Abstract for talk given by Gina Kuperberg at the Tufts University Cognitive and Brain Sciences Series, in Medford, MA on October 15, 2013

The Neurobiology of Language Comprehension:
A Hierarchical Bayesian Architecture
Gina R Kuperberg MD PhD

Language comprehension presents us with an enormous challenge. To determine what message a speaker or writer intended to communicate, we must decode a sequence of letters or sounds that unfold rapidly in real time, often in noisy environments, in the face of ambiguity at all levels of the language code.

Fortunately, to help us accomplish this task, we can draw upon all the real-world and semantic knowledge stored within our memories. There is now robust evidence that our brains can use context to mobilize this stored knowledge very quickly, sometimes ahead of the bottom-up input: information that is predicted correctly can facilitate semantic processing of incoming words, reflected by reduced activity within the left anterior temporal cortex within 300ms of their onset.

This is all well and good if our predictions are correct. In the real world, however, we often encounter input that is unexpected. And when this happens, we will pay a price if we have predicted incorrectly. I will discuss a large body of evidence from multimodal neuroimaging studies suggesting (a) that this price manifests as prolonged neural activity within a fronto-parietal network between 500-800ms after the onset of words that disconfirm our predictions, and (b) that different components of this network are recruited depending on both the representational level and certainty with which these predictions were generated.

I will argue that this literature tells us something fundamental about the neurobiology of language comprehension—that it can be conceptualized within a hierarchical Bayesian framework. First, it suggests that our predictions are generated in a probabilistic manner, with quite different consequences depending on whether, at any given moment, they are violated by input that falls within or outside the hypothesis space that defines our current beliefs. Second, it suggests that we predict at multiple representational levels: sometimes just the semantic features and syntactic properties of upcoming words, and at other times, down to the level of specific lexical form. Finally, it suggests that the costs incurred when our predictions are violated can be conceptualized as learning signals, enabling us to update our beliefs in rapid response to bottom-up input. This, in turn, can explain why and how our brains are able to respond so dynamically and flexibly to the ever-changing demands of our communicative environment.