Morphologically-conditioned tonotactics in multilevel Maximum Entropy grammar

Stephanie S Shih
University of California, Merced

Sharon Inkelas
University of California, Berkeley

This paper presents a novel approach to probabilistic lexically-conditioned tonotactics, featuring a case study of Mende in which tonotactics vary by lexical category. The study contributes to the understanding of morphologically-conditioned phonology in several ways. First, the observed part-of-speech sensitivity goes beyond the noun-adjective-verb distinctions noted by e.g., Smith 2011, more closely resembling the complexity of morphophonological variation, as addressed in both single grammar (e.g., indexed constraints: Itô & Mester 199; Alderete 2001) and multiple grammar (e.g., cophonologies: Anttila 2002; Inkelas & Zoll 2005) theories. Second, the variation is not just a matter of differential faithfulness; it involves markedness reversals of the kind that Alderete's 2001 'grammar dependence' hypothesis predicts impossible (cf. Pater 2009). Third, the study models not just the space of variation but also the frequency of variation. This is accomplished by indexed weight adjustments for each constraint (i.e., 'varying slopes'; see also Coetzee & Pater 2011) in a Maximum Entropy Harmonic Grammar (MaxEnt HG; Goldwater & Johnson 2003). Couched in multilevel statistical models, the approach presented here unites the treatment of lexical class-sensitive phonotactics with the treatment of morphophonology, and directly addresses the overarching issue in morphophonology of how to quantify the heterogeneity that morphological conditioning can engender in a phonological system.

Early generative accounts of Mende surface tonotactics dealt only with nouns and focused on the majority tone patterns, modeled by a pre-specified, limited set of five surface tone 'melodies' (H, L, HL, LH, LHL), mapped onto syllables by universal autosegmental processes (e.g., Leben 1978). However, as discussed by subsequent studies (Dwyer 1978; Conteh et al. 1983; Zoll 2003; Zhang 2007), many surface patterns deviate from the supposed five melodies and their universal autosegmental association principles. Inkelas & Shih 2015 argue for abandoning the original autosegmental insights and modeling tonal patterns in Mende nouns with general similarity- and proximity-driven surface correspondence, an approach that has recently gained traction for both phonotactics and phonological alternations (e.g., Frisch et al. 2004; Hansson 2001; Rose & Walker 2004; Wayment 2009; Bennett 2013). The basic (violable) insights for Mende tonotactics, as formalised in ABC+Q by Inkelas & Shih 2015, are as follows:

(1) *Contour tones are avoided.* CORR-q::q, IDENT-XX [tone] mandate agreement between subparts of a segment (*q*). E.g., *[ǎ], ✓[à].

(2) *If necessary, contour tones are tolerated at the right edge.* CORR-$[q_w::q_w]_\sigma$, IDENT-XX [tone] mandate agreement between subsegments (*q*) within nonfinal ('weak') syllables. E.g., *[ǎ.à], ✓[à.ǎ].

(3) *Tone changes align with syllable boundaries (more syllables leads to more non-level tone patterns).* qq-EDGE σ, CORR-q::q prevent correspondence (and resulting tone agreement) across syllable boundaries. E.g., *[á.á], ✓[à.á].

(4) *HLH troughs are avoided.* CORR-q[H]q[H], q[H]q[H]-qADJ mandate correspondence and adjacency of H tones. E.g., *[á.à.á], ✓[à.á.á].

(5) *Words preferably have at least one H tone.* HAVE H mandates the presence of one H tone. E.g., *[á.á], ✓[à.á].

This study argues that variation across part of speech in Mende can be captured in MaxEnt HG in terms of the degree to which the basic tonotactic principles in (1)-(5) (the 'Base Grammar') are followed in each lexical class (inspired by e.g., Anttila 2002 in classic OT). The data come from a cor-

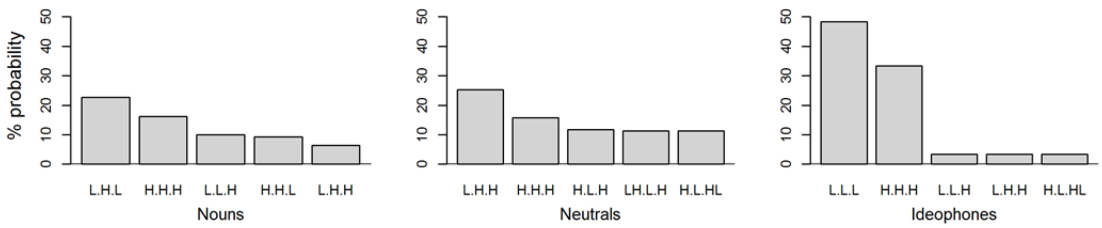

Figure A. Observed % probability for top 5 most frequent trisyllabic surface tone patterns per lexical class.

pus lexicon developed from Innes' 1969 Mende dictionary. Results are reported here from the three largest lexical classes: nouns (*n*=2707), neutrals (i.e., verbs/adjectives) (*n*=792), and ideophones (*n*=546). Figure A illustrates the top 5 most frequent observed tone patterns for trisyllabic words. Relative tone pattern frequencies vary by lexical class: e.g., Nouns prefer the L.H.L pattern; Neutrals, the L.H.H pattern; Ideophones, the L.L.L pattern.

We provide an analysis in Maximum Entropy grammar, for which the output is a probability distribution over all possible surface tone patterns, per the number of syllables and lexical class of the word. Morphological conditioning is modeled as an additive, lexical class-sensitive weight adjustment for each constraint in the Base Grammar: e.g., w_1·CORR-q::q + w_2·(CORR-q::q×NOUN) + w_3·(CORR-q::q×NEUT). In essence, each lexical class is allowed varying slopes for every model parameter, formally executed here as interaction terms (cf. e.g., Gelman & Hill 2007), and overall tonotactics are predicted by the main base weights for the constraints.

The results of the varying-slope approach accurately capture both shared features across lexical classes and class-specific morphological conditioning on the distribution of surface tone patterns. The Base Grammar for Mende tonotactics reveals the importance of contour tone avoidance and contour tone alignment to the right edge (see 1–2): *w*(CORR-q::q)=1.175, *w*(CORR-$[q_w::q_w]_\sigma$)=0.815, other constraints are *w*=0. This is true across all lexical classes, reflecting universal dispreferences for (nonfinal) contour tones (Gordon 2001; Zhang 2004). Constraint weights also vary by class. Nouns and neutrals are more similar to each other than to ideophone tonotactics. Both nouns and neutrals exhibit tone disagreement across syllable boundaries (see 3; *w*(qq-EDGE×NOUN)=0.16, *w*(qq-EDGE×NEUT)=0.18), whereas ideophones preferentially feature more level tones across the board (*w*(qq-EDGE×ID)=0): this pattern can be observed in Figure A. Nouns and neutrals also show greater affinity for the requisite H tone than ideophones (5), with neutrals leading the trend: *w*(HAVEH×NEUT)=1.32, *w*(HAVEH×NOUN)=0.41, *w*(HAVEH×ID)=0. Differences between nouns and neutrals include a greater avoidance of HLH troughs for nouns (4) (*w*(q[H]q[H]-qADJ×NOUN)=1.51, *w*(q[H]q[H]-qADJ×NEUT)=0) and a greater preference for H tones and transitions at the word-final syllable for neutrals: in fact, the adjusted weighting of CORR-q::q and CORR-$[q_w::q_w]_\sigma$ for neutrals is reverse that of the base grammar.

Reinterpreting cophonology subgrammars as indexed weight adjustments of the basic grammar captures the insights of indexed constraints and cophonology approaches in the same system. Our approach shifts the burden of deciding which constraints require exceptional indexation to the grammar itself (cf. Pater 2009): this is necessary in particular for probabilistic phonotactic applications (see e.g., Coetzee & Pater 2011). While previous approaches to language-internal morphophonological variation have focused on constraining it (Kiparsky 1982; Alderete 2001; Smith 2011), we still have only a rudimentary understanding of the quantitative *extent* to which variation (i.e., entropy) is possible within a coherent grammar. This case study offers a way to quantitatively probe the heterogeneity, and suggests the potential value of examining probabilistic morpho-phonotactic variation, which is almost certainly not unique to Mende (see e.g., Arabic, Japanese). Morphologically-conditioned phonotactics are potentially an important cue for part of speech, with consequent implications for language processing and acquisition.