Synchronous gain and predictive coding

An obvious candidate for controlling post-synaptic gain is synchronization of pre-synaptic inputs—a phenomena referred to as *synchronous gain* (Chawla, Lumer, & Friston, 1999). This means that the selection of prediction errors—that drive higher-level representations—almost certainly involves synchronization. Indeed, there is current interest in the possibility that bottom-up messages—from superficial pyramidal cells —are mediated by fast (gamma) frequencies, while top-down messages from deep pyramidal cells may be mediated by slower (beta) frequencies (Buffalo, Fries, Landman, Buschman, & Desimone, 2011). It is this hypothesis that current collaborations with Pascal Fries and colleagues hope to test—using dynamic causal modeling (Bastos et al., 2012).

Empirical predictions

Finally, I will reiterate the importance of formal theories and modeling—as emphasized by Gotts et al.— by commenting on the empirical predictions made by predictive coding. First, repetition suppression rests on optimizing connection strengths that mediate predictions. Crucially, these change (anti-symmetrically) the efficacy of *both forward and backward* connections (Friston, 2008). Second, because predictive coding minimizes prediction error, it is based upon feedback dynamics. This means that either forward or backward connections must be (effectively) inhibitory. The fact that both forward and backward connections are excitatory (Glutamatergic) has exercised us a little. Current thinking is that *explaining away* is mediated by local inhibitory interneurons (Bastos et al., 2012). Finally, repetition suppression is expressed throughout the hierarchy (in high and low areas) *at the same time*. This is because message-passing is recurrent and suppression of prediction error emerges concurrently at all levels. Repetition suppression to high-level attributes will clearly occur later but it will be expressed at lower levels. This phenomenon has been studied extensively in the context of the simplest repetition suppression— namely the mismatch negativity (Garrido, Kilner, Stephan, & Friston, 2009).

In conclusion, I think Gotts et al. raise a number of fascinating questions that may herald some important advances in our understanding of computational architectures in the brain, over the next few years.

* * *

Explaining away repetition effects via predictive coding

Michael P. Ewbank and Richard N. Henson

MRC Cognition and Brain Sciences Unit, Cambridge, UK
E-mail: michael.ewbank@mrc-cbu.cam.ac.uk

Abstract: Gotts, Chow and Martin summarize Predictive Coding models in which repetition-related decreases in neural activity reflect an "Explaining Away" of stimulus-driven neural activity. Here we elaborate the subtleties of testing such models, particularly with fMRI.

The "Explaining Away" model described by Gotts et al. is really the application of a more general doctrine in neuroscience—that of "predictive coding" (Friston, 2012)—to the case of repetition effects. The key idea is that neurons receive predictions from higher layers of a hierarchical network, with any difference between those predictions and the input from lower layers producing a prediction error in that layer. Synaptic change serves to reduce future prediction error (i.e., improve predictions), resulting in reduced activity in those neurons coding the prediction error when a stimulus is repeated.

In the specific instantiation of predictive coding discussed by Gotts et al., each layer contains three types of neurons: Not just those coding prediction error, but also those coding predictions (from higher layers) and input (prediction errors from lower layers). Yet the relative contribution of these different types of neurons to a hemodynamic measure like BOLD is uncertain (see Egner, Monti, & Summerfield, 2010), making such models difficult to test with fMRI. Testing may be easier with EEG/MEG though, given that Friston (2008) makes a specific claim that the cortical neurons coding prediction error are the large, supra-granular pyramidal neurons, thought to make the dominant contribution to the EEG/MEG signal.

Regarding experimental paradigms to test predictive coding, it is important to note that the recent debate about whether expectation of repetition does, on the basis of human fMRI and EEG (e.g., Summerfield, Wyart, Johnen, & de Gardelle, 2011), or does not, on

the basis of monkey single-cell recording (Kaliukhovich & Vogels, 2011), modulate repetition suppression is actually somewhat parenthetical to predictive coding. This is because the "predictions" manipulated in the Summerfield et al. paradigm are likely to be conscious/strategic (and so may be less prevalent in monkeys). Yet the "predictions" in predictive coding theory are automatic, intrinsic properties of the brain networks that do not necessarily depend on conscious expectation. Thus while the effects of higher-order expectancy are clearly interesting and important (and probably generated by prefrontal regions that act on the ventral stream), the lack of such expectancy effects in other paradigms (Kaliukhovich & Vogels, 2011; Larsson & Smith, 2012) should not be used to reject predictive-coding models.

Another approach used to support predictive coding models of repetition suppression is to examine changes in connectivity between brain regions. Our own work, for example, has used Dynamic Causal Modelling (DCM) of fMRI data to show that repetition of bodies (Ewbank et al., 2011) or faces (Ewbank, Henson, Rowe, Stoyanova, & Calder, in press), at least across different images, modulates backward connections from "higher" regions in fusiform cortex to "lower" regions in extrastriate occipital cortex. Gotts et al. wondered why this modulation by repetition reflected a more positive coupling parameter in the DCM, when according to predictive coding, one might expect a more negative coupling associated with the suppression of prediction error in lower regions by higher regions. Again, however, the precise interpretation is more subtle because we do not know which types of excitatory/inhibitory neurons contribute to the BOLD signal. Moreover, due to high interdependency between parameters in such recurrent DCMs, inference is often more appropriate at the level of model selection rather than model parameters (Rowe, Hughes, Barker, & Owen, 2010). Thus, although we discussed our results in terms of predictive coding, the main conclusion of the Ewbank et al. papers (which were based on model selection) is that repetition suppression is not purely a local phenomenon (such as sharpening or even neuronal fatigue; Grill-Spector, Henson, & Martin, 2006), but also entails interactions between brain regions. This claim is consistent with both predictive coding and synchrony theories.

A further reason why DCM for fMRI may be limited in its ability to distinguish theories like predictive coding and synchrony is that the modulatory inputs (repetition in this case) need to be sustained over several seconds in order to have an appreciable impact on the network dynamics (Henson, Wakeman, Phillips, & Rowe, 2012). This is why we used a blocked design in the Ewbank et al. studies, where the modulation was assumed to operate throughout blocks. As Gotts et al. observe, such designs are undesirable from a behavioral perspective (e.g., encouraging use of conscious expectancies like those discussed above). Randomized designs (e.g., Henson, 2012) are clearly preferable, but in order to test for changes in effective connectivity as defined by dynamic models like DCM, data with higher temporal resolution are needed (e.g., Garrido, Kilner, Stephan, & Friston, 2009). Thus we agree with Gotts et al. that an exciting future direction is to examine connectivity, perhaps via synchrony, between regions using methods like EEG/MEG.

* * *

Repetition accelerates neural dynamics: In defense of facilitation models

Richard N. Henson

MRC Cognition and Brain Sciences Unit, Cambridge, UK
E-mail: rik.henson@mrc-cbu.cam.ac.uk

Abstract: Gotts, Chow and Martin give an excellent contemporary summary of the neural mechanisms that have been proposed to underlie the effects of stimulus repetition on brain and behavior. Here I comment on their Facilitation mechanism, and provide EEG evidence that repetition can accelerate neural processing.

Gotts et al. (2012) review four types of neural mechanism that might underlie the reduced brain response associated with repetition of a stimulus: Facilitation, Sharpening, Synchrony and Explaining Away. In particular, they make a case for mechanisms based on