

leaving open the possibility that granule cells sculpt M/T responses to specific odors (Fig. 1). In addition, higher frequency olfactory bulb oscillations, which reflect underlying network activity and are linked to olfactory discrimination learning, have been localized to the EPL¹⁵. To what extent these oscillations, and the activity of the underlying cell assemblies, are coordinated by granule cell activity versus that of other interneuron networks remains to be addressed. Perhaps the tighter control of behavioral state might reveal unexpected results in future studies, such as stronger respiratory phase coupling during particular phases of behavioral tasks, that may reveal more about how granule cells influence M/T cell output.

By showing that granule cells are strongly influenced by waking state, the authors have confirmed this network of inhibitory interneurons as a viable source of dynamic information processing in the bulb. The door is finally opened to empirically addressing granule cell function *in vivo*.

COMPETING FINANCIAL INTERESTS

The authors declare no competing financial interests.

- Restrepo, D., Doucette, W., Whitesell, J.D., McTavish, T.S. & Salcedo, E. *Trends Neurosci.* **32**, 525–531 (2009).
- Markopoulos, F., Rokni, D., Gire, D.H. & Murthy, V.N. *Neuron* **76**, 1175–1188 (2012).
- Boyd, A.M., Sturgill, J.F., Poo, C. & Isaacson, J.S. *Neuron* **76**, 1161–1174 (2012).
- Shipley, M. T. & Ennis, M. *J. Neurobiol.* **30**, 123–176 (1996).

- Koulakov, A.A. & Rinberg, D. *Neuron* **72**, 124–136 (2011).
- Devore, S. & Linster, C. *Front. Behav. Neurosci.* **6**, 52 (2012).
- Cazakoff, B.N., Lau, B.Y.B., Crump, K.L., Demmer, H.S. & Shea, S.D. *Nat. Neurosci.* **17**, 569–576 (2014).
- Shepherd, G.M. *The Synaptic Organization of the Brain* (Oxford University Press, USA, 2004).
- Shusterman, R., Smear, M.C., Koulakov, A.A. & Rinberg, D. *Nat. Neurosci.* **14**, 1039–1044 (2011).
- Cury, K.M. & Uchida, N. *Neuron* **68**, 570–585 (2010).
- Kapoor, V. & Urban, N.N. *J. Neurosci.* **26**, 11709–11719 (2006).
- Abraham, N.M. *et al. Neuron* **65**, 399–411 (2010).
- Kato, H.K., Gillet, S.N., Peters, A.J., Isaacson, J.S. & Komiyama, T. *Neuron* **80**, 1218–1231 (2013).
- Miyamichi, K. *et al. Neuron* **80**, 1232–1245 (2013).
- Kay, L.M. *et al. Trends Neurosci.* **32**, 207–214 (2009).

Useful dynamic regimes emerge in recurrent networks

Vishwa Goudar & Dean V Buonomano

The internal dynamics of recurrent cortical circuits is crucial to brain function. We now learn that simply increasing the strengths of recurrent connections shifts neural dynamics to a potentially powerful computational regime.

As a pianist plays, spatiotemporal patterns of action potentials in the brain are transformed into changing patterns in the musculoskeletal system. The resulting sounds in turn produce spatiotemporal patterns of neural activity in the auditory cortex of the listener. Similarly, a sequence of words on a page is the product of spatiotemporal patterns of neural activity in a writer's brain, and those words in turn generate patterns of activity in the brains of readers. Ultimately, time-varying patterns of neural activity underlie just about everything we do. It is generally accepted that these spatiotemporal patterns of activity arise in part from the internal dynamics of recurrent cortical circuits. For this reason, considerable efforts have been devoted to understanding how patterns of activity emerge from recurrent neural networks. Although this work has generated many insights, it has also proved humbling. But in this issue of *Nature Neuroscience*, Ostojic¹ takes an important step—in what will undoubtedly be a long walk—toward better understanding neural

dynamics in computational models of recurrent neural networks. His work demonstrates that increasing the synaptic strengths between recurrently connected units induces a transition from a regime with little computational potential to one with high potential.

Ostojic¹ based his simulations on a previous model of recurrent neural networks² composed of simple spiking units called integrate-and-fire neurons. A typical simulation is composed of 10,000 units, 80% excitatory and 20% inhibitory. All units are randomly connected with a connection probability of 0.1, reflecting the experimental observation that the connection probability between nearby cortical pyramidal neurons is 0.1 to 0.2. Each unit also receives a large tonic input, resulting in spontaneous firing in the absence of any recurrent connections.

One of the most important parameters in this class of models is the connection strength between the units. This value, often denoted by J , determines the strength of both the excitatory and inhibitory weights in the networks (we are making the simplifying assumption that all weights are the same). But the inhibitory weights are further governed by a factor of g . Thus, the total input to a neuron would be proportional to $N_{\text{Ex}}J - N_{\text{Inh}}gJ$, where J represents the strength of an excitatory synapse, gJ represents the strength of an inhibitory synapse, and N_{Ex} and N_{Inh} represent the

number of excitatory and inhibitory synapses, respectively. When $g = 4$, the net input will on average be 0 because there are four times more excitatory than inhibitory neurons; this would represent a perfectly balanced case. Ostojic focused on $g = 5$, meaning that the recurrent connections are dominated by inhibition, but the network is nevertheless said to be 'balanced' because the ratio of excitation and inhibition is constant across a range of activity levels.

Previous studies have demonstrated that balanced excitatory-inhibitory networks exhibit a dynamic regime in which neurons fire irregularly in a manner that resembles the activity of cortical neurons during stationary conditions^{2–6}. This type of pattern is referred to as an irregular asynchronous, or a homogeneous asynchronous, regime (Fig. 1a). The irregular activity arises because excitation and inhibition mostly cancel each other out, but spikes are generated from time to time because of voltage fluctuations.

Ostojic¹ used simulations and an analytical approach to extend these previous results. He found that, as synaptic strength parameter J is increased, there is a transition from a homogeneous to a heterogeneous regime (Fig. 1b). In this state, the firing rates vary substantially in time because the units become bursty. This burstiness resembles experimentally observed responses while animals

Vishwa Goudar and Dean V. Buonomano are in the Departments of Neurobiology and Psychology, Integrative Center for Learning and Memory, Brain Research Institute, University of California, Los Angeles, Los Angeles, California, USA.
e-mail: dbuono@ucla.edu




Figure 1 Regimes of neural dynamics in a spiking recurrent neural network. (a) Top, the pattern of activity in a network with moderately strong connections generates a homogeneous regime in which the units spike irregularly, maintain more-or-less constant firing rates and do not burst. Each row represents the spiking activity of a single neuron. In this regime, the overall pattern of activity is fairly homogeneous. For example, the patterns at 0.5 and 0.8 s (gray bars) have a high correlation (0.7). Bottom, the correlation matrix shows the correlations between all possible time-bin pairs (20-ms Gaussian convolution). The intersection of the columns and rows at 0.5 and 0.8 s reflects the correlation between the two highlighted patterns (*). (b) Top, when the coupling between the neurons is strong, the spatiotemporal pattern of action potentials shifts to a heterogeneous regime in which the firing rates of units can vary substantially over time and the spiking is bursty. In this regime, the correlation between patterns at any two points is smaller. For example, the correlation at 0.5 and 0.8 s is 0.3. Bottom, the correlation matrix shows that the mean correlation is much lower over all possible time bins, reflecting an increase in the ability of the network to encode information.

perform tasks and is, in one respect, surprising, as it is an emergent property of the network as opposed to an intrinsic property of the units. The transition from a homogeneous to a heterogeneous regime resonates well with the transition described in the seminal work of Sompolinsky and colleagues in firing-rate networks⁷ (simulations in which the units do not spike).

The use of relatively simple integrate-and-fire units along with a few simplifying assumptions allowed Ostojic¹ to mathematically demonstrate that this heterogeneous regime is a function of the inherent instabilities of the system: small fluctuations lead to differences in the firing rate of individual neurons, which in turn generate further fluctuations. In the language of dynamical systems, in the heterogeneous regime, the network has eigenvalues larger than 1, meaning that it will not converge to a stable state. On the intuitive level, it is more difficult to understand the nature of the transition, but one possibility is that the stronger coupling generates longer-lasting inhibitory postsynaptic potentials. This longer-lasting

inhibition, particularly of the inhibitory units, could provide a longer window of opportunity for cells to continue to spike as a result of their tonic activity.

Why is the transition from a homogeneous to a heterogeneous regime of interest? Both are similar to experimentally observed regimes, but they have very different computational properties. The homogeneous regime is well suited to transmitting information about the firing rate of external inputs: given a time-varying input, neurons in the network will simply increase or decrease their mean firing rate in relation to the input. But, computationally speaking, this state is boring because the network itself is primarily functioning as a relay. In contrast, the complex dynamics of heterogeneous regimes are more interesting. This can be intuitively understood in terms of a simple computation: encoding time. The notion that the brain might encode time in the changing patterns of neural activity was first proposed in the context of the Marr-Albus-Mauk model of the cerebellum⁸ and

is one of the main theories of how cortical networks may encode time and perform complex time-varying computations^{9–11}. This framework, however, requires that the pattern of activity (the neural trajectory) of a circuit vary appreciably over time. Look at the patterns of activity in a homogeneous regime (Fig. 1a) and a heterogeneous regime (Fig. 1b) and imagine that you needed to use a pattern of a given bin to mark that point in time. For this to work, the population vector at each time bin must be different from that at all other time bins. Thus, the more distinct the patterns, the easier it will be to use the evolving pattern of activity to tell time. The correlation matrices indeed demonstrate that the patterns are much more distinct and information-rich in the heterogeneous versus homogeneous regime.

Ostojic¹ highlights the computational potential of the heterogeneous regime by feeding two different time-varying inputs into a network in the homogeneous regime and demonstrating that the resulting spatiotemporal pattern of activity is low dimensional, meaning that the network cannot distinguish between the inputs. In contrast, in the heterogeneous regime, the same inputs generate a higher dimensional representation, allowing the network to discriminate the inputs or to hold a memory of past events. This behavior is highly desirable in frameworks in which computations arise from the interaction of external inputs and the inherently time-varying internal state of the circuits¹⁰.

It must be stressed that, at this point, the computational potential of the network is hypothetical, as the network is chaotic. That is, if the simulation shown in Figure 1b were run again in the presence of a bit of noise, the pattern of activity would be entirely different, both at the level of the spike times and firing rates of the units. This voids the ability of the network to encode time or maintain a memory of past events. Indeed, it is well established that spiking recurrent networks are generally chaotic^{6,12}. Thus, Ostojic's results¹ further highlight the need to understand how to control the dynamics of spiking recurrent neural networks. One potential way to do this is to incorporate plasticity into the recurrent synapses. Indeed, it has recently been demonstrated, using simpler firing-rate models, that supervised learning rules can be used to tune the recurrent weights so as to get the best of both worlds: complex heterogeneous patterns without the chaos¹³. However, the learning rule used in that study is not biologically plausible and is not directly applicable to spiking networks. Another potential approach takes advantage

of a study demonstrating that incorporating clustering into the connectivity can also introduce spontaneous transitions between different firing regimes¹⁴.

As with all computational models the network studied by Ostojic¹ incorporates potentially important simplifications, including the fact that the synapses sum linearly (there is no saturating driving force). But perhaps one of the most important assumptions of the model is that the units are spontaneously active. Thus, in a sense, the burstiness is not an emergent property of the network because each unit in isolation is essentially 'tonically bursting'. Rather, it may be the pattern of inhibition-induced pauses in tonic activity that creates the heterogeneous state. In addition,

strictly speaking, the patterns do not arise autonomously from the internal dynamics of the network as with actual cortical networks, in which self-perpetuating patterns of activity can be observed even *in vitro*¹⁵. Thus, another question for future studies will be to determine whether Ostojic's results¹ hold for excitation-dominated regimes and whether the network can generate heterogeneous patterns of activity that are driven entirely by the internal dynamics.

COMPETING FINANCIAL INTERESTS

The authors declare no competing financial interests.

1. Ostojic, S. *Nat. Neurosci.* **17**, 594–600 (2014).
2. Brunel, N. *J. Comput. Neurosci.* **8**, 183–208 (2000).
3. Troyer, T.W. & Miller, K.D. *Neural Comput.* **9**, 971–983 (1997).

4. Renart, A. *et al. Science* **327**, 587–590 (2010).
5. Shadlen, M.N. & Newsome, W.T. *Curr. Opin. Neurobiol.* **4**, 569–579 (1994).
6. van Vreeswijk, C. & Sompolinsky, H. *Science* **274**, 1724–1726 (1996).
7. Sompolinsky, H., Crisanti, A. & Sommers, H.J. *Phys. Rev. Lett.* **61**, 259–262 (1988).
8. Mauk, M.D. & Donegan, N.H. *Learn. Mem.* **4**, 130–158 (1997).
9. Mauk, M.D. & Buonomano, D.V. *Annu. Rev. Neurosci.* **27**, 307–340 (2004).
10. Buonomano, D.V. & Maass, W. *Nat. Rev. Neurosci.* **10**, 113–125 (2009).
11. Rabinovich, M., Huerta, R. & Laurent, G. *Science* **321**, 48–50 (2008).
12. Wallace, E., Maei, H.R. & Latham, P.E. *Neural Comput.* **25**, 1408–1439 (2013).
13. Laje, R. & Buonomano, D.V. *Nat. Neurosci.* **16**, 925–933 (2013).
14. Litwin-Kumar, A. & Doiron, B. *Nat. Neurosci.* **15**, 1498–1505 (2012).
15. Kruskal, P.B., Li, L. & MacLean, J.N. *Nat. Commun.* **4**, <http://dx.doi.org/10.1038/ncomms3574> (2013).

Go means green

Joseph T McGuire & Joseph W Kable

A simple cued-approach training procedure can bias economic choices toward specific goods. It appears to work by drawing overt attention toward trained items, scaling up their judged value.

A bottle of Napa Valley Cabernet is probably tasty, but is it worth \$70 to me? Is a fancy cup of coffee worth \$7? As a would-be purchaser, I need to translate my subjective and intuitive valuation of a good into units of cold, hard cash. Even for nonmonetary choices, such as whether to read a book or watch TV, I need to be able to assess and compare the values of dissimilar options on a common scale.

Schonberg *et al.*¹ describe a simple manipulation that boosts the value people place on individual goods. The goods in their experiments were snack items (for example, a Butterfinger bar) and the manipulation consisted of 'Go' training. Pictures of 60 different foods appeared one by one on a computer screen, 12 times each over the course of 48 min. Most items just had to be viewed passively, but about a quarter were designated by the experimenters as 'Go items'. Every appearance of a Go item was quickly followed by an auditory signal to make a speeded key press.

In the next stage of the study, participants made binary decisions, picking which of two snack items they would rather receive at the end of the experiment. In pairs consisting of a Go item and a control item, matched for their pretraining value to the participant, Go items were selected 60–65% of the time. Similar

effects were apparent in a subsequent auction task, in which participants bid an average of about 12 cents more for Go items than for originally equivalent control items.

These are surprising results, as the Go-training manipulation differed markedly from other common strategies for modifying people's preferences. Items were not associated with any new incentive, participants received no additional information, and the Go and control items did not differ in their familiarity or duration of exposure². The manipulation does not appear to target habitual responding^{3,4}, nor does it alter the framing of decisions⁵ or the architecture of the choice environment⁶. Given the unexpectedness of the observed effects, it is important and commendable that the researchers present multiple replications⁷, documenting the influence of Go training on binary choice in a total of five independent samples.

A potentially revealing wrinkle is that not all snack items were equally susceptible to the effects of Go training. There was a 'rich-get-richer' effect across items, such that an individual's favorite items received the biggest boost. The strongest effects were seen when comparing initially high-valued items assigned to the Go versus control condition. For initially less-favored items, the differential effect of Go training was small to nonexistent.

The dependence on initial item value helps to argue against the uninteresting possibility that participants might, for some reason, have

selected the Go items on purpose instead of expressing their true preferences. It also helps shed light on how Go training works. One way to make sense of this pattern of effects is to suppose that, instead of adding a fixed increment to an item's subjective value, Go training had an effect more akin to scaling item value by a multiplicative factor (**Fig. 1**). A differential scaling or amplification of subjective value would have the biggest effect if both items were highly valued to begin with.

Amplification-like influences on subjective valuation have been discussed before in characterizing the effects of overt visual attention. A recent theoretical proposal by Krajbich *et al.*⁸ holds that choice alternatives loom larger in moment-by-moment value comparisons if they are being visually fixated. Schonberg *et al.*¹ present evidence for a direct connection between Go training and visual fixation. Eye-tracking data revealed that participants spent more time fixating Go items than control items during binary choices, even controlling for the fact that Go items were more frequently selected. This raises the possibility that overt attentional capture may have mediated the effect of Go training on expressed preferences.

In situations in which preferences shift, it is potentially instructive to investigate whether there is a corresponding shift in preference-related brain activity. A natural target for such an investigation is ventromedial prefrontal cortex (VMPFC). VMPFC is one in a small

Joseph T. McGuire and Joseph W. Kable are in the Department of Psychology, University of Pennsylvania, Philadelphia, Pennsylvania, USA. e-mail: kable@psych.upenn.edu