Phonetic universals – abstract structure vs fine detail

Anonymous
Erewhon

Whether or not one believes in strong universals in the Chomskyan sense, or only in the statistical preponderance of weak universals, one naturally seeks grounding for such universals in phonetic facts – if facts they be – of articulation or perception. For some decades now, researchers have sought to provide evidence for universals by conducting computer simulations. One of the best known earlier works along these lines is de Boer's (2001) work on the evolution of vowel systems, in which a system of agents, with fairly detailed phonetic and perceptual models of vowel space, 'evolved' vowel systems from scratch by playing imitation games, occasionally inventing new 'phonemes', and otherwise trying to keep phonemes apart. De Boer found that reasonably 'natural' vowel systems evolved. More recently, simulations with more explicit phonological content have become popular, such as Boersma and Chladková's (2010) work on features in vowel systems, or J. Kirby's (forthcoming) work on probabilistic phonologization.

When designing such simulations, an apparently critical decision is how much phonetic detail to include. For example, if dealing with vowels, does one simply deal with points in a two or three-dimensional abstract Euclidean space, corresponding to a position in the IPA vowel chart? Does one build as detailed a model as possible, with an articulatory synthesizer at one end, and a state of the art speech recognition model at the other? Or does one compromise, and use a crude but tractable model of these things, such as identifying vowels by two or three formants, and using simple functional models of perception. Generally researchers put as much detail as they can – but it is questionable to what extent the results require this detail. De Boer, for example, started off (as related in his thesis, but not in the (2010) book) with a detailed model; he got no results, so fell back to a compromise model of the kind just sketched.

Moreover, our own experiments in replicating de Boer's work indicate that in so far as his results do match naturally occurring vowel systems (which has been challenged in a review (Donegan 2004)), the results are unpleasantly sensitive to a number of rather arbitrary parameters of the model.

We have therefore engaged in a small project, on whose progress this presentation will report, to look at the dependence of such works upon the level of detail employed in the model. The first part of the programme adopts de Boer's framework, but pulls all the way back to a purely abstract model, no more than the vowel quadrilateral (or rather cuboid). Strikingly, we can obtain very similar results to de Boer, though we have discarded almost all his phonetic modelling. Because of the simplicity of the model, it is much easier to simulate and analyse the effects of varying the different parameters (such as how close together two 'phonemes' may be before merging).

In the second part, on-going and due to complete by October, we explore the effects of gradually enriching the model of 'phoneme', from the simple point used by de Boer to a model of convex polyhedra, which gives most of the power of exemplar models without the memory overload. This model also allows us to investigate whether de Boer and Chladkavá's features can exist as stable but emergent properties (cf. Mielke 2008) of a simple model, rather than being built in to universal phonetic grammar.

The final part will gradually re-introduce more detailed phonetic modelling, using regression testing (in the sense of software engineering rather than statistics) techniques, in order to investigate what influences there are upon the robustness or otherwise of results.