

authors as to this new role is that neurotoxin-effected cleavage of VAMP-2, the synaptic vesicle SNARE, also reduces evoked release without lowering the calcium sensitivity of the release apparatus. Does this mean that syt2 and VAMP-2 combine to localize vesicles to calcium channels? Do they associate independently with calcium channels, or must they bind to the other SNAREs first? When syt2 or VAMP-2 is compromised, is the mispositioning due to loosened tethering of vesicles to active zones or to calcium channels drifting away from active zones? Is there an increase in the physical distance between calcium channels and primed vesicles or perhaps an increase in the endogenous calcium buffer concentration, which would limit long-range calcium diffusion? It would be of

great interest to see how these questions are addressed in the future.

**REFERENCES**

Bollmann, J.H., Sakmann, B., and Borst, J.G. (2000). *Science* 289, 953–957.

Geppert, M., Goda, Y., Hammer, R.E., Li, C., Rosahl, T.W., Stevens, C.F., and Sudhof, T.C. (1994). *Cell* 79, 717–727.

He, L., Wu, X.S., Mohan, R., and Wu, L.G. (2006). *Nature* 444, 102–105.

Hosoi, N., Holt, M., and Sakaba, T. (2009). *Neuron* 63, 216–229.

Meinrenken, C.J., Borst, J.G., and Sakmann, B. (2002). *J. Neurosci.* 22, 1648–1667.

Schneggenburger, R., and Neher, E. (2000). *Nature* 406, 889–893.

Sun, J.Y., and Wu, L.G. (2001). *Neuron* 30, 171–182.

Sun, J., Pang, Z.P., Qin, D., Fahim, A.T., Adachi, R., and Sudhof, T.C. (2007). *Nature* 450, 676–682.

Wadel, K., Neher, E., and Sakaba, T. (2007). *Neuron* 53, 563–575.

Wimmer, V.C., Nevian, T., and Kuner, T. (2004). *Pflugers Arch.* 449, 319–333.

Wu, L.G., and Borst, J.G.G. (1999). *Neuron* 23, 821–832.

Wu, X.S., McNeil, B.D., Xu, J., Fan, J., Xue, L., Melicoff, E., Adachi, R., Bai, L., and Wu, L.G. (2009). *Nat. Neurosci.* 12, 1003–1010.

Xu, J., Mashimo, T., and Sudhof, T.C. (2007). *Neuron* 54, 567–581.

Xue, M., Ma, C., Craig, T.K., Rosenmund, C., and Rizo, J. (2008). *Nat. Struct. Mol. Biol.* 15, 1160–1168.

Young, S., and Neher, E. (2009). *Neuron* 63, this issue, 482–496.

## Imagining the Possibilities: Ripples, Routes, and Reactivation

James J. Knierim<sup>1,\*</sup>

<sup>1</sup>Zanvyl Krieger Mind/Brain Institute and Department of Neuroscience, Johns Hopkins University, Baltimore, MD 21218, USA

\*Correspondence: [jknierim@jhu.edu](mailto:jknierim@jhu.edu)

DOI 10.1016/j.neuron.2009.08.002

**Hippocampal place cells fire selectively when a rat occupies a particular location. Under certain conditions, the cells briefly represent trajectories along locations away from the rat’s current location. New results lend important insight into this phenomenon and demonstrate spatiotemporally coherent, cognitive representations that are independent of current sensory input.**

Two key goals of cognitive neuroscience are to understand (1) how the brain constructs internal representations of the world “out there” and (2) how the animal uses these representations to guide adaptive behavior, store memories, form plans, evaluate potential outcomes, and decide on a proper course of action. For many years, place cells of the rat hippocampus have been a superb system for addressing these issues. Place cells are the building blocks of context-specific, map-like representations of the animal’s location in an environment. These cells are influenced by exteroceptive senses (such as vision) and interoceptive senses

(such as the vestibular system). However, the hippocampus is many synapses removed from the primary sensory areas. Accordingly, the activity of place cells is controlled at least as much by the internal dynamics of hippocampal networks as by the external stimuli that an animal experiences as it moves through an environment. Recent years have witnessed a dramatic increase in our understanding of these internal dynamics—how they can be divorced from explicit sensory input and how they may underlie the role of the hippocampus in memory consolidation and cognition. A new study by Davidson and colleagues (2009; this

issue of *Neuron*) provides important new insights into this field of investigation.

As good hippocampologists, we need first to put these findings into context. A major advance in understanding the intrinsic activity of the hippocampus was the discovery of “sleep reactivation” of place cells (Wilson and McNaughton, 1994). When a rat is in slow-wave sleep (or resting in a period of quiet wakefulness), the hippocampus enters a state characterized by large, irregular activity (LIA) in the hippocampal EEG. This state is punctuated by prominent events called sharp-wave/ripple (SW-R) complexes in the CA1 region, which are thought to be

generated by bursts of activity in the CA3 network that drive population bursts of CA1 pyramidal cells (Buzsaki, 1986). Wilson and McNaughton (1994) showed that, when an animal is in the LIA state after exploration of an environment, the patterns of neural activity in these SW-R events are similar to the patterns of activity that occurred during the recent exploratory behavior itself. Thus, the CA1 region reactivates the spatial representations that were active during behavior. Subsequent studies demonstrated a similar reactivation phenomenon in REM sleep (Louie and Wilson, 2001), in which the LIA activity in the EEG disappears and is replaced by the rhythmic, slow activity of the theta rhythm (the state that also dominates during exploratory behavior). Initial speculation centered on the notion that this reactivation was a mechanism of offline memory consolidation, by which the hippocampus “teaches” the neocortex the memories that it stored during the actual experience.

These reactivation events, and other similar phenomena, occur not just during sleep or restfulness—they occur when the animal is active and alert. SW-R events occur during behavior (O’Neill et al., 2006; Jackson et al., 2006), sometimes even when theta rhythm is present, and neural firing patterns during these exploratory SW-R events are similar to those during sleep or rest. Moreover, the frequency of occurrence of exploratory SW-R events increases with experience, suggesting that they are a result of synaptic strengthening during exploration. Whereas the replay during sleep and quiet restfulness retains the same temporal order that occurred during the behavioral experience (Skaggs and McNaughton, 1996; Lee and Wilson, 2002), “reverse replay” can occur when rats pause at the end of a linear track and consume reward (Foster and Wilson, 2006). That is, if the rat runs through points A, B, and C to reach the reward site (activating place cells a, b, and c), the replay during the subsequent SW-R event at the reward site consists of the place cells firing in the reverse order (c-b-a). This result was interpreted as a potential mechanism for associating the reward location with representations of the trajectories that started elsewhere but eventually led to the reward (the so-called credit assignment problem). In

contrast to the reverse replay that occurs at the end of a run, forward replay occurs just prior to the start of the next run (Diba and Buzsaki, 2007). In all cases of alert SW-R reactivation, the hippocampus can momentarily activate representations of locations that do not correspond to the physical location of the animal.

The study by Davidson and colleagues (2009) provides two major advances. The first advance is a demonstration that the replay events can extend over much larger spatiotemporal sequences than previously thought. In prior studies, the experiments were performed on the small tracks and mazes that are typical of most place-cell recordings. Consequently, only brief sequences could be represented (in a time-compressed fashion) during single SW-R events. It was unclear from these studies whether the hippocampus was capable of reactivating representations of longer trajectories or sequences of events. An answer to this question would have important consequences for theories of the role of this reactivation in the consolidation of episodic memories or in the linking of brief sequences to create unified representations of an extended region (or experience). Davidson et al. show that, when an animal pauses during exploration of a large, extended track (~10 m), the hippocampus can chain together multiple SW-R events to create a coordinated reactivation of a much larger trajectory (many meters) than previously demonstrated. Each SW-R is associated with a smaller subset of the trajectory (as in previous studies), but these representations are reactivated in the proper sequence during a chain of SW-Rs such that the extended trajectory is replayed almost continuously. The extended replays can either be in the forward or reverse direction.

The second advance was the demonstration that the extended replays can be initiated at locations distant from the current location of the animal. That is, the animal may be at one location on the track when it pauses, but the reactivation event may start at a representation of another location on the track and then proceed either toward or away from the animal. Thus, the reactivation is not necessarily initiated by a set of sensory cues that the animal experiences at its current location. This finding reinforces

the importance of intrinsic network dynamics, unrelated to current sensory input, in driving hippocampal activity. It is hard to resist the speculation that these reactivation events are a correlate of the rat’s “thinking” (to the extent that rats can think) about other parts of the track and about its recent experiences at locations other than its current position.

These results bear a strong resemblance to a phenomenon reported by Johnson and Redish (2007) during a multiple T-maze task. As the animal learned a series of left-right turns in order to navigate to the reward location, the animals would often pause at the choice points, as if it was considering its choices and deciding which way to turn. This behavior has been termed vicarious trial and error (VTE). During these pauses, the hippocampal cells fired in patterns that reflected trajectories down the left or right choice arms (in irregularly alternating sweeps). These sweeps may be a neural correlate of the VTE behavior. The animal’s choice was not predictable from these sweeps, which suggests that a downstream area (such as the prefrontal cortex) may be the ultimate “decider” on which direction to turn each trial.

Unlike the extended reactivations of Davidson et al. (2009), the putative VTE sweeps of Johnson and Redish (2007) were not associated with SW-R events. Nonetheless, both studies demonstrate that during pauses in movement, the hippocampus briefly activates representations of extended behavioral sequences that are far removed from the current location of the animal. Thus, the hippocampus does not just represent the current location of the animal; it is capable of transiently representing nonlocal trajectories, which may be useful for route planning, decision-making, or other forms of cognition in which imagining other locations or events removed from the current sensory experience of the organism may be adaptive. It is noteworthy that human hippocampal amnesic patients have been shown to have a deficit in creating detailed narratives of imaginary events (Hassabis et al., 2007). Could the nonlocal representations of past experience shown in the rodent studies be a precursor of the hippocampus-dependent ability of humans to imagine completely novel experiences?

Of course, these interesting speculations about the functions of extended, nonlocal reactivations are just that: speculations. After many years, it has still not been demonstrated that the sleep reactivation is involved in memory consolidation; it has not been shown that reverse replay is involved in the credit assignment problem; it has not been shown that the extended sweeps at choice points are involved in VTE behavior. It is possible that these results are all epiphenomena of hippocampal circuitry and plasticity mechanisms, which may be necessary for proper functioning during active locomotion but which have no well-defined functions during sleep, restfulness, or pauses. Proving that these phenomena serve specific functional roles in cognition will be a difficult task that requires clever behavioral tasks and neuronal activity markers that may be beyond current technological limitations. For example,

it would be interesting to see if one could train the rat to produce specific nonlocal representations during pauses and SW-Rs, which would suggest that these events are under the control of the rat. Further, the increasing development of sophisticated imaging tools and molecular-genetic techniques may someday allow the manipulation of the specific neural circuits that generate these events. Future research may ultimately expand on these results to decipher the neural mechanisms underlying internally generated, conscious thought that is removed from external sensory stimulation. Imagine the possibilities!

**REFERENCES**

Buzsaki, G. (1986). *Brain Res.* 398, 242–252.

Davidson, T.J., Kloosterman, F., and Wilson, M.A. (2009). *Neuron* 63, this issue, 497–507.

Diba, K., and Buzsaki, G. (2007). *Nat. Neurosci.* 10, 1241–1242.

Foster, D.J., and Wilson, M.A. (2006). *Nature* 440, 680–683.

Hassabis, D., Kumaran, D., Vann, S.D., and Maguire, E.A. (2007). *Proc. Natl. Acad. Sci. USA* 104, 1726–1731.

Jackson, J.C., Johnson, A., and Redish, A.D. (2006). *J. Neurosci.* 26, 12415–12426.

Johnson, A., and Redish, A.D. (2007). *J. Neurosci.* 27, 12176–12189.

Lee, A.K., and Wilson, M.A. (2002). *Neuron* 36, 1183–1194.

Louie, K., and Wilson, M.A. (2001). *Neuron* 29, 145–156.

O'Neill, J., Senior, T., and Csicsvari, J. (2006). *Neuron* 49, 143–155.

Skaggs, W.E., and McNaughton, B.L. (1996). *Science* 271, 1870–1873.

Wilson, M.A., and McNaughton, B.L. (1994). *Science* 265, 676–679.

## Harnessing Chaos in Recurrent Neural Networks

Dean V. Buonomano<sup>1,2,3,\*</sup>

<sup>1</sup>Department of Neurobiology

<sup>2</sup>Department of Psychology

<sup>3</sup>Brain Research Institute

University of California, Los Angeles, Los Angeles, CA 90095, USA

\*Correspondence: [dbuono@ucla.edu](mailto:dbuono@ucla.edu)

DOI 10.1016/j.neuron.2009.08.003

In this issue of *Neuron*, Sussillo and Abbott describe a new learning rule that helps harness the computational power of recurrent neural networks.

One of the ironies of the human condition is that while many of our cognitive abilities derive from the complex dynamics of recurrent networks of neurons, we are quite inept at understanding neural dynamics in these same networks. The challenge of understanding the behavior of complex systems is, of course, not unique to neuroscience. Scientists have long struggled with how macroscopic properties emerge from interactions between a large number of simple elements, such as how complex social structures emerge from the interaction of individual members of insect colonies, or

how biomolecules give rise to life itself. However, the problem seems particularly perverse in neuroscience, since no one would describe a neuron as a simple element, and the synaptic strengths—which ultimately determine the nature of the interactions—continuously vary as a result of short- and long-term forms of plasticity.

Over the past 3 decades, neuroscientists have sought to understand the rules governing synaptic weights and have examined how computations can be achieved by appropriately adjusting these weights. These studies have resulted

in experimentally derived learning rules, such as spike-timing-dependent plasticity (Dan and Poo, 2004), and biologically implausible but computationally powerful algorithms, such as backpropagation (Rumelhart et al., 1986). The computational potential of these rules has focused primarily on feedforward networks, largely because they are not burdened with the complexities that can result in chaotic dynamics or epileptic-like “runaway excitation.”

Despite some progress with artificial (Hopfield, 1982) and experimentally derived (Izhikevich, 2006; Legenstein