

Hou, Q., Zhang, D., Jarzylo, L., Huganir, R.L., and Man, H.Y. (2008). *Proc. Natl. Acad. Sci. USA* 105, 775–780.

Hou, J., Gilbert, J., and Man, H.-Y. (2011). *Neuron* 72, this issue, 806–818.

Lee, M.C., Yasuda, R., and Ehlers, M.D. (2010). *Neuron* 66, 859–870.

Sacktor, T.C. (2008). *Prog. Brain Res.* 169, 27–40.

Sutton, M.A., Ito, H.T., Cressy, P., Kempf, C., Woo, J.C., and Schuman, E.M. (2006). *Cell* 125, 785–799.

Szobota, S., Gorostiza, P., Del Bene, F., Wyart, C., Fortin, D.L., Kolstad, K.D., Tulyathan, O., Volgraf,

M., Numano, R., Aaron, H.L., et al. (2007). *Neuron* 54, 535–545.

Turrigiano, G.G. (2008). *Cell* 135, 422–435.

Turrigiano, G.G., Leslie, K.R., Desai, N.S., Ruthford, L.C., and Nelson, S.B. (1998). *Nature* 391, 892–896.

Mechanisms of Selective Attention: Response Enhancement, Noise Reduction, and Efficient Pooling of Sensory Responses

John T. Serences^{1,*}

¹Department of Psychology and Neuroscience Graduate Program, University of California, San Diego, La Jolla, CA 92093-0109, USA

*Correspondence: jserences@ucsd.edu

DOI [10.1016/j.neuron.2011.11.005](https://doi.org/10.1016/j.neuron.2011.11.005)

In this issue of *Neuron*, Pestilli and coworkers provide evidence that response gain and noise reduction are insufficient to account for attention-induced changes in perception. Instead, selection may critically depend on the biased pooling of sensory signals during decision making.

The philosopher Malebranche noted in 1674 that “the mind does not pay equal attention to everything that it perceives. For it applies itself infinitely more to those things that affect it, that modify it, and that penetrate it, than to those that do not affect it and do not belong to it” (p. 412) (Malebranche, 1997). In the ensuing 300+ years, research on selective attention has continually progressed, and although we have made careful behavioral measurements using the tools of psychophysics, poked and prodded neural circuits with electrodes, and taken fancy pictures of human brains in action, we still have a vague understanding of how neuronal networks work in concert so that the mind “...applies itself infinitely more to those things that affect it...” Thus, we are rich in our knowledge of what and where, but poor in our understanding of how the brain prioritizes relevant over irrelevant sensory inputs. Here, Pestilli et al. (2011) use well-validated experimental and quantitative frameworks to evaluate the relative contribution of three candidate mechanisms by which selective information processing might operate: response enhancement,

noise reduction, and the efficient selection of sensory responses during decision making.

Response Enhancement

Over the last 35 years, most research has focused on the notion that selective attention operates by increasing the firing rate of neurons that are tuned to relevant spatial locations, objects, or features. Computationally, response gain should improve the reliability of neural signals as long as the variance of the firing rate does not increase faster than the mean. Attention-induced gain is also ubiquitous, extending from the earliest stages of cortical processing in the lateral geniculate nucleus (LGN) all the way through areas of frontal cortex, with the degree of response enhancement progressively increasing across the cortical hierarchy (from about 20%–30% in midlevel areas such as V4 to almost 100% in prefrontal cortex; Serences and Yantis, 2006; Treue, 2003).

Noise Reduction

More reliable encoding of relevant sensory inputs can also be achieved by

decreasing the variance of single neurons and by decreasing the degree of correlated noise across neural populations. Mitchell et al. (2007) showed that attending to an object reduced the ratio of the variance of the firing rate to the mean firing rate (the fano factor) by approximately 10%–20%. This reduction in relative variability should magnify any concurrent effects of response gain to further increase the reliability of neural codes. Ultimately, however, single neurons are too noisy to support perception: responses must be pooled from many neurons to achieve a stable representation. Unfortunately, averaging across multiple neurons will not attenuate biases induced by correlated noise, so decreasing moment-to-moment noise correlations between similarly tuned sensory neurons is generally thought to be beneficial. Although the issue is complex and still debated, several recent reports show that attention decreases pairwise correlations between neurons in midlevel areas V4 and MT and that these reductions are associated with improvements in behavior (Cohen and Kohn, 2011; Cohen and Maunsell, 2009, 2011; Mitchell et al., 2009).

Efficient Selection

Relying primarily on psychophysics and mathematical models, a parallel line of research has shown that many of the behavioral effects ascribed to selective attention can also be explained without resorting to response enhancement or to reduced neural noise. Instead, efficient selection can be achieved by assuming that decision mechanisms pool information only from those neural populations that are optimally tuned to discriminate the attended stimulus (Eckstein et al., 2009; Palmer et al., 2000; Shaw, 1984). These models are particularly effective at explaining how attention can greatly attenuate (or even eliminate) the influence of irrelevant distracting items that are simultaneously present in the scene (Palmer and Moore, 2009). Because information is only pooled from sensory neurons that optimally discriminate the relevant feature, the influence of irrelevant distracting items is naturally attenuated. In this sense, efficient selection operates via a form of noise reduction, albeit not at the level of variability (or covariability) in the firing rates of sensory neurons as discussed in the preceding section. Instead, selective pooling shunts interference from populations of sensory neurons that encode irrelevant features, thereby preserving the fidelity of neural signals associated with behaviorally relevant items.

Present Study

Few studies have formally linked the effects of attention on neural activity directly with the effects of attention on perception and behavior. Filling this void is obviously critical to understand the relative contributions from the three candidate mechanisms discussed above. To address this issue, Pestilli et al. investigated the influence of spatial attention on contrast-detection thresholds (i.e., the smallest noticeable difference between the contrast of two stimuli, or Δc) and on neural responses measured indirectly via functional MRI and the blood oxygenation level dependent (BOLD)

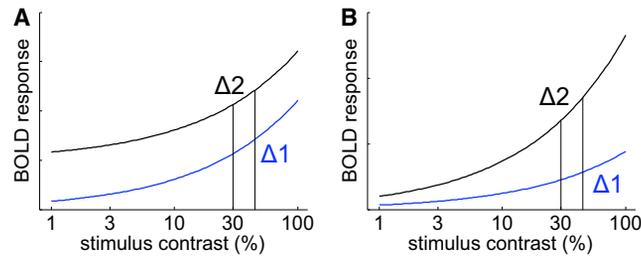


Figure 1. Simulated BOLD Contrast Response Functions Based on Equation 3 from Pestilli et al.

The relative change in response levels evoked by stimuli rendered at different contrasts (e.g., vertical lines in A) can be used to predict psychophysical contrast discrimination thresholds, forming a link between neural activity and behavior. Larger differential neural responses should correspond to a decrease in contrast discrimination thresholds.

(A) Simulation of BOLD contrast response functions observed by Pestilli et al. in the distributed attention condition (blue line) and in the focal cue condition (black line). Attention induced a purely additive shift of the contrast function, which should not, by itself, change contrast discrimination thresholds ($\Delta 1 = \Delta 2$).

(B) Same as (A), but focused attention induces response gain (black line). In this case, attention should decrease contrast discrimination thresholds ($\Delta 2 > \Delta 1$).

signal. The authors used contrast as the critical visual feature because a well-validated linkage hypothesis relates neural activity in early visual areas measured using either single-unit or BOLD signals with psychophysical data (Boynton et al., 1999; see Figure 1).

Here, the authors used a variant of a two-interval forced choice (2IFC) procedure. In the first temporal interval, one disk was presented in each quadrant of the visual field, and each disk was assigned a contrast from a range of “pedestal” values extending from 0% to 84%. This was followed by a blank period of 200 ms, and then a second array of four disks was presented in the previously occupied spatial locations. The contrast of a single disk was either slightly lower or slightly higher in the second interval, and the subject’s task was to indicate whether the first or the second display had the higher contrast disk. In half of the trials, subjects were given a spatial precue that indicated the target quadrant (focal attention cue), and in the remaining trials, a distributed attention cue indicated that all locations were equally likely to contain the target.

In this context, quantitative models posit that decisions are based on the application of a “max” rule that computes the temporal interval that contained the higher overall contrast level. In focal-cue trials, this max rule is applied only to stimuli

presented at the target location: the interval with the higher contrast determines the response. However, in distributed cue trials, the max rule is applied to a pooled estimate of the total contrast level across all stimuli in each interval. Not surprisingly, the authors found that subjects could detect a smaller contrast change (Δc) on focal-cue trials across the full range of pedestal contrast levels (see their Figure 3). Consistent with previous data, the authors also observed that focal attention increased the BOLD response at each contrast level by a constant amount (Figure 1A; Buracas and Boynton, 2007).

To account for improved behavioral performance, the authors largely discount response gain because the observed additive shift in the BOLD contrast response function should not improve discriminability (compare Figures 1A and 1B). However, the contribution of response enhancement to the observed increase in behavioral performance is nuanced, and I’ll return to this issue below. Next, a quantitative model that was constrained by the psychophysical data was used to show that neural responses would need to undergo not only an additive increase but also an unreasonably high 400% reduction in noise to adequately fit the BOLD data. Thus, response enhancement and noise reduction do not appear to be sufficient to account for observed improvements in behavior.

The authors then move on to show that the data can be explained by a relatively simple pooling framework. The overall pooled neural response in each stimulus presentation interval (R_p) was estimated as:

$$R_p = \frac{1}{4} \sqrt[k]{\sum_{i=1}^4 r_i^k},$$

where the parameter k determines the contribution that the response evoked by each stimulus (r_i) makes to the overall pooled response. When k equals 1, then all responses are given equal weight, and

this pooling operation is equivalent to averaging. As k increases, the most active neural populations will increasingly dominate the pooled response.

The model fits suggested that a single value of k ($k = 68$, toward the maximizing side of the spectrum) could account for behavioral performance on both distributed and focal cue trials because the stimulus location evoking the highest response will dominate the pooling. Recall that the contrast of the stimulus in each quadrant was assigned a random pedestal value ranging from 0% to 84%. Thus, on distributed cue trials, one of the nontarget locations should evoke the largest response on average, leading to an increase in the contrast change required for accurate discrimination (a larger Δc). On focused cue trials, the location evoking the largest response almost always corresponded to the target because of the attention-induced additive shift in the BOLD response (Figure 1a).

Although this pooling rule could account for the results, it is critical to note that like response enhancement and noise reduction, selective pooling is not sufficient. Rather, it was the pooling rule combined with additive response enhancement that led to improved perceptual acuity with focused attention (and the same principle would apply, given other forms of response enhancement as well; see main text, last paragraph starting on page 843). This finding is particularly exciting because it suggests that biased pooling rules might enable attentional gating by amplifying relatively modest changes in metabolically expensive response enhancement, thus maximizing perceptual selectivity while minimizing energy expenditure.

One major remaining question concerns the extent to which the value of k is systematically tied to the properties of the stimulus array. In a simple case in which only one stimulus is presented in a known position in the visual field,

pooling is largely irrelevant because there is only one location associated with an evoked response during each interval. In such sparse stimulus arrays, response enhancement and noise reduction probably play a dominant role. However, k should grow with the number of competing stimuli, because maximizing the influence of attention-enhanced responses should become increasingly important as distractor-evoked responses threaten to drown out relevant neural signals. Thus, a key avenue for future research will be to determine how k changes with the size and complexity of the search set and to understand whether and when k reaches asymptote (which may determine the upper limit on the effectiveness of pooling as a means of facilitating selection). Future studies should also determine whether k is strictly constrained by the expected scene statistics or whether there are situations in which attention directly operates by changing the pooling rule, perhaps as a function of cue reliability or learned priors about the probable location of relevant stimuli.

Conclusions

By focusing on the algorithm that pools information from the sensory neurons that are targeted by attentional gain and noise reduction, the authors provide exciting new empirical data regarding how selective information processing is implemented. Given that the same value of k fit the data on both focused and distributed cue trials, these results suggest that attention doesn't operate directly via manipulating the pooling of sensory information (at least in this context). Instead, a separate process may determine the value of k based on perceptual priors to optimally weight sensory inputs so that relatively modest changes in attention-induced gain and noise reduction can have a disproportionately large impact on perceptual decisions. Ultimately, the approach employed

by Pestilli et al. opens up many new avenues of inquiry, primarily because they laudably integrated branches of psychophysics, neurophysiology, and mathematical modeling that have unfortunately remained largely distinct; hopefully many other such efforts are soon to follow.

REFERENCES

- Boynton, G.M., Demb, J.B., Glover, G.H., and Heeger, D.J. (1999). *Vision Res.* 39, 257–269.
- Buracas, G.T., and Boynton, G.M. (2007). *J. Neurosci.* 27, 93–97.
- Cohen, M.R., and Kohn, A. (2011). *Nat. Neurosci.* 14, 811–819.
- Cohen, M.R., and Maunsell, J.H. (2009). *Nat. Neurosci.* 12, 1594–1600.
- Cohen, M.R., and Maunsell, J.H. (2011). *Neuron* 70, 1192–1204.
- Eckstein, M.P., Peterson, M.F., Pham, B.T., and Droll, J.A. (2009). *Vision Res.* 49, 1097–1128.
- Malebranche, N. (1997). *Malebranche: The Search after Truth* (Cambridge: Cambridge University Press).
- Mitchell, J.F., Sundberg, K.A., and Reynolds, J.H. (2007). *Neuron* 55, 131–141.
- Mitchell, J.F., Sundberg, K.A., and Reynolds, J.H. (2009). *Neuron* 63, 879–888.
- Palmer, J., and Moore, C.M. (2009). *Vision Res.* 49, 1045–1064.
- Palmer, J., Verghese, P., and Pavel, M. (2000). *Vision Res.* 40, 1227–1268.
- Pestilli, F., Carrasco, M., Heeger, D.J., and Gardner, J.L. (2011). *Neuron* 72, this issue, 832–846.
- Serences, J.T., and Yantis, S. (2006). *Trends Cogn. Sci. (Regul. Ed.)* 10, 38–45.
- Shaw, M.L. (1984). Division of attention among spatial locations: A fundamental difference between detection of letters and detection of luminance increments. In *Attention and Performance X*, H. Bouma and D.G. Bouwhuis, eds. (Hillsdale, NJ: Erlbaum).
- Treue, S. (2003). *Curr. Opin. Neurobiol.* 13, 428–432.