Modeling inter-speaker variation in Velar Softening

In English velar softening (VS), underlying /k/ becomes [s] when it occurs in the environment: /_+{ɪ, a͡ɪ}, where + indicates a morpheme boundary. This is a non-automatic process as lexically-specified exceptions abound. For example, the *-ity* [ɪti] suffix triggers velar softening: *electri*[k]∼*electri*[s]*ity*. However, the suffix *-ish* [ɪʃ] does not: *electri*[k]∼*electri*[k]*ish*. Velar softening is not driven purely by phonotactics. The [kɪ] and [ka͡ɪ] sequences are legal in English. But when {ɪ, a͡ɪ} belong to a VS-triggering suffix, [k] is severely under-represented relative to [s] in the /_ (+){ɪ, a͡ɪ} environment: the %[k] (out of all [k] or [s]) is 31.9% in this environment in general, but only 7.2% with VS-triggering suffixes. (Data from CMU Pronouncing Dictionary; counts based on type frequency.) I will model this asymmetry by utilizing non-automatic markedness constraints that prohibit the [kɪ] and [ka͡ɪ] sequences only when their segments arise from a VS-triggering suffix. Such constraints are non-automatic as they make reference to specific affixes (Alderete 1999, Itô & Mester 1999). One implication of lexically-specified constraints is that learning when /k/ becomes [s] goes hand in hand with learning morphology (*e.g.* To segment a word into stem+suffix? Which suffixes are VS triggers?). Hence, I will propose a model that learns them simultaneously.

When presented with the same surface information, different learners may arrive at different analyses (*i.e.* inter-speaker variation in analyses). In a production task with the *-ity* suffix and nonce stems, Pierrehumbert (2006) found that 2 in 10 participants never applied velar softening. This suggests that approximately 20% of learners may not learn velar softening, but may instead memorize full underlying forms (*e.g.* /ɪlɛktɹɪsɪti/) for existing words.

The present study models the underlying representations (UR) and grammars that give rise to inter-speaker variation in velar softening by modeling how velar softening is acquired. I assume that word segmentation has taken place by this point in learning, so the child has access to pairs like *electricity* (*word*) and [ɪlɛktɹɪsɪti] (SR), and has to figure out the URs and grammar that govern the *word*-to-SR mapping. There are three questions to figure out for the intermediate *word*-to-UR mapping: (1) Is there a morpheme boundary in *electric*(+)*ity*, (2) If so, is the suffix one that triggers VS, and (3) Is the consonant of interest /k/ or /s/? For the subsequent UR-to-SR mapping, the child has learn whether and in which contexts /k/ becomes [s].

The model consists of two Maximum Entropy (Goldwater & Johnson 2003) sub-models that are chained together via the product rule: $P(word, ur, sr) = P(ur, word) \times P(sr|ur, word)$. The representation sub-model, $P(ur, word)$, models the *word*-to-UR mapping. The grammar sub-model, $P(sr|ur, word)$, models whether and when /k/→[s] applies. The two sub-models interact in order to match the observed *word*-SR frequencies. For instance, a highly improbable /k/→[s] in the grammar sub-model influences the representation sub-model towards a UR like /ɪlɛktɹɪsɪti/ (can produce the correct SR without needing /k/→[s]), and away from a UR like /ɪlɛktɹɪk+ɪti/, for the *word electricity*. Likewise, the representation sub-model can influence the grammar sub-model. Highly probable /ɪlɛktɹɪk+ɪti$_{\text{VS-suffix}}$/ and /kɪti/ (*kitty*) will push the grammar sub-model towards applying /k/→[s] only when a VS-triggering suffix is present.

The representation sub-model is parameterized by UR constraints (Zuraw 2000), whose weights are adjusted during UR learning. I use two types of UR constraints: (1) Constraints that specify the presence of a particular morpheme in a word. Such constraints model the segmentation into morphemes. *e.g.* The UR constraint ELECTRIC specifies that the morpheme <electric> is present in the *word electricity*. It is satisfied by the UR /ɪlɛktɹɪk+ɪti/, but not by /ɪlɛktɹɪsɪti/. (2) Constraints that specify relevant properties of the UR of a morpheme. *e.g.* The UR constraint (ELECTRIC, /k/) specifies that the morpheme <electric> should have a morpheme-final /k/. It is satisfied by the UR /ɪlɛktɹɪk+ɪti/, but not by /ɪlɛktɹɪs+ɪti/.

The grammar sub-model has constraints like *kɪ that disfavour all [kɪ] sequences, and a faithfulness constraint favouring the retention of the underlying segment.

The CMU data suggested that velar softening was driven by the presence of particular suffixes, so I included the suffix-specific constraint $^*\text{kɪ}_{\text{VS-suffix}}$, which disfavoured the [kɪ] sequence if any of its segments arose from a VS-triggering suffix. Each suffix was associated with a UR constraint (*e.g.* (-ITY, VS-suffix)) to model its VS-triggering status.

The training data consisted of the frequencies of *word*-SR pairs (*e.g. electricity*∼[ɪlɛktɹɪsɪti]: 1, *electricity*∼[ɪlɛktɹɪkɪti]: 0, *kitty*∼[kɪti]: 1, *kitty*∼[sɪti]: 0, *etc.*) to mirror the information that a child received. URs are missing from the observed data (*word*-SR pairs) that the full *word*-UR-SR model predicts. I use Expectation-Maximization (Jarosz 2015) to simulate learning in this "missing information" scenario. The child/learner's goal was to match the training data as closely as possible. 11 of 125 randomly initialized runs were able to match the training data perfectly (as measured by hitting the ceiling likelihood). In other words, only these 11 trained models produced the correct surface forms for words they had encountered before.

I subjected these 11 models to a wug task. I included only models that matched the training data perfectly because actual English speakers have correct surface forms for existing words that they've heard before like *electricity*, *kitty*, *etc.* I assume humans have a way to avoid the local optima that fail to match the training data, such as by further exploring the learning space. The novel words consisted of novel stems ending in /k/ and the *-ity* suffix. I supplied the UR of the nonce stem, while the trained model filled in what it had learned about (1) The *-ity* suffix (*i.e.* is *-ity* a VS-triggering suffix?), and (2) When to change /k/ to [s]. 10 of the 11 generalized the alternation to novel stems (each at rates >95%). The one that didn't predicted that the novel stem alternates to [s] upon *-ity* suffixation at a rate <.01%. Overall, my model of simultaneous morphology-phonology learning predicted that a small but non-negligible minority will fail to learn the URs and grammar needed to extend the non-automatic alternation to novel stems. This prediction finds a parallel in the results of Pierrehumbert's wug study.

The non-generalizing model learned the UR /ɪlɛktɹɪsɪti/ for *electricity*, and a grammar where /k/ always surfaced as [k]. (Threshold for a "learned UR" was set at $P(ur|word) >$ 99%.) In contrast, all generalizing models had the UR $/\text{ɪlɛktɹɪk+ɪti}_{\text{VS-suffix}}/$, and the grammar for /k/→[s] $/\_+\{\text{ɪ}, \widehat{\text{aɪ}}\}_{\text{VS-suffix}}$. Thus, speakers may split the labour between the lexicon and morphophonology differently for the same surface VS pattern *(i.e.* inter-speaker variation).

My model, which simultaneously learns morphology and phonology, predicted: (1) When morphology isn't learned, neither is the non-automatic phonological alternation; instead, the unanalyzed form (*e.g.* /ɪlɛktɹɪsɪti/) is stored in the lexicon. (2) This occurs at a low yet non-negligible rate for velar softening. The Pierrehumbert study provided supporting evidence that a small minority of humans may indeed fall into this category. For these individuals, not acquiring the required morphology and grammar for the non-automatic [k]∼[s] alternation results in them losing the generalization to novel stems. This in turn could contribute to the diachronic instability of velar softening.

Alderete, J. 1999. *Morphologically governed accent in Optimality Theory*. Ph.D. thesis, UMass, Amherst. **Goldwater, S; Johnson, M**. 2003. Learning OT constraint rankings using a maximum entropy model. In: J Spenader, A Eriksson, Ö Dahl (ed.), *Proceedings of the Stockholm Workshop on Variation within Optimality Theory*. 111–20. **Itô, J; Mester, A**. 1999. The Phonological Lexicon. In: N Tsujimura (ed.), *The Handbook of Japanese Linguistics.* Oxford: Blackwell. 62-100. **Jarosz, G**. 2015. *Expectation Driven Learning of Phonology*, Ms. UMass, Amherst. **Pierrehumbert, J**. 2006. The statistical basis of an unnatural alternation. In L Goldstein, DH Whalen, CT Best (ed.), *Laboratory Phonology VIII*. Berlin: Mouton de Gruyter. 81–107. **Zuraw, K**. 2000. Patterned exceptions in phonology. Ph.D. thesis, UCLA.