PHYS.ORG

Of noise and neurons: Sensory coding, representation and short-term memory

October 31 2012, by Stuart Mason Dambrot

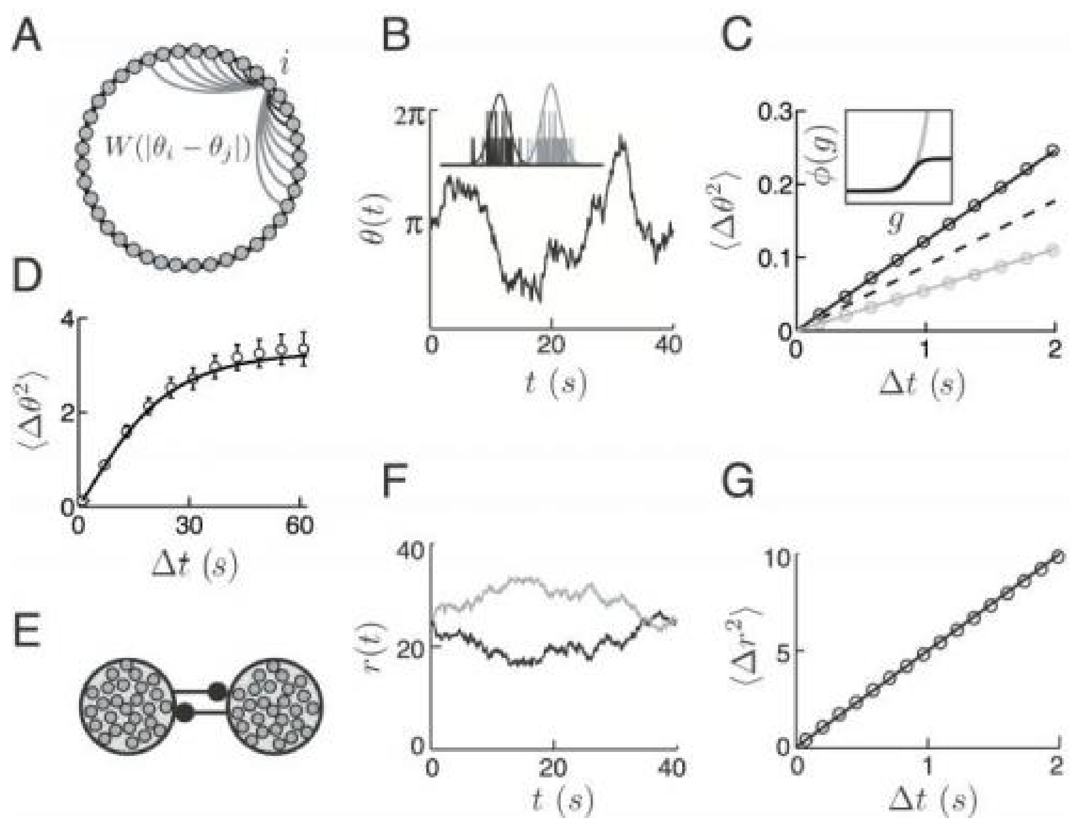

Numerical validation of the diffusivity result in different attractor networks. (A) Ring network: Schematic of a memory network representing a periodic variable such as orientation. (B) Simulation of a 1,024 neuron network, with exponential neural transfer function. Inset: snapshots of population activity at two times separated by 10 s. (C, D) Mean squared displacement of the attractor state over elapsed time intervals Δt]. Circles: numerical simulations with exponential (gray) or sigmoidal (black) neural transfer function (see Inset). Solid lines: theoretical prediction of Eq. 2. Dashed lines: bound on D from the information-diffusion inequality for sigmoidal nonlinearity. (E) Mutual-inhibition network. Schematic of a memory network representing a nonperiodic 1-d variable. Two neurons or two populations, inhibit each other with equal weights. The difference of the two firing rates represents the stored variable. (F) Random drift of the firing rates (in 1/ms) in the mutual-inhibition network. (G) Mean squared displacement of the attractor state over an elapsed time Δt. Circles: simulation of mutual-inhibition network, with linear transfer function. Solid line: theoretical prediction of Eq. 2. Credit: PNAS.

(Phys.org)—While much is known about the limiting effect of neural noise on the fidelity of sensory coding representation, knowledge about the impact of noise in short-term memory and integrator networks has remained more elusive. (Integrator networks are networks of nodes – in this case neurons in a biological network – often recurrently connected, whose time dynamics settle to stable stationary, cyclic, or chaotic patterns, that can integrate or store memories of external inputs.)

Recently, however, scientists at The Hebrew University of Jerusalem, Harvard University and University of Texas, Austin used statistical and dynamical approaches to investigate how neural noise interacts with neural and network parameters to limit memory. They derived a series of unanticipated results – including the implications that short-term memory may be co-localized with sensory representation – by establishing a fundamental limit on the network's ability to maintain a persistent neural state.

Assistant Professors Ila R. Fiete and Yoram Burak described the various challenges they encountered. "The dynamics of spiking neural networks are in general highly nonlinear and involve a very large number of degrees of freedom," Fiete tells *Phys.org*, addressing their analysis of how stored memory in continuous attractor networks will stochastically degrade over time. She adds that most work on such networks is focused on deterministic dynamics. "*A priori*," she continues, "it wasn't obvious that one could evaluate precisely how noise affects the state of the system," pointing out that investigations into noise affecting a memory state in such networks was previously done for very simple systems with linear neurons (those with no nonlinearity in the neural response), with noise being externally injected and having simple statistical properties.

"By contrast," Burak explains, "we wanted to understand the role of noise that originates within the network – that is, noise intrinsic to single neurons or synapses, as opposed to simple external noise." Intrinsic neuronal noise has a more complicated form, and its properties vary in each neuron, based on the neuron's firing rate at that moment in time. "Unlike external noise, which is assumed to directly affect the memory state, internal noise must be passed through the nonlinear dynamics of the system, to derive its effects on the memory state. We wanted to obtain a general theory, without making particular assumptions about network connectivity and neural nonlinearity."

Their work began with an intuitive idea they had before doing calculations about continuous attractor networks, says Fiete. "The network's limited ability to read its own past state from its spikes, so to speak, must limit its ability to ***maintain*** that past state into the future. This must limit the accuracy of persistent activity. The biggest challenge here was to translate this intuitive idea into a rigorous formal statement about a concrete model of spiking neurons. The formal statement is given by a combination of a statistical limit with a dynamical property, in the form of an *information-diffusion inequality*."

Among the study's unexpected consequences, Fiete continues, was that despite the long persistence time of short-term memory networks, it does not pay to accumulate spikes for much longer than the short time-constant of individual neurons, to read out the contents of the network. "This result was born out of our attempt to understand the consequences of the gradual loss of accuracy in storing a variable in a memory network due to diffusive dynamics. Our initial intuition was that in a network with persistent memory, the longer one observes the spikes generated by the network, the better one should be able to infer its state to arbitrary precision. However, while one is collecting spikes to improve statistical precision in estimating the network's state, the state itself drifts due to diffusion. As a result, the state of the network cannot be inferred to arbitrary precision." What surprised the researchers was that there was actually no benefit to collecting spikes for any appreciable length of time beyond a very short time scale – that is, the intrinsic time constant of single neurons.

"This is actually quite a satisfying result," Burak points out, "because it means that one does not need a separate memory network that collects and remembers spikes over a long time to read out what is encoded in the memory network. Moreover, while the derivation is quite straightforward, in order to see that the relevant time scale is short, it is necessary to use the information-diffusion inequality that was derived in our study."

Another surprise was that for certain neural transfer functions, the conditions for optimal sensory coding coincide with those for optimal storage. This suggests that short-term memory may be co-localized with sensory representation. "This is a direct outcome of the information-diffusion inequality and the observation that in some networks, the inequality is saturated," Burak explains. "One of the goals of our work was to understand how the structure of the network affects its ability to maintain a persistent state. The relationship with the internal Fisher Information allows us to immediately use previously derived results, on optimality of tuning curves for sensory coding." (*Fisher Information* is a particular Riemannian metric, definable on a smooth statistical manifold, that quantifies how much uncertainty remains about the parameters of a probabilistic process after a finite number of observations.)

An interesting outcome of this result – i.e., that the same conditions (for example, tuning curve shapes in the neural responses) are needed for accurate sensory representation and for maximally persistent memory – is that the same network may be performing both sensing and memory functions.

Fiete and Burak devised a number of innovations to derive their results. "The first innovation in our work is that we were able to calculate *exactly* the stochastic dynamics of the attractor state under very general conditions – specifically, arbitrary weights and neural transfer function." Fiete recounts. "The key insight of the work is that the degradation of the stored memory, a dynamic property of the network, is intimately related to the noisiness in network spikes as seen by an external observer that considers these spikes to be encoding a variable." In fact, the latter point of view – looking at spikes generated in a certain brain area and asking how much information they carry about a stimulus or a stored memory – is very common in computational neuroscience. The researchers showed that the ambiguity in the encoded variable, as seen by an external observer, is linked to the network's own ability to maintain a persistent state through an *information-diffusion inequality*.

The scientists have also identified other innovations they might develop and apply to the current experimental design, as well as the planned next steps in their research. "We derived our results for a fairly simple model of spiking neurons, so it will be very interesting to see whether our key insights apply also to more realistic neural models," Fiete continues. "We're also interested in applying these ideas to concrete brain areas where the underlying dynamics are believed to be governed by a continuous attractor, such as the occulomotor system that maintains a stable position of the eye." More broadly, they're interested in encoding continuous variables in noisy, spiking neural networks. "We also want to see if our results – in particular the information-diffusion inequality – applies more generally to any dynamical system with attractors, continuous or not, and related to neuroscience or not."

In addition, adds Burak, other areas of research might benefit from their findings. "There's a long history in physics of relating concepts from thermodynamics, such as noise and fluctuations, to memory," In biology this is an important issue, since biological systems require memory for their function – and they're often noisy:—the limited resources in a biological system often dictate a certain amount of stochasticity. "Therefore," Burak concludes, "it's very likely that our ideas can be applied to memory storage outside the context of neural systems – for example, in gene regulatory networks within a living cell. In fact, the type of noise occurring in these systems is quite similar from a mathematical point of view to the Poisson noise assumed in our model neural networks."

More information: Fundamental limits on persistent activity in networks of noisy neurons, *PNAS* October 23, 2012, vol. 109 no. 43 17645-17650, doi:10.1073/pnas.1117386109