Nonlocal phonological interactions and trigram constraints

Gillian Gallagher and Maria Gouskova (New York University)

Inducing nonlocal representations Nonlocal phonological patterns such as vowel harmony and long-distance consonant assimilation and dissimilation motivate representations that include only the interacting segments—e.g., autosegmental tiers/projections (Clements 1976 et seq.) or correspondence relations (Hansson 2001; Rose and Walker 2004; Bennett 2015). Learning nonlocal interactions inductively is much harder than learning local (bigram) co-occurrence constraints, since there are many more possible trigrams and tetragrams than bigrams. Some heuristics are needed for identifying nonlocal interactions and narrowing down the search space. We present an implemented computational model that induces projections based on phonotactic properties that are observable without nonlocal representations, in real data. We illustrate the success of our model with case studies of Quechua and Aymara, and present the interesting challenges posed by additional patterns in Shona and Kinyarwanda.

Inductive Projection Learner Our learner capitalizes on three observations about the phonological properties of nonlocal interactions: (a) they often turn up in languages that have simple syllable structure (McCarthy 1989), (b) they involve segments whose distributions are sufficiently varied that they end up being trigram-local, (c) they involve segments that belong to a natural class—often a small one (liquids, stridents, plosives). We build on the UCLA Phonotactic Learner (Hayes & Wilson 2008), which induces a weighted constraint grammar given some learning data and a feature set. Our model searches this baseline grammar for *placeholder constraints*: trigram constraints where the middle gram is "any segment" X, as in *[+low]-X-[-high,-low] (as in Hayes & Wilson's Shona case study). These constraints are a clue that [+low] and [-high,-low] interact nonlocally, since the identity of the middle segment does not matter. Based on these placeholder constraints, our model induces a nonlocal projection of the *smallest natural class* that includes the classes surrounding the placeholder X (representing the most general hypothesis about what segments interact nonlocally).

Quechua laryngeal phonotactics One key insight of our learner is that segments that interact in nonlocal phonological patterns are frequently separated by just one intervening segment. For example, in Quechua, pairs of plain, ejective and aspirated stops are subject to nonlocal restrictions: ejectives and aspirates may not be preceded, at any distance, by another stop (Gallagher 2016). While this generalization holds at any distance (e.g., *[k'amp'i], *[k'amip'a]), interacting stops are often separated by just one vowel (e.g., *[k'ap'i]). Our training data is a corpus of 10,848 phonological words from the *Conosur Ñawpaqman* newspaper. In this corpus, there are 742 stop…stop pairs, but 313 (42%) of these are separated by just a single vowel.

Interactions between consonants in a CVC configuration can be captured via a trigram constraint in the baseline grammar. The UCLA Phonotactic Learner induces a baseline grammar for Quechua with two placeholder constraints: *[-continuant, -sonorant]X[+cg] and *[-continuant, -sonorant]X[+sg]. These constraints account for part of the generalization: they penalize ungrammatical forms like *[k'ap'i], but incorrectly allow equally ungrammatical forms like *[k'amp'i] or *[k'amip'a]. Based on these constraints, our learner builds a stop projection, [-continuant, -sonorant]--the smallest class that includes the natural classes in the placeholder constraints. On this projection, the model learns two constraints that concisely capture the distribution of ejectives and aspirates: *X-[+cg] and *X-[+sg], "*ejectives and aspirates should* *be the first stop in the word”*. The figure below plots the harmony scores that the final grammar assigned to a set of 24,352 disyllabic nonce words. Legal and illegal forms are clearly separated, showing that the model captures the phonotactic generalizations.

Three more case studies The model is similarly successful at finding trigram placeholder constraints in Aymara, Shona and Kinyarwanda. ***Aymara*** shows laryngeal restrictions similar to those in Quechua, but there are fewer restricted combinations. As a result, the relevant constraints cover smaller natural classes. ***Shona*** has vowel harmony alternations and a static height harmony pattern in its verbal stems (Beckman 1997, Hayes & Wilson 2008): mid vowels and high vowels generally agree in height [-pe̲r-e̲r-a] ‘end in’ vs. [-i̲p-i̲r-a] ‘be evil for’, and [a] cannot be followed by mid vowels, [-pofoma̲dz-i̲r-a] ‘blind for’. We trained the learner on a list of 4,688 verbs (Chimhundu 1996). As in the other languages, a large proportion of nonlocally interacting segments appear in trigrams: CC and CCC clusters occur, but 79% of vowels are separated by just one consonant, VCV. In the baseline grammar without projections, there are several placeholder constraints, e.g., *[+lo]X[-hi, -lo] and *[+hi]X[-hi,-lo]. The final model induced by our learner captures most of the vowel co-occurrence restrictions, but on four different projections (the generalizations are spread over projections including mid vowels, nonlow vocoids, and nonhigh vowels). ***Kinyarwanda*** (Walker, Byrd and Mpiranya 2008), has sibilant harmony (/-sas-i/ [-ʂaʂi] ‘bed maker’, cf. [-sas-a] ‘make the bed’). Our model notices the transvocalic interactions between retroflex and alveolar sibilants, and induces a sibilant projection to account for the harmony. In both Shona and Kinyarwanda, however, there are opaque segments that reportedly block harmony (for Kinyarwanda, coronal stops: /-siːtaːʐ-e/ not *[-ʂiːtaːʐ-e] ‘make stub’). If these opaque segments are not present on the induced projection, our models miss this aspect of the pattern. As we show, however, the statistical evidence for opacity in Shona and Kinyarwanda is rather weak. We discuss methods for searching through larger projections for opaque patterns for languages that present a stronger case for needing such expanded projections.

Discussion Our computational model learns nonlocal phonological generalizations based on baseline phonotactics, capitalizing on the empirical observation that nonlocal interactions often show up as trigram constraints in a model with no nonlocal projections. Unlike other approaches, our learner does not stipulate the representations, assume a priori that nonlocal interactions exist in a language (Hayes & Wilson 2008; Futrell et al. 2015), or artificially idealize the learning set to include only inviolable distributions (Heinz 2010, Jardine 2015).