

On computational and behavioral evidence regarding Hebbian transcortical cell assemblies

Michael Spivey, Mark Andrews, and Daniel Richardson

Department of Psychology, Cornell University, Ithaca, NY 14853.

{spivey;mwa1;dc18}@cornell.edu

www.psych.cornell.edu/Psychology/Spivey/homepage.html

www.psych.cornell.edu/graduates/people/Mark W. Andrews.htm

www.psych.cornell.edu/graduates/people/Daniel C. Richardson.htm

Abstract: Pulvermüller restricts himself to an unnecessarily narrow range of evidence to support his claims. Evidence from neural modeling and behavioral experiments provides further support for an account of words encoded as transcortical cell assemblies. A cognitive neuroscience of language must include a range of methodologies (e.g., neural, computational, and behavioral) and will need to focus on the on-line processes of real-time language processing in more natural contexts.

Understanding the interaction between the perceptual modalities and the cortices that subserve them is a goal that has too long played second fiddle to the more popular goal of “finding the boxes in the head.” Although many perceptual/cognitive functions exhibit some degree of cortical localization, the continuous communication between these “modules” is at least as important as their anatomical separation. Indeed, if Pulvermüller is right about the Hebbian transcortical cell assembly being the brain's primary “unit of representation,” then a solitary punctate region of cortex whose averaged activity correlates with a broadly defined class of perceptual/cognitive tasks may *not* be especially relevant for understanding perception and cognition. Instead, what we should study are more specific environmental instances (ecologically valid combinations of stimuli) and sequences thereof, and *the multiple cortical regions that participate* in representing those environmental instances.

Pulvermüller's shortcoming is that he unnecessarily limits his evidence predominantly to one methodology (neuroimaging), when a wide range of methodologies provide insight into “words in the brain's language.” This commentary points to the relevance of two areas of cognitive neuroscience that Pulvermüller barely mentions: neural modeling and behavioral experiments. By also looking at computational and behavioral results, one can see a more complete picture of how a proposal like Pulvermüller's may be manifested at multiple levels of description: from the individ-

ual synaptic weights that *produce* the cell assemblies to the motor output that *is produced by* the cell assemblies.

Neural networks. If Pulvermüller's proposal is correct, it indicates that one's representation of a word includes multi-modal information not traditionally considered linguistic (e.g., perceptual features, motor routines). He suggests that our mental representation of a word might best be described as a pattern of activation widely distributed across a disparate set of information processing units. In the field of neural networks, this notion of a “distributed representation” has made significant advances since Hebb's original thesis (far more than is implied by Pulvermüller's brief nod to the neural network literature; Gutfreund & Toulouse 1994).

As far as they go, the neuroimaging results are substantial. However, the descriptions of the physiology of these cell assemblies, the spread of neural activation throughout them, and their temporal dynamics, are left somewhat vague and impressionistic. This is, of course, less a criticism of Pulvermüller's review than a reflection on the lack of precision of current neuroimaging techniques. These limitations will undoubtedly be somewhat remedied with improved technology. For now, though, further insight into these more specific issues can come from the study of biologically plausible artificial neural networks. Far more in-depth analysis and experimentation can be performed on artificial neural networks than will ever be possible with real brains.

Of particular interest are attractor networks, which represent words as distributed representations and have some temporal dynamics leading to the full ignition of a cell assembly (McRae et al. 1997). Moreover, models that can integrate sensory information from separate modalities during learning (de Sa & Ballard 1997) can provide examples of possible developmental trajectories for Hebb-inspired accounts of lexical acquisition.

Behavioral experiments. As important as it is to “zoom in” on the subcomponents of the cell assembly, we also need to see the “wide angle” view that includes what neural representations are good for: behavior. Of course, some behaviors are better than others at elucidating the temporal dynamics of a cell assembly attempting full ignition. For example, Munakata et al. (1997) use a developing distributed representation to explain how some infants who are still making the classic A-not-B error in their *reaching* behavior actually show signs of being aware of the correct solution in their *eye movement* behavior. Behaviors, such as eye movements, that are sensitive to probabilistic information can provide converging evidence on the partially activated representations (cell assemblies) that are computed in real time while the brain gravitates toward a stable state to encode an environmental instance.

As another example, when adult participants are instructed to “pick up the candy,” they never *reach* for incorrect objects, but they often *look* at incorrect objects whose names have similar phonology such as a candle (Tanenhaus & Spivey-Knowlton 1996). Additionally, as Hebb himself predicted, oculomotor representations are activated when participants are imagining an event or remembering an object. While listening to a story about a train going past, participants made predominantly horizontal eye movements; and when asked to recall a missing object on a grid, there were eye movements toward the vacated square (Spivey & Geng 1998).

Summary. To study anything in “the brain's language,” we must be as opportunistic with *our* information sources as the brain is with *its*. We need convergence of methodologies. We need to map out in detail the cycle from perception to action. In contrast to neuroimaging, neural modeling simultaneously allows the spatial and temporal resolution necessary to explore how a particular cell assembly might approach its stable state over dozens of timesteps (Tabor et al. 1997). Moreover, behavioral experiments allow us to observe the pragmatic consequences of these cell assemblies, realized as motor output grounded in the same environment as the original perceptual input, thus completing the cycle.