

Variation-driven speech perception

Meghan Sumner
Stanford University

A single word is produced differently each time it is uttered. This results in a wide range of phonetic and phonological variation across multiple productions of a single word. Across accents, variation adds complexity for listeners – sounds, each with their own range of phonetic variation, are substituted for one another. Consider the word *card*. This word may sound like the word *cod* when produced by a non-rhotic speaker from Boston, the word *cawed* when produced by a non-rhotic speaker from New York City, the word *guard* when produced by a native speaker of Spanish, or the word *cart* when produced by a native speaker of German. As listeners, we must navigate through this variation to understand a number of acoustic forms as an instance of one word, and not another word. This is particularly challenging considering that minimal differences are oftentimes meaningful. A central goal in the field of speech perception is to understand how listeners take such a variable speech signal and map it to meaning.

Over the past fifteen years, evidence has grown in support of the notion that lexical representations are rich with acoustic detail. This evidence has supported multiple exemplar theories of speech perception utilizing exemplar dynamics. In other words, this research has yielded a variety of theories that, one way or another, account for variability not by discarding it and mapping a variable acoustic signal onto discrete, abstract categories, but by storing this information and utilizing it when a similar signal is encountered.

This talk focuses on effects of phonological and phonetic variation and how variation is integrated with well-established frequency effects. I review a handful of recent studies with data that exhibit the pattern expected by frequency-based accounts of speech perception; namely, how often a listener hears a word with a particular phonological variant predicts how well a listener understands that word. Then, I discuss in greater detail the subtle differences that arise once we consider the interaction of phonological variant frequency with both low-level (word-level phonetics) and high-level (listener attitudes toward an accent) factors. I present the results from two ongoing projects that show phonological variant frequency is tightly tied to immediate (within a word) and more distant (surrounding words) contexts within accents and that, once controlled, raw frequency is insufficient to account for word recognition effects across accents. I discuss the implications and predictions that fallout from these projects for theories of speech perception.