

Gibbs Sampling as an Experimental Design: Prospects and Opportunities in Phonetics

Bill Thompson^{1,2} & Bart de Boer¹

¹Artificial Intelligence Laboratory, Vrije Universiteit Brussel

²Language and Cognition Department, Max Planck Institute for Psycholinguistics

Computational and experimental approaches to phonetics entail complementary strengths and weaknesses. Computational methods provide rigorous inferential frameworks for drawing statistical insights from data, but cannot yet provide adequate formal structures for many of the most interesting cognitive abilities underpinning human phonetic capabilities. Experimental methods make available the full suite of human phonetic abilities, but typically lack a formal framework for organizing data collection in a manner that guarantees convergence on valid posterior inferences in complex or high-dimensional problems. An emerging paradigm in cognitive science combines the strengths of these two approaches by constructing experiments that implement stochastic sampling algorithms using human minds as the conditional samplers. In this talk, I'll introduce this paradigm and explore the ways in which the phonetic sciences might benefit by doing something similar using human ears. As an example, I'll propose a use case for this paradigm in a substantial open problem at the intersection of phonetics and machine learning: identifying sets of primitive acoustic units in unfamiliar streams of connected speech, or *unsupervised induction of phonemic inventories via joint segmentation and clustering*. I'll suggest that although this joint inference problem may be too difficult for individual humans to solve efficiently in experimental settings, this emerging paradigm may provide a solution. Contemporary machine learning methods provide an explicit statistical model of this problem and a Gibbs sampling algorithm for posterior inference, but are constrained in their usefulness by limitations to the acoustic models that represent phonemes. I propose to implement this posterior inference algorithm in humans, using iterated learning to approximate the stochastic process; to construct experiments in which chains of participants alternate between performing segmentation tasks conditional on the previous participant's clustering decisions, and clustering decisions based on the previous participant's segmentations. More generally, I hope to stimulate discussion around the idea that human experimental participants can often be understood as samplers: as such, posterior sampling methods provide a rich collection of blueprints for experimental design in phonetics.