case of similarity avoidance, because it is supported by perceptual and articulatory grounding, lexical items that conform to the OCP are easier to process and produce. As a result, these forms will be favored in language acquisition, lexical borrowing, and the coining of new words. Over the course of many generations, the OCP patterns can better survive waves of language changes and become typologically well-attested (Frisch et al. 2004).

The main contribution of the current study is that similarity avoidance (OCP) belongs to the set of synchronic learning biases, and it should be a substantive bias since the OCP is phonetically well-motivated. By favoring the learning of OCP-based patterns, similarity avoidance can be easily and efficiently learned through generations, and thus OCP is a typologically frequent pattern. This argument does not preclude the effects of diachronic sound changes on phonological typology. Both synchronic and diachronic factors can contribute to the typological asymmetries in the world's languages.

Together, this paper contributes to our understanding of the nature of the OCP-Place and how this principle may serve as a learning bias in phonological acquisition and thus influence phonological typology. In a broader sense, using OCP-Place as a test case, this paper provides further evidence to support that phonological learning is a biased process. This helped us move one small step forward in understanding the cognitive mechanisms of how speakers internalize the linguistic input they receive and how they use that knowledge to produce language.

Appendix 1: Stimuli for the Experiment

All stimuli are transcribed in Pinyin. The correspondence between Pinyin and IPA is as follows: p = [p^h], b = [p], k = [k^h], g = [k], h = [x]. All other symbols are the same in both systems.

Training Stimuli in OCP Condition

pukou	pugou	bakei	bouha	feigu	keipa	keifu	geibou	fagu	hubou
pouka	peihou	beikou	beihu	fahei	koupei	koufu	goubu	fouga	heibu
paku	puha	bagou	bouhu	fuhei	kouba	kufa	gabou	fugei	habou
peikou	pahou	bagei	fakou	feihou	kabu	geipu	gubei	heipu	heiba
pakei	puhou	bugei	foukei	fuhou	koubu	gupou	gafou	heipou	hafu
peigou	peiha	buga	feiku	fouha	kabou	geipa	goufu	houpu	hufei
puga	bukei	bouga	fouka	koupu	koubei	gapu	gafei	hapei	houfa
pagu	beiku	bahou	fukei	kupa	kafei	goupa	goufa	hupei	houfu
pouga	bukou	buhou	fuga	keipou	keifou	gabei	geifa	habei	hafei

Training Stimuli in Harmony Condition

pupou	pubou	bapei	boufa	feibu	keika	kahou	gouga	fabou	hukou
poupa	peifou	beipou	beifu	fafu	koukei	kouhu	gougu	fouba	heigu
papu	pufa	babou	bafei	fufa	kouga	kuha	gagou	fubei	hagou
peipou	pafou	babei	fapou	feifou	kagu	geiku	gugei	heiku	heiga
papei	pufou	bubei	foupei	fufou	kougu	gukou	gahou	heikou	hahu
peibou	peifa	buba	feipa	foufa	kagou	geika	gouhu	huka	huhei
puba	bupei	bouba	foupa	kouku	kagei	gaku	gahei	hakei	houha
pabu	beipu	bafou	fupei	kuka	kahei	gouka	gouha	hukei	heiha
pouba	bupou	bufou	fuba	keikou	keihou	gagei	geiha	hagei	hahei

Training Stimuli in Arbitrary Condition

pukou	peiha	fouka	koubu	hupei	fapu	beipa	fouba	heipou	peifou
pouka	bukei	fukei	kabou	hafu	fupou	bapou	feibou	houpu	peifa
paku	beiku	koupu	koubei	hufei	kahei	bapei	feiba	hapei	poufa
peikou	bukou	kupa	gabei	houfa	kahu	gukei	fabou	boubu	pafu
pakei	bakei	keipou	geibou	houfu	kuhou	gouku	geihou	babu	gugei
peihou	beikou	keipa	goubu	hafei	kahou	gakei	geiha	beibou	gouga
puha	fakou	koupei	gabou	foupu	kouhei	gouka	gahou	buba	guga
pahou	foukei	kouba	gubei	feipou	bupei	geikou	guha	boubei	geiga
puhou	feiku	kabu	heipu	foupa	bupou	fuba	gouha	poufu	geigu

Test Stimuli

papou	peigu	boufu	feipu	feigou	kufei	gapei	geigou	hafou
peibu	pahei	bouku	fabu	fouhei	kukei	goubei	gahu	houku
poufei	bapu	beigou	fafei	kapu	kougei	goufei	houpei	houga
peiku	bubou	bahei	feikou	kubou	keihu	gakou	heibou	houhu

Supplementary files

The raw judgment data of our experiments and all the codes for statistical analyses can be found at: https://osf.io/9gwps/.

Ethics and consent

All experimental procedures were reviewed and approved by the Human Subjects Committee at the University of Kansas (ID: STUDY00147485).

Funding information

This work was supported by Discipline Team Support Program of Beijing Language and Culture University (2023YGF06), Science Foundation of Beijing Language and Culture University (supported by “the Fundamental Research Funds for the Central Universities”) (23YBB06), and a Graduate Research Award from the University of Kansas.

Acknowledgements

For helpful discussions and feedback, we thank Robert Fiorentino, Allard Jongman, Yan Li, Joan Sereno, Annie Tremblay, James White, two anonymous reviewers of *Glossa*, and Associate Editor Sara Finley. All remaining errors are our own. We thank Ren Cai, Chulong Liu, and Wenting Tang for their assistance in advertising the online experiments in China. Finally, we thank all the participants for their contribution.

Competing interests

The authors have no competing interests to declare.

References

Anwyl-Irvine, Alexander L. & Massonnié, Jessica & Flitton, Adam & Kirkham, Natasha & Evershed, Jo K. 2020. Gorilla in our midst: An online behavioral experiment builder. *Behavior Research Methods* 52(1). 388–407. DOI: https://doi.org/10.3758/s13428-019-01237-x

Archangeli, Diana & Pulleyblank, Douglas. 1994. *Grounded phonology*. MIT Press.

Bates, Douglas & Mächler, Martin & Bolker, Ben & Walker, Steve. 2015. Fitting linear mixed-effects models using lme4. *Journal of Statistical Software* 67. 1–48. DOI: https://doi.org/10.18637/jss.v067.i01

Becker, Michael & Ketrez, Nihan & Nevins, Andrew. 2011. The Surfeit of the Stimulus: Analytic biases filter lexical statistics in Turkish laryngeal alternations. *Language* 87(1). 84–125. DOI: https://doi.org/10.1353/lan.2011.0016

Beguš, Gašper. 2020. Estimating historical probabilities of natural and unnatural processes. *Phonology* 37(4). 515–549. DOI: https://doi.org/10.1017/S0952675720000263

Beguš, Gašper. 2022. Distinguishing cognitive from historical influences in phonology. *Language* 98(1). 1–34. DOI: https://doi.org/10.1353/lan.2021.0084

Berent, Iris & Shimron, Joseph. 1997. The representation of Hebrew words: Evidence from the obligatory contour principle. *Cognition* 64. 39–72. DOI: https://doi.org/10.1016/S0010-0277(97)00016-4

Berent, Iris & Shimron, Joseph & Vaknin, Vered. 2001. Phonological constraints on reading: evidence from the obligatory contour principle. *Journal of Memory and Language* 44(4). 644–665. DOI: https://doi.org/10.1006/jmla.2000.2760

Berent, Iris & Wilson, Colin & Bemis, Douglas K. & Marcus, Gary F. 2012. On the role of variables in phonology: Remarks on Hayes and Wilson 2008. *Linguistic Inquiry* 43(1). 97–119. DOI: https://doi.org/10.1162/LING_a_00075

Berkley, Deborah. 1994. The OCP and gradient data. *Studies in the Linguistic Sciences* 24. 59–72.

Blevins, Juliette. 2004. *Evolutionary phonology*. Cambridge: Cambridge University Press.

Boersma, Paul. 2000. The OCP in the perception grammar. *Rutgers Optimality Archive* 435. 1–52.

Boersma, Paul & Weenink, David. 2017. Praat: doing phonetics by computer.

Boll-Avetisyan, Natalie. 2012. *Phonotactics and its acquisition, representation, and use: An experimental-phonological study*. Utrecht: LOT Publications.

Boll-Avetisyan, Natalie & Kager, René. 2014. OCP-PLACE in speech segmentation. *Language and Speech* 57(3). 394–421. DOI: https://doi.org/10.1177/0023830913508074

Boll-Avetisyan, Natalie & Kager, René. 2016. Is speech processing influenced by abstract or detailed phonotactic representations? The case of the Obligatory Contour Principle. *Lingua* 171. 74–91. DOI: https://doi.org/10.1016/j.lingua.2015.11.008

Brohan, Anthony. 2014. *Analytic bias in coocurrence restrictions*. MIT dissertation.

Chen, Tsung-ying. 2022. On the learnability of level-based and unit-based tonal OCP constraints: An artificial grammar learning study. *Glossa*. DOI: https://doi.org/10.16995/glossa.5795

Chong, Adam Junxiang. 2021. The effect of phonotactics on alternation learning. *Language* 97(2). 213–244. DOI: https://doi.org/10.1353/lan.2021.0017

Clements, George N. & Hume, Elizabeth V. 1995. The internal organization of speech sounds. In *The handbook of phonological theory*, 245–306. Blackwell. DOI: https://doi.org/10.1111/b.9780631201267.1996.00009.x

Coetzee, Andries W. & Pater, Joe. 2008. Weighted constraints and gradient restrictions on place co-occurrence in Muna and Arabic. *Natural Language & Linguistic Theory* 26(2). 289–337. DOI: https://doi.org/10.1007/s11049-008-9039-z

Cser, András. 2010. The -alis/-aris allomorphy revisited. In Rainer, Franz & Dressler, Wolfgang U. & Kastovsky, Dieter & Luschützky, Hans Christian & Peters, Elisabeth (eds.), *Variation and change in morphology: Selected papers from the 13th International Morphology Meeting, Vienna, February 2008*, 33–52. John Benjamins Publishing Company. DOI: https://doi.org/10.1075/cilt.310.02cse

Danis, Nick. 2019. Long-distance major place harmony. *Phonology* 36(4). 573–604. DOI: https://doi.org/10.1017/S0952675719000307

Dell, Gary S. & Burger, Lisa K. & Svec, William R. 1997. Language production and serial order: A functional analysis and a model. *Psychological Review* 104(1). 123–147. DOI: https://doi.org/10.1037//0033-295X.104.1.123

Do, Youngah & Havenhill, Jonathan. 2021. Production and substantive bias in phonological learning. *Proceedings of the Annual Meetings on Phonology* 9. DOI: https://doi.org/10.3765/amp.v9i0.4925

Duanmu, San. 1994. Against contour tone units. *Linguistic Inquiry* 25(4). 555–608.

Finley, Sara. 2011. The privileged status of locality in consonant harmony. *Journal of Memory and Language*. DOI: https://doi.org/10.1016/j.jml.2011.02.006

Finley, Sara. 2012. Testing the limits of long-distance learning: Learning beyond a three-segment window. *Cognitive Science* 36(4). DOI: https://doi.org/10.1111/j.1551-6709.2011.01227.x

Finley, Sara. 2017. Locality and harmony: Perspectives from artificial grammar learning. *Language and Linguistics Compass* 11(1). e12233. DOI: https://doi.org/10.1111/lnc3.12233

Finley, Sara & Badecker, William. 2009. Artificial language learning and feature-based generalization. *Journal of Memory and Language* 61(3). 423–437. DOI: https://doi.org/10.1016/j.jml.2009.05.002

Frisch, Stefan A. 2004. Language processing and segmental OCP effects. In Hayes, Bruce & Kirchner, Robert & Steriade, Donca (eds.), *Phonetically based phonology*, 346–371. Cambridge: Cambridge University Press. DOI: https://doi.org/10.1017/CBO9780511486401.011

Frisch, Stefan A. & Pierrehumbert, Janet B. & Broe, Michael B. 2004. Similarity avoidance and the OCP. *Natural Language and Linguistic Theory* 22(1). 179–228. DOI: https://doi.org/10.1023/B:NALA.0000005557.78535.3c

Frisch, Stefan A. & Zawaydeh, Bushra Adnan. 2001. The psychological reality of OCP-place in Arabic. *Language* 77(1). 91–106. DOI: https://doi.org/10.1353/lan.2001.0014

Gallagher, Gillian. 2010. Perceptual distinctness and long-distance laryngeal restrictions. *Phonology* 27(3). 435–480. DOI: https://doi.org/10.1016/j.lingua.2011.11.012

Gallagher, Gillian. 2013. Learning the identity effect as an artificial language: bias and generalisation. *Phonology* 30(2). 253–295. DOI: https://doi.org/10.1017/S0952675713000134

Gallagher, Gillian. 2016. Asymmetries in the representation of categorical phonotactics. *Language* 92(3). 557–590. DOI: https://doi.org/10.1353/lan.2016.0048

Gallagher, Gillian. 2019. Phonotactic knowledge and phonetically unnatural classes: the plain uvular in Cochabamba Quechua. *Phonology* 36(1). 37–60. DOI: https://doi.org/10.1017/S0952675719000034

Gallagher, Gillian & Coon, Jessica. 2009. Distinguishing total and partial identity: Evidence from Chol. *Natural Language and Linguistic Theory* 27(3). 545–582. DOI: https://doi.org/10.1007/s11049-009-9075-3

Garrett, Andrew & Johnson, Keith. 2013. Phonetic bias in sound change. *Origins of Sound Change: Approaches to Phonologization* 1. 51–97. DOI: https://doi.org/10.1093/acprof:oso/9780199573745.003.0003

Glewwe, Eleanor. 2019. *Bias in phonotactic learning: Experimental studies of phonotactic implicationals (PhD thesis)*. UCLA dissertation.

Glewwe, Eleanor. 2022. Substantive bias and the positional extension of major place contrasts. *Glossa: A Journal of General Linguistics* 7(1). DOI: https://doi.org/10.16995/glossa.6537

Gomez, Rebecca L. & Gerken, Louann. 1999. Artificial grammar learning by 1-year-olds leads to specific and abstract knowledge. *Cognition* 70(2). 109–135. DOI: https://doi.org/10.1016/S0010-0277(99)00003-7

Gong, Shuxiao & Zhang, Jie. 2021. Modelling Mandarin speakers' phonotactic knowledge. *Phonology* 38(2). 241–275. DOI: https://doi.org/10.1017/S0952675721000166

Gong, Shuxiao & Zhang, Jie & Fiorentino, Robert. 2023. Phonological well-formedness constraints in Mandarin phonotactics: Evidence from lexical decision. *Language and Speech* 67(3). 676–691. DOI: https://doi.org/10.1177/00238309231182363

Gordon, Matthew. 2016. *Phonological Typology*. Oxford: Oxford University Press. DOI: https://doi.org/10.1093/acprof:oso/9780199669004.001.0001

Graff, Peter. 2012. *Communicative efficiency in the lexicon (PhD thesis)*. MIT dissertation.

Gwilliams, L. & Marantz, A. 2015. Non-linear processing of a linear speech stream: The influence of morphological structure on the recognition of spoken Arabic words. *Brain and Language* 147. 1–13. DOI: https://doi.org/10.1016/j.bandl.2015.04.006

Hansson, Gunnar. 2010. *Consonant harmony: long-distance interaction in phonology*. Berkeley: University of California Press.

Hansson, Gunnar. 2020. Consonant Harmony. In Hansson, Gunnar (ed.), *Oxford Research Encyclopedia of Linguistics*. Oxford University Press. DOI: https://doi.org/10.1093/acrefore/9780199384655.013.369

Hayes, Bruce. 2004. Phonological acquisition in Optimality Theory: The early stages. In Kager, René & Pater, Joe & Zonneveld, Wim (eds.), *Constraints in phonological acquisition*, 158–203. Cambridge: Cambridge University Press. DOI: https://doi.org/10.1017/CBO9780511486418.006

Hayes, Bruce. 2009. *Introductory Phonology*. Chichester: John Wiley & Sons, Inc.

Hayes, Bruce & Kirchner, Robert & Steriade, Donca. 2004. *Phonetically based phonology*. (Hayes, Bruce & Kirchner, Robert & Steriade, Donca, eds.). Cambridge: Cambridge University Press. DOI: https://doi.org/10.1017/CBO9780511486401

Hayes, Bruce & Siptar, Peter & Zuraw, Kie & Londe, Zsuzsa. 2009. Natural and unnatural constraints in Hungarian vowel harmony. *Language* 85(4). 822–863. DOI: https://doi.org/10.1353/lan.0.0169

Hayes, Bruce & White, James. 2013. Phonological naturalness and phonotactic learning. *Linguistic Inquiry* 44(1). 45–75. DOI: https://doi.org/10.1162/LING_a_00119

Hayes, Bruce & White, James. 2015. Saltation and the P-map. *Phonology* 32(2). 267–302. DOI: https://doi.org/10.1017/S0952675715000159

Hughto, Coral. 2018. Investigating the consequences of iterated learning in phonological typology. DOI: https://doi.org/10.7275/R5WH2N63

Ito, Chiyuki. 2007. Morpheme structure and co-occurrence restrictions in Korean monosyllabic stems. *Studies in Phonetics, Phonology, and Morphology* 13(3). 373–394. DOI: https://doi.org/10.17959/SPPM.2007.13.3.373

Jin, Shao-Jie & Wang, Sheng-Fu & Lu, Yu-An. 2023. Identifying generalizable knowledge from the distribution of tonotactic accidental gaps in Mandarin. *Laboratory Phonology* 14(1). DOI: https://doi.org/10.16995/labphon.6455

Kager, René & Pater, Joe. 2012. Phonotactics as phonology: knowledge of a complex restriction in Dutch. *Phonology* 29(1). 81–111. DOI: https://doi.org/10.1017/S0952675712000048

Kager, René & Shatzman, Keren. 2007. Phonological constraints in speech processing. *Linguistics in the Netherlands* 24(1). 99–111. DOI: https://doi.org/10.1075/avt.24.11kag

Kawahara, Shigeto. 2006. A faithfulness ranking projected from a perceptibility scale: The case of [+voice] in Japanese. *Language* 82(3). 536–574. DOI: https://doi.org/10.1353/lan.2006.0146

Kenstowicz, Michael John & Kisseberth, Charles Wayne. 1977. *Topics in phonological theory*. New York: Academic Press.

Koo, Hahn & Callahan, Lydia. 2012a. Tier-adjacency is not a necessary condition for learning phonotactic dependencies. *Language and Cognitive Processes* 27(10). DOI: https://doi.org/10.1080/01690965.2011.603933

Koo, Hahn & Callahan, Lydia. 2012b. Tier-adjacency is not a necessary condition for learning phonotactic dependencies. *Language and Cognitive Processes* 27(10). DOI: https://doi.org/10.1080/01690965.2011.603933

Kuo, Li Jen. 2009. The role of natural class features in the acquisition of phonotactic regularities. *Journal of Psycholinguistic Research* 38(2). 129–150. DOI: https://doi.org/10.1007/s10936-008-9090-2

Lai, Regine. 2015. Learnable vs. unlearnable harmony patterns. *Linguistic Inquiry* 46(3). 425–451. DOI: https://doi.org/10.1162/LING_a_00188

Leben, William. 1973. *Suprasegmental phonology (PhD thesis)*. MIT dissertation.

Li, A. & Xu, B. 2001. *Chinese annotated dialogue and conversation corpus*. Beijing: Institute of Linguistics, Chinese Academy of Social Sciences.

Lin, Isabelle. 2017. *Tone sequences in lexical processing of Beijing Mandarin (MA thesis)*. University of California, Los Angeles dissertation.

Linzen, Tal & Gallagher, Gillian. 2017. Rapid generalization in phonotactic learning. *Laboratory Phonology: Journal of the Association for Laboratory Phonology* 8(1). 24. DOI: https://doi.org/10.5334/labphon.44

Marcus, G. F. & Vijayan, S. & Bandi Rao, S. & Vishton, P. M. 1999. Rule learning by seven-month-old infants. *Science* 283(5398). 77–80. DOI: https://doi.org/10.1126/science.283.5398.77

Martin, Alexander. 2017. *Biases in phonological processing and learning (PhD thesis)*. Université Paris sciences et lettres dissertation. Retrieved from https://theses.hal.science/tel-01939096

Martin, Alexander & Peperkamp, Sharon. 2020. Phonetically natural rules benefit from a learning bias: A re-examination of vowel harmony and disharmony. *Phonology*. DOI: https://doi.org/10.1017/S0952675720000044

Martin, Alexander & White, James. 2021. Vowel harmony and disharmony are not equivalent in learning. *Linguistic Inquiry* 52(1). 227–239. DOI: https://doi.org/10.1162/ling_a_00375

Martin, Andrew. 2007. *The evolving lexicon (PhD thesis)*. University of California, Los Angeles dissertation.

McCarthy, John J. 1986. OCP Effects: Gemination and antigemination. *Linguistic Inquiry* 17(2). 207–263.

McMullin, Kevin & Hansson, Gunnar Ólafur. 2019. Inductive learning of locality relations in segmental phonology. *Laboratory Phonology: Journal of the Association for Laboratory Phonology* 10(1). 14. DOI: https://doi.org/10.5334/labphon.150

Mester, Armin. 1988. *Studies in Tier Structure*. London: Routledge.

Moreton, Elliott. 2008. Analytic bias and phonological typology. *Phonology* 25(1). 83–127. DOI: https://doi.org/10.1017/S0952675708001413

Moreton, Elliott. 2012. Inter- and intra-dimensional dependencies in implicit phonotactic learning. *Journal of Memory and Language* 67(1). 165–183. DOI: https://doi.org/10.1016/j.jml.2011.12.003

Moreton, Elliott & Pater, Joe. 2012a. Structure and substance in artificial-phonology learning, Part I: Structure. *Language and Linguistics Compass* 6(11). 686–701. DOI: https://doi.org/10.1002/lnc3.363

Moreton, Elliott & Pater, Joe. 2012b. Structure and substance in artificial-phonology learning, Part II: Substance. *Language and Linguistics Compass* 6(11). 702–718. DOI: https://doi.org/10.1002/lnc3.366

Moreton, Elliott & Pertsova, Katya. 2014. Pastry phonotactics: Is phonological learning special. In *Proceedings of the 43rd Annual Meeting of the Northeast Linguistic Society, City University of New York*, Vol. 2, 1–14. Amherst, MA: Graduate Linguistics Students' Association.

Myers, Scott & Padgett, Jaye. 2014. Domain generalisation in artificial language learning. *Phonology* 31(3). 399–433. DOI: https://doi.org/10.1017/S0952675714000207

Newport, Elissa L. 2016. Statistical language learning: computational, maturational, and linguistic constraints. *Language and Cognition* 8(3). 447–461. DOI: https://doi.org/10.1017/langcog.2016.20

Ohala, John J. 1981. The listener as a source of sound change. *Papers from the Parasession on Language and Behavior Chicago Linguistic Society*.

Ohala, John J. 1993. The phonetics of sound change. In Jones, Charles (ed.), *Historical Linguistics: Problems and Perspectives*, 237 – 278. London: Longman Academic.

Ohala, John J. 1994. Towards a universal, phonetically-based, theory of vowel harmony. In *Proceedings of the 3rd international conference on spoken language processing*, 491–494. DOI: https://doi.org/10.21437/ICSLP.1994-113

Peperkamp, Sharon & Dupoux, Emmanuel. 2007. Learning the mapping from surface to underlying representations in an artificial language. In *Laboratory Phonology 9*, 315–338. Berlin: Mouton de Gruyter.

Prickett, Brandon. 2018. Complexity and naturalness biases in phonotactics: Hayes and White (2013) revisited. *Proceedings of the Annual Meetings on Phonology*. DOI: https://doi.org/10.3765/amp.v5i0.4230

Pycha, Anne & Nowak, Pawel & Shin, Eurie & Shosted, Ryan. 2003. Phonological rule-learing and its implications for a theory of vowel harmony. *Proceedings of the 22nd West Coast Conference on Formal Linguistics (WCCFL 22)* (January 2003). 101–114.

R Core Team. 2018. *R: A language and environment for statistical computing*. Vienna: R Foundation for Statistical Computing.

Rose, Sharon & Walker, Rachel. 2004. A typology of consonant agreement as correspondence. *Language* 80(3). DOI: https://doi.org/10.1353/lan.2004.0144

Rose, Sharon & Walker, Rachel. 2011. Harmony systems. In *The handbook of phonological theory*, 240–290. Malden, MA: Wiley-Blackwell. DOI: https://doi.org/10.1002/9781444343069.ch8

Seidl, Amanda & Buckley, Eugene. 2005. On the learning of arbitrary phonological rules. *Language Learning and Development* 1(3–4). 289–316. DOI: https://doi.org/10.1080/15475441.2005.9671950

Skoruppa, Katrin & Peperkamp, Sharon. 2011. Adaptation to novel accents: Feature-based learning of context-sensitive phonological regularities. *Cognitive Science* 35. 348–366. DOI: https://doi.org/10.1111/j.1551-6709.2010.01152.x

Stanton, Juliet. 2016. Learnability shapes typology: The case of the midpoint pathology. *Language* 92(4). 753–791. DOI: https://doi.org/10.1353/lan.2016.0071

Suzuki, Keiichiro. 1998. *A typological investigation of dissimilation (PhD Thesis)*. The University of Arizona dissertation.

Tanaka, Yu. 2023. Learning biases in proper nouns. *Phonology* 40(1–2). 155–186. DOI: https://doi.org/10.1017/s0952675724000046

Tesar, Bruce & Prince, Alan. 2003. Using phonotactics to learn phonological alternations. In *Proceedings from the Annual Meeting of the Chicago Linguistic Society*, 241–269.

Tesar, Bruce & Prince, Alan. 2007. Using phonotactics to learn phonological alternations. In *Proceedings of the Thirty-Ninth Conference of the Chicago Linguistics Society*, 209–237.

Walter, Mary. 2007. *Repetition avoidance in human language (PhD Thesis)*. MIT dissertation.

White, James. 2017. Accounting for the learnability of saltation in phonological theory: A maximum entropy model with a P-map bias. *Language* 93(1). 1–36. DOI: https://doi.org/10.1353/lan.2017.0001

Wilson, Colin. 2006. Learning phonology with substantive bias: an experimental and computational study of velar palatalization. *Cognitive Science* 30(5). 945–982. DOI: https://doi.org/10.1207/s15516709cog0000_89

Woods, David. L. & Yund E. William & Herron, Timothy J. & Cruadhlaoich, Matthew A. I. Ua. 2010. Consonant identification in consonant-vowel-consonant syllables in speech-spectrum noise. *The Journal of the Acoustical Society of America* 127(3). 1609–1623. DOI: https://doi.org/10.1121/1.3293005

Yi, Li & Duanmu, San. 2015. Phonemes, features, and syllables: Converting onset and rime inventories to consonants and vowels. *Language and Linguistics* 16(6). 819–842. DOI: https://doi.org/10.1177/1606822X15602610

Yip, Moira. 1982. Reduplication and CV skeleta in Chinese secret languages. *Linguistic Inquiry* 13(4). 637–661.

Yip, Moira. 1988. The obligatory contour principle and phonological rules: A loss of identity. *Linguistic Inquiry* 19(1). 65–100.

Zheng, Shuang & Do, Youngah. 2025. Substantive bias in artificial phonology learning. *Language and Linguistics Compass* 19(1). e70005. DOI: https://doi.org/10.1111/lnc3.70005

Zuraw, Kie. 2007. The role of phonetic knowledge in phonological patterning: Corpus and survey evidence from Tagalog infixation. *Language* 83(2). 277–316. DOI: https://doi.org/10.1353/lan.2007.0105