

Learning-theoretic generative phonetics with maxent: a Hausa example

Bruce Hayes

(University of California Los Angeles)

In generative phonetics we seek to form grammars that can quantitatively predict speakers' phonetic outputs. The field has achieved a number of successes, notably Beckman & Pierrehumbert's (1988) grammar for Japanese pitch contours, but research has been difficult and progress slow. A plausible route to improvement would be to let learning algorithms play a role in the construction of grammars (Boersma 2001), a strategy that might be especially effective with maxent grammars (Smolensky 1986, Goldwater and Johnson 2003), whose learning algorithm is known to be convergent. Flemming (2001) contributed another important observation: that phonetic Harmonic Grammars (of which maxent grammars are a stochastic variant) can capture "compromise effects"; in Flemming's case the well-known locus effect for CV sequences, whereby the F2 value observed at consonant release is a compromise between conflicting requirements of the vowel and consonant targets.

I report a small-scale study employing all of the above ideas, whose goal is to model the durations of syllables in sung quantitative verse in Hausa (Hayes and Schuh, in progress). This is a simple, manageable domain, forming a useful testing bed for the theoretical ideas. Our weighted constraints follow Flemming in penalizing squared divergences from a target value. Targets are estimated as the smoothed mode of the statistical distribution of a phonological or rhythmic category; and the four categories employed are syllable, mora, eighth note, and metrical foot. Our model shows appropriate qualitative behavior; like Flemming's it captures the outcome that phonetic surface values are compromises between conflicting demands. It also is reasonably accurate in capturing the quantitative data, including the patterning of variation.