

Phone Informativity

Using the amount of information encoded in different phones helps predict phonological phenomena such as deletion and change. I present the results of a number of corpus studies that demonstrate this claim. Markedness accounts would have difficulty accounting for phone deletion rates and phone durations in American English, since /p/ is longer and less likely to delete than /k/, but /b/ is shorter and more likely to delete than /g/, requiring a reversal of the markedness hierarchies for voiced and voiceless stops. Phone frequency-based accounts do not fare much better, since /ŋ/ is less frequent than /m/ but is shorter and more likely to delete. Moreover, these effects cannot be reduced to local predictability, as uninformative phones tend to delete more even in contexts where they are unexpected. Phone informativity, measured as the mean surprisal of seeing a phone in context across the entire language, handles these asymmetries with ease, as the different deletion rates and durations are correlated with the informativity of these pairs.