From input distributions to graded internal structure: artificial learning of novel phonetic categories

Paul Olejarczuk & Vsevolod Kapatsinski
University of Oregon
paulo@uoregon.edu vkapatsi@uoregon.edu

This paper presents the results of a phonetic learning experiment investigating the relationship between the shape of input distributions and the internal structure of the categories inferred from these distributions. Specifically, we ask whether category goodness is (a) proportional to token frequency, suggesting veridical exemplar storage, (b) centered around the mode of the input distribution, suggesting a symmetry bias and the abstraction of a 'modal prototype', or (c) symmetrical around the mean of the input distribution, reflecting a process of perceptual averaging.

Two groups of native-English listeners received passive exposure to a novel phonetic category: a monosyllable with a rising-falling (LHL) tone. The magnitude of the pitch excursion varied, and each group was exposed to a different training distribution: the Left distribution featured negative skew, while the Right distribution featured positive skew. Crucially, both training distributions had identical means and variances. Following training, both groups rated the category goodness of tokens that fell within and outside the range of experienced exemplars.

The figure below presents the results superimposed over the training distributions. Overall, the ratings distributions were less peaky than the input, suggesting regions of perceptual equivalence (c.f. Kuhl, 1991; Pierrehumbert, 2001). Both groups were also biased against low-magnitude excursions. As for the hypotheses, the ratings were neither directly proportional to token frequency nor centered around the modes. Category goodness was roughly symmetrical, but the two distributions were not identical, arguing against a strong version of the perceptual averaging account. Surprisingly, the Right group's ratings were shifted *rightward* relative to the Left group ($\beta = 2.02$, $t(6.4) = 4.19$, $p < .001$), suggesting at least two possible mechanisms: (a) the subjects overestimated the probability of rare exemplars, reflecting an 'atypicality bias' (c.f. Tanaka et al., 1998 in face categorization), or (b) the steep drop-offs near each mode cued a category boundary, resulting in a shift of the entire goodness distribution away from this induced boundary. Taken together, the results question both the view that exemplar influence is always proportional to frequency, and the view that categories are stored as parametric summaries of input distributions.

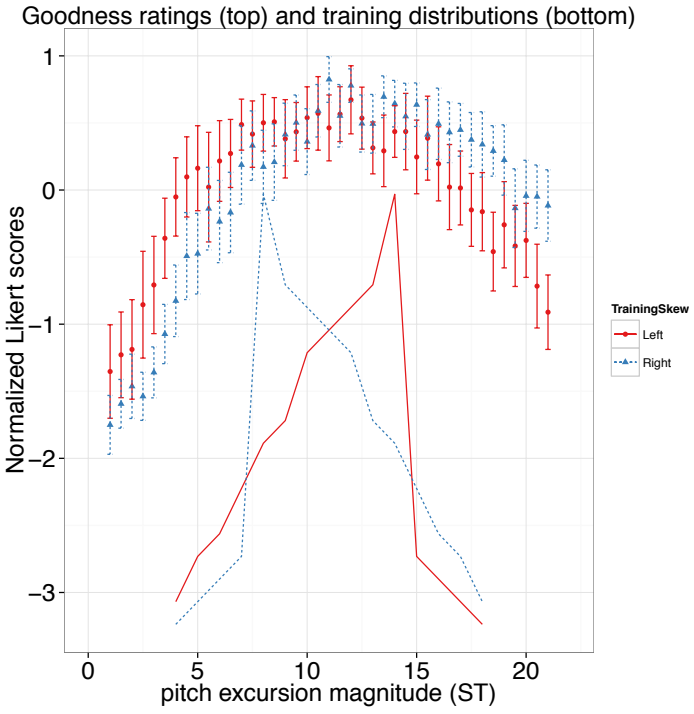

References:

Kuhl, P.K. (1991). Human adults and human infants show a "perceptual magnet effect" for the prototypes of speech categories, monkeys do not. *Perception and Psychophysics, 50,* 93-107**.**

Pierrehumbert, J. B. (2001). Exemplar dynamics: Word frequency, lenition, and contrast. In Joan L. Bybee and Paul Hopper, (eds.), *Frequency and the Emergence of Linguistic Structure*, 137-157. John Benjamins.

Tanaka, J., Giles, M., Kremen, S., & Simon, V. (1998). Mapping attractor fields in face space: the atypicality bias inface recognition. *Cognition, 68(3)*, 199–220.