Inducing nonlocal constraints from baseline phonotactics

Gillian Gallagher and Maria Gouskova, NYU

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 phonology inductively from the data requires either stipulating the representations (Hayes & Wilson 2008; Futrell et al. 2015) or an idealized learning set (Heinz 2010, Jardine 2015). We present an implemented computational model that induces projections based on phonotactic properties that are observable without nonlocal representations, in real data.

Inductive Projection Learner. 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 builds a search space for nonlocal projections starting from the smallest natural class that includes the classes surrounding the placeholder X, and continues searching through larger such classes. The learner returns the projection(s) encoding the strongest generalizations, based on constraint weights and the size of the projections.

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'<u>amp</u>'i], *[k'<u>amip</u>'a]), interacting stops are usually separated by just one vowel (e.g., *[k'<u>a</u>p'i]). Our training data for Quechua is a corpus of 10,848 phonological words from the *Conosur Ñawpaqman* newspaper (www.cenda.org/periodico-conosur). In this corpus, there are 742 stop…stop pairs, but 313 (42%) of the stop pairs are separated by just one vowel.

Interactions between consonants in a <u>C</u>V<u>C</u> configuration (or vowels in a <u>V</u>C<u>V</u> 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][][+cg] and *[-continuant, -sonorant][][+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 search space starting with 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 learner returns this projection after searching larger projections such as [-syllabic], which yield weaker generalizations.

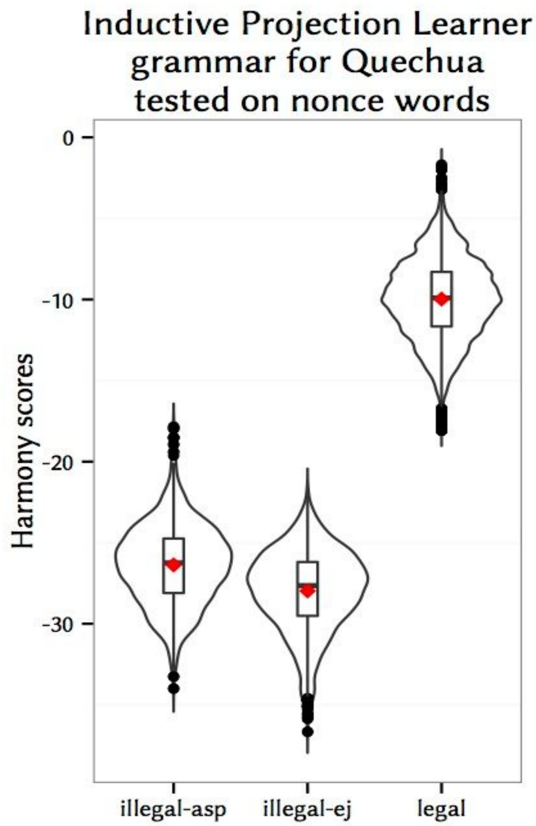

The figure plots the harmony scores the final grammar assigned to a set of 24,352 disyllabic nonce words. Legal and illegal forms are clearly separated, showing that the model has captured the phonotactic generalizations of the language.

Shona vowel harmony. 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 [-per-er-a] 'end in' vs. [-ip-ir-a] 'be evil for', and [a] cannot be followed by mid vowels, [-pofomadz-ir-a] 'blind for'. We trained the learner on a list of 4,688 verbs (Chimhundu 1996), where 79% of vowels appear in VCV trigrams. The baseline grammar includes several placeholder constraints, e.g., *[+lo]X[-hi, -lo] and *[+hi]X[-hi,-lo]. Our learner constructs search paths starting with [-high] and [-low], and eventually converges on [+syllabic], the smallest projection that includes all and only the interacting segments. Larger projections such as [-consonantal] (glides) are correctly rejected, based on lower constraint weights.

Discussion Our computational model learns nonlocal phonological generalizations via a structured search through a space of projections. Our model capitalizes on the empirical observation that nonlocal interactions often show up as trigram constraints. It also incorporates the traditional insight that natural class structure restricts the types of nonlocal phonological dependencies found in languages.