

Ghosts in the Decision Machine

Christopher Summerfield^{1,*} and Annabelle Blangero¹

¹Department of Experimental Psychology, University of Oxford, South Parks Road, Oxford, UK

*Correspondence: christopher.summerfield@psy.ox.ac.uk

<http://dx.doi.org/10.1016/j.neuron.2015.05.009>

Humans and monkeys occasionally report the presence of a stimulus that has not occurred. A new study by Carnevale et al. (2015) sheds light on the nature and timing of the neural mechanisms that give rise to false detections.

Most of the time, our subjective experience of the world accurately reflects the information arriving at our sense receptors. When I see a face in front of me, it is because another person is there; when my arm is itchy, it is because a mosquito has bitten me there. However, many of us will have experienced fleeting misperceptions: a subjective experience that does not have an obvious sensory precursor. These illusory perceptions often follow the receipt of information that predicts the forthcoming experience (Summerfield and de Lange, 2014). Hearing a ghost story late at night might change our interpretation of unexpected sights or sounds, or being told to watch out for bedbugs might make us spontaneously start scratching an imaginary insect bite. These experiences might be heightened when one is tired, or after the consumption of psychoactive substances, or in some psychiatric disease states, but they can also occur even in the healthy, sober, rested state. Why do these misperceptions arise, and what features of our neural circuitry make them possible? In this issue of *Neuron*, Carnevale et al. (2015) offer important new insights into the neural mechanisms that give rise to “illusory” reported perceptions.

The mechanisms that underlie perceptual decisions can be studied in the laboratory using psychophysical methods, where the intensity of sensory stimulation can be carefully controlled and the reported perception can be accurately measured. In one influential paradigm, macaque monkeys report the presence or absence of a mechanical vibration applied to the skin of one fingertip, while neuronal recordings are made from somatosensory, parietal, and premotor cortical regions. Vibration first excites neurons in early somatosensory cortices,

where firing rates correlate with the intensity of stimulation, but subsequent activity in secondary somatosensory, parietal, and premotor cortices correlates with the decision made by the monkey. In these regions, many neurons respond to the vibration only when the monkey correctly reports it had occurred (“hits”) but not when the monkey fails to report it (“misses”) (de Lafuente and Romo, 2005). This transition from sensation to a perceptual decision signal seems to occur gradually, both across time and along different stages of the cortical hierarchy (de Lafuente and Romo, 2006).

However, like every psychophysical detection paradigm, this approach yields another class of error trial: those on which the monkey reported a perception, even in the absence of stimulation (“false alarms”). Studying false alarms can help us to arbitrate among competing theories of how choices are formed. In the visual domain, reported misperceptions are associated with levels of sensory activity that resemble those evoked by veridical sensory input. For example, blurry images that are perceived as faces elicit comparable levels of fMRI activity in face-sensitive regions of the fusiform gyrus (Summerfield et al., 2006). This implies that false alarms are not just guesses, but occur when decision regions read out sensory activity of a comparable amplitude to that occurring during veridical perception.

In the study reported here, monkeys reported the presence or absence of vibrotactile stimuli that could occur at any point in a decision period lasting from 1.5 to 3.5 s after a key down (KD) event initiating each trial. In a previous analysis of this data set, Carnevale and colleagues showed that false alarm trials are associated with two unique neural signatures in macaque premotor cortex, each with a

distinct temporal profile. First, firing rates on false alarm trials begin to grow at 1.5 s after KD and are sustained at higher levels than on other trial types throughout the decision period and into a subsequent delay. Second, cross-correlations between spike counts in simultaneously recorded neurons wax and wane across the decision period, peaking roughly in the middle (~2.5 s after KD). In this earlier paper, the authors use a computational model to show that detection judgments can be understood as depending on both amplitude of the applied stimulus and on a slowly fluctuating signal that provides correlated input to the neuronal population (Carnevale et al., 2012).

Here, the authors go a step further, introducing an innovative new analysis technique that capitalizes on the heterogeneity of single neuron responses in the premotor cortex. The responses of single neurons within decision-related brain regions are often very variable, and researchers are often obliged to carefully sub-select neurons of a given class or to attempt to cluster neurons by their response profiles. Instead, Carnevale and colleagues leverage this variability, averaging single neuron responses on hit trials to form a “template” that they then used to detect signal-like events on false alarm trials. Using the correct rejection trials (on which no stimulus was applied and none was reported) as a control, the authors show that “FA events,” (i.e., cell-specific neural signatures that matched this template) occurred most frequently when the monkey made false alarm. Moreover, these FA events tended to cluster within decision period, when the monkey was presumably expecting the stimulus to occur.

Next, using a multivariate analysis technique (Stokes et al., 2013) to estimate how

slowly varying patterns of reconstructed population activity encoded the intensity of stimulation, and the evolution of detection, they found that hits and false alarms shared similar trajectories, whereas misses and correct rejections showed limited encoding of these variables. In other words, false alarms occurred not only because of elevated baseline activity, as previously reported, but because specific neural signatures (“ghosts” in the decision process) resembling those evoked by veridical sensory stimulation were present in the premotor signal.

Together, these findings reveal that during the detection task, the monkeys are not passive recipients of somatosensory information. Rather, they capitalize on temporal expectations—learned information about the likely timing of sensory inputs—to maximize the probability of detecting the stimulus at each point in time (Nobre et al., 2007). One likely interpretation of the slowly fluctuating input signal that drives the premotor neurons in common is thus that it reflects temporal expectations about when the stimulus will arrive. Previous studies in the visual domain have suggested that widespread, slowly fluctuating signals evolve to match the timing of stimulation, perhaps acting as a time-varying gain control signal in perceptual decision making (Lakatos et al., 2008; Wyart et al., 2012). Indeed, the observation that the network evolves through characteristic states is consistent with the authors’ view that neural responses are partly driven by a slow, time-vary latent state that can encode information about the likely timing of events (Goel and Buonomano, 2014).

Nevertheless, there are some features of the data described by Carnevale and colleagues that are curious and merit further investigation. When onset times are variable, the function that tracks the probability of stimulation over time is

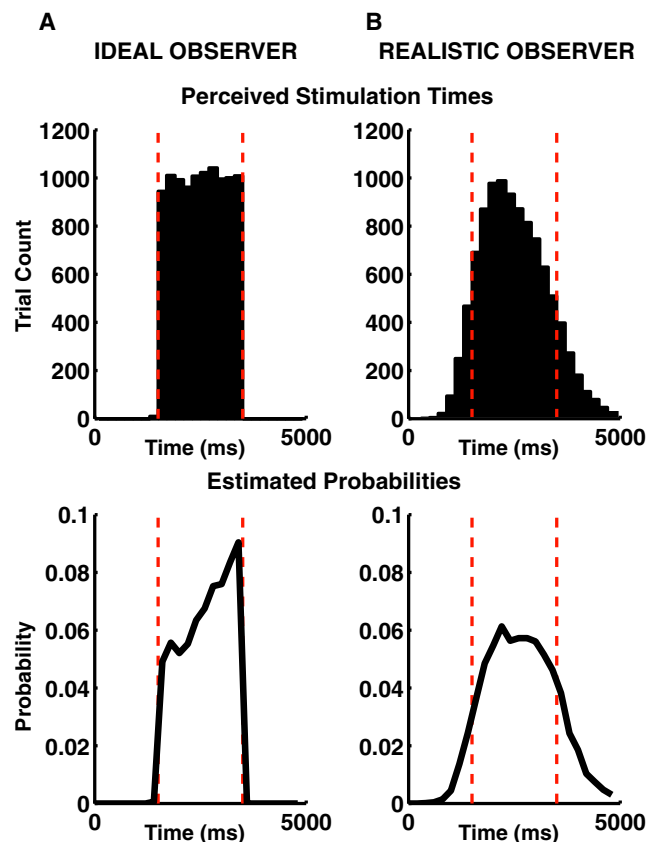


Figure 1. Simulation of the Estimated Stimulation Distributions and Estimated Hazard Functions

(A and B) Simulation of the estimated stimulation distributions (top row) and estimated hazard functions (bottom row) for an ideal observer (A) and a realistic observer with temporal estimation that becomes noisier with time (B). For an ideal observer, the hazard function increases as time elapses. For the realistic observer, the noise in the estimation process smooths the hazard function, and the Weber compression shifts the peak earlier in time, similar to the temporal profile of FA events reported by Carnevale et al. (2015).

known as the hazard function. Where the stimulus probability is uniform (as was the case here), the hazard function grows steadily across the decision period. Consider, for example, an ideal observer detecting a noiseless stimulus that could occur at any point across a 2-s interval, which for convenience we divide into ten time bins of 200 ms each (Figure 1A, upper panel). Assuming that the stimulus is present on half of all trials, the probability of the stimulus occurring in the first time bin is exactly 1 in 20, or 5%. However, given that no stimulus has been detected in the first nine time bins, the probability that it will occur in the final bin is 1 in 11, or about 9.1%. The hazard function thus grows over time, in a fashion similar to that shown in Figure 1A (lower panel). In

a visual detection task, the buildup activity of neurons in the lateral intraparietal area (LIP) has been shown to track a “smooth” version of the hazard function, as if the monkeys were basing their choices on a noisy time-varying estimate of stimulus probability (Janssen and Shadlen, 2005). Indeed, in the authors’ previous analyses of these data, the overall firing rate on FA trials increases steadily after KD and is sustained into the subsequent delay period, apparently tracking the optimal hazard function (Carnevale et al., 2012).

In the new study, the probability of FA events is maximal within the expected decision period, showing that the monkeys exhibit temporal expectations. However, its peak occurs early, and falls off as the decision period progresses—the opposite of what might be expected given the hazard function for stimulation. One possibility is that unlike in previous studies where stimulation was guaranteed to occur on every trial, here the hazard function rises only modestly across the decision epoch. Noise in either sensory processing or temporal estimation would flatten and smooth

the distribution of perceived onsets yet further (Figure 1B, upper panel), as the monkeys use prior information about the central tendency of stimulus onset to guide their temporal expectations (Jazayeri and Shadlen, 2010). Finally, uncertainty in the estimation of time interval increases proportionally to its duration (i.e., Weber’s law) (Gibbon et al., 1997; Janssen and Shadlen, 2005). This induces a compression of time perception that shifts temporal expectations forward, toward the start of the decision period (Figure 1B, lower panel). Indeed, that the monkeys seem to expect stimuli early in the decision period is also suggested by behavioral data; the hit rate steadily declines as a function of stimulus onset, both for high- and low-amplitude stimulation trials.

The authors argue that their results show a dynamic modulation of response criterion, expressed as a boundary in state space (or “separatrix”) that segregates the “yes” from the “no” choices. Using a dynamic, recurrent neural network simulation, they were able to recreate key features of their data, including the early peak of FA events and the state space trajectories of the network on each trial type. One interesting next step for Carnevale and colleagues could be to try to relate the three neural phenomena they have now described that accompany somatosensory false alarms in premotor cortex: sustained increases in overall firing rate, a late peak in spike correlations, and an early peak in FA events. They will also need to relate their findings to elegant psychophysical work that suggested that temporal expectations might do more than modulate the proximity of the network state to a choice boundary—they might actively increase signal-to-noise ratios,

resulting in more sensitive detection of sensory events (Rohenkohl et al., 2012). Precisely how temporal expectations modulate perceptual decisions remains an open question, but Carnavale and colleagues offer new insights and new tools for unravelling this interesting question and for deciphering the neural signatures of false detection in the primate brain.

REFERENCES

Carnevale, F., de Lafuente, V., Romo, R., and Parga, N. (2012). *Proc. Natl. Acad. Sci. USA* *109*, 18938–18943.

Carnevale, F., de Lafuente, V., Romo, R., Barak, O., and Parga, N. (2015). *Neuron* *86*, this issue, 1067–1077.

de Lafuente, V., and Romo, R. (2005). *Nat. Neurosci.* *8*, 1698–1703.

de Lafuente, V., and Romo, R. (2006). *Proc. Natl. Acad. Sci. USA* *103*, 14266–14271.

Gibbon, J., Malapani, C., Dale, C.L., and Gallistel, C. (1997). *Curr. Opin. Neurobiol.* *7*, 170–184.

Goel, A., and Buonomano, D.V. (2014). *Philos. Trans. R. Soc. Lond. B Biol. Sci.* *369*, 20120460.

Janssen, P., and Shadlen, M.N. (2005). *Nat. Neurosci.* *8*, 234–241.

Jazayeri, M., and Shadlen, M.N. (2010). *Nat. Neurosci.* *13*, 1020–1026.

Lakatos, P., Karmos, G., Mehta, A.D., Ulbert, I., and Schroeder, C.E. (2008). *Science* *320*, 110–113.

Nobre, A., Correa, A., and Coull, J. (2007). *Curr. Opin. Neurobiol.* *17*, 465–470.

Rohenkohl, G., Cravo, A.M., Wyart, V., and Nobre, A.C. (2012). *J. Neurosci.* *32*, 8424–8428.

Stokes, M.G., Kusunoki, M., Sigala, N., Nili, H., Gaffan, D., and Duncan, J. (2013). *Neuron* *78*, 364–375.

Summerfield, C., and de Lange, F.P. (2014). *Nat. Rev. Neurosci.* *15*, 745–756.

Summerfield, C., Egner, T., Mangels, J., and Hirsch, J. (2006). *Cereb. Cortex* *16*, 500–508.

Wyart, V., de Gardelle, V., Scholl, J., and Summerfield, C. (2012). *Neuron* *76*, 847–858.