Cellular/Molecular

Short-Term Memory in Networks of Dissociated Cortical Neurons

Mark R. Dranias, Han Ju, Ezhilarasan Rajaram, and Antonius M. J. VanDongen

Program in Neuroscience and Behavioral Disorders, Duke-NUS Graduate Medical School, Singapore 169857

Short-term memory refers to the ability to store small amounts of stimulus-specific information for a short period of time. It is supported by both fading and hidden memory processes. Fading memory relies on recurrent activity patterns in a neuronal network, whereas hidden memory is encoded using synaptic mechanisms, such as facilitation, which persist even when neurons fall silent. We have used a novel computational and optogenetic approach to investigate whether these same memory processes hypothesized to support pattern recognition and short-term memory *in vivo*, exist *in vitro*. Electrophysiological activity was recorded from primary cultures of dissociated rat cortical neurons plated on multielectrode arrays. Cultures were transfected with *ChannelRhodopsin-2* and optically stimulated using random dot stimuli. The pattern of neuronal activity resulting from this stimulation was analyzed using classification algorithms that enabled the identification of stimulus-specific memories. Fading memories for different stimuli, encoded in ongoing neural activity, persisted and could be distinguished from each other for as long as 1 s after stimulation was terminated. Hidden memories were detected by altered responses of neurons to additional stimulation, and this effect persisted longer than 1 s. Interestingly, network bursts seem to eliminate hidden memories. These results are similar to those that have been reported from similar experiments *in vivo* and demonstrate that mechanisms of information processing and short-term memory can be studied using cultured neuronal networks, thereby setting the stage for therapeutic applications using this platform.

Introduction

Short-term memory enables animals to preserve stimulus information across brief delays (Baddeley, 1996). The mechanisms underlying these stimulus-memory processes involve transient changes in the active and hidden states of neurons in a neuronal network (Buonomano and Maass, 2009). This raises the question whether random networks formed by cultures of dissociated cortical neurons also have the ability to store stimulus information for short periods of time. To test whether the same mechanisms that support short-term memory *in vivo* can encode stimulusspecific memories *in vitro*, we employed a novel combination of optogenetic and computational approaches.

Active or fading short-term memories are encoded in the action potential activity reverberating in a neuronal network. These signals preserve stimulus identity information until the activity subsides (Wang, 2001). Hidden memory traces store information as a spatial pattern of synaptic adaptation, which persists even

Received June 6, 2012; revised Nov. 11, 2012; accepted Dec. 7, 2012.

The authors declare no competing financial interests.

Correspondence should be addressed to Dr. Antonius M. J. VanDongen, Program in Neuroscience and Behavioral Disorders, Duke-NUS Graduate Medical School, 8 College Road, Singapore 169857. E-mail: antonius.vandongen@duke-nus.edu.sg.

DOI:10.1523/JNEUROSCI.2718-12.2013

Copyright © 2013 the authors 0270-6474/13/331940-14\$15.00/0

when neurons fall silent (Buonomano and Maass, 2009). These hidden memories temporarily alter the timing or synaptic transmission of action potentials, and the presence of these memories can be revealed by the altered electrophysiological response of a network to additional stimulation (Mongillo et al., 2008).

Previous investigations into the information processing capabilities of neuronal networks formed from dissociated rat cortical neurons plated on multielectrode arrays (MEAs) have utilized electrical stimulation and shown that these networks can actively represent stimulus information (Bakkum et al., 2008b). In these studies, the response of the network to electrical stimulation is characterized by two phases: an early phase of short latency action potentials (3-20 ms) that are directly evoked by stimulation, followed by a burst of firing activity across the network that peaks between 20 and 200 ms (Eytan and Marom, 2006; Gal et al., 2010; Kermany et al., 2010). These network bursts are characterized by a simultaneous, highly correlated firing of neurons across the culture, after which activity dissipates to baseline (Beggs and Plenz, 2003; Wagenaar et al., 2006a). A strong relationship exists between the distance of a neuron to the stimulating electrode and the latency of its earliest response, enabling researchers to reconstruct the location of the stimulating electrode, often using latency information alone (Eytan et al., 2003; Dockendorf et al., 2009; Kermany et al., 2010).

The capacity of cultured neuronal networks for storing stimulus information or performing state-dependent computations using active and hidden processes has not been studied *in vitro*. We have therefore investigated the responses of cortical neurons to complex stimuli in order to answer the following questions regarding short-term memory mechanisms in randomly orga-

Author contributions: M.R.D. and A.M.J.V. designed research; M.R.D. and E.R. performed research; M.R.D., H.J., E.R., and A.M.J.V. contributed unpublished reagents/analytic tools; M.R.D., H.J., and A.M.J.V. analyzed data; M.R.D. and A.M.J.V. wrote the paper.

This work was supported by United States National Science Foundation Grant ECCS-0925407 and an award from the Singapore Ministry of Health and A*STAR, the Agency for Science, Technology and Research. We thank Hendrika VanDongen for expert technical advice, Niamh Higgins for technical support, and Dr. Karl Deisseroth for the ChR2-EYFP plasmids.



Figure 1. Optical path for the custom laser projection system. The beam of a 25 mW 488 nm laser passes through an acoustooptic tunable filter onto a reflective LCoS SLM. Blue light patterns reflecting off the SLM were projected onto the neuronal network using the lenses as shown. Components enclosed by the dotted outline are located inside the inverted microscope.

nized cultured neuronal networks: 1) What are the response dynamics to the optical presentation of complex spatial stimuli? 2) How long do short-term memories of stimuli persist? 3) Can cultured neuronal networks store stimulus-specific information using both active and hidden memory mechanisms? Clarifying the mechanisms of short-term memory has implications for treating memory disorders ranging from transient cognitive impairment to Alzheimer's disease.

Materials and Methods

Primary neuron cell culture. Cortical neurons obtained from rat E18 embryonic brain of either sex were dissociated and plated as described previously (Van de Ven et al., 2005; Dranias et al., 2011) on MEA-containing culture dishes (Multi-Channel Systems). Two different electrode configurations were used: a standard 60 electrode dish arranged in an 8 × 8 grid and a 252 electrode dish with electrodes arranged in a 16×16 grid. Prior to plating, MEAs were cleaned with detergent, rinsed, sterilized overnight under UV lighting, and coated with poly-D-lysine and fibronectin (Invitrogen). Cortical tissue from E18 rats is dissociated using papain and 500 μ l of DNase. The resulting cell suspension was rinsed, centrifuged, and then the pellet was resuspended in NB-Active 4 cell medium (BrainBits), in some experiments supplemented with 10% FBS. Neurons slated to be transfected underwent electroporation at this point and then were plated onto the poly-D-lysine and fibronectin-coated MEA dishes. Cell culture was maintained in NB-Active 4 cell medium to suppress astrocyte growth.

Optogenetic transfection. Plasmid DNA encoding Channel Rhodopsin-2 (ChR2) was a kind gift from Karl Deisseroth. ChR2 was fused to EYFP for visualization, and mutations H134R and T159C were introduced to boost current (Nagel et al., 2005; Berndt et al., 2011). The plasmid included a cytomegalovirus promoter. Two transfection methods were used: electroporation (Amaxa nucleofector II kit, Lonza) performed before plating, or lipofection (Lipofectamine 2000, Invitrogen) at least 24 h before recordings on 7–14 days *in vitro* (DIV). ChR2-EYFP expression was observed within 24 h after transfection.

MEA recordings. Extracellular electrophysiological recordings of neurons were made from 60 and 252 electrode MEA dishes using the MEA1060 and USB-MEA256 hardware systems, respectively (Multi Channel Systems). Recordings were performed on an antivibration table and in a Faraday cage. Signals were amplified and sent to a personal computer using a PCI-based data acquisition card for the MEA1060 system or via high speed USB for the USB-MEA256. During experimental recordings, the cell culture medium of each MEA dish was replaced

with Dulbecco's phosphate-buffered saline solution containing glucose and pyruvate (Sigma).

Data analysis. MC_Rack software (Multichannel Systems) was used to acquire electrophysiological signals in real-time. Extracellular signals were high pass filtered at 300 Hz and low pass filtered at 3 kHz with second order Butterworth filters. Action potentials were detected using a voltage threshold rule whose value was determined by the user for each dish based on the amount of channel noise. When an action potential was detected, the timing of the spike was recorded together with a 2 ms sample of the waveform, sampled at 120 kHz.

Optical stimulus presentation and imaging. The MEA system was mounted on an inverted microscope (Eclipse Ti, Nikon) during recordings. A 25 mW 488 nm laser (Spectra-Physics) beam was projected through an acousto-optic tunable filter (AA Opto-Electronic), which allows the laser light to be controlled by a TTL pulse generated by the personal computer. The beam was optically expanded, polarized, and projected at a 45° angle onto a reflective LCoS Spatial Light Modulator microdisplay (SLM,

Holoeye Photonics AG) with a resolution of 1920×1200 pixels and a 60 Hz refresh rate. The status of the pixels is controlled using the DVI port of a personal computer allowing images to be generated in MATLAB. Blue light patterns reflecting off the SLM were passed through a second polarizing filter and projected onto the neuronal network growing on top of the MEA. Fluorescent and bright-field images were captured from the MEA dishes via a cooled CCD camera (Orca, Hamamatsu) and analyzed using NIS Elements software (Nikon). This enabled verification of optical stimuli and visual inspection of electrophysiologically active channels. Figure 1 illustrates the optical path of the custom laser projection system. Custom drivers and MATLAB software were used to generate the stimulus patterns, control the onset and offset of optical stimuli, and generate TTL pulses to synchronize electrophysiological recordings with the presentation of laser stimuli.

Stimulus material. Two sets of stimuli were used in these experiments. The first set of stimuli consisted of a sequence of lighted squares from a 2×2 grid or a 5 \times 5 grid. These stimuli were presented to assess the efficacy of transfection. The grid occupied a region of ~ 1.5 mm², and the individual squares in the grid were of variable size $(0.3-0.75 \text{ mm}^2)$. In the experiments testing short-term memory for complex stimuli, random dot stimuli were presented. The random dot stimuli were constructed from 18 to 22 randomly positioned squares on a 10×10 grid. When projected onto the MEA dish, the 10×10 grid occupied an area of \sim 1.25 mm². Effective light intensity was measured as 0.1 mW/mm². The light was projected through a low-power objective (Fig. 1), which allowed us to stimulate neurons located anywhere in a 1.5 mm² field of view. A consequence of this approach is that relatively long (100–200 ms) optical stimuli are required to elicit light responses. At this intensity, the mean latency to the first recorded neural action potential was 165 ± 4.7 ms (N = 115 neurons).

Preprocessing of network activity. Electrophysiological data were imported from MC_Rack into MATLAB using the Neuroshare API library (www.neuroshare.org). Data were grouped into single trial spike timing vectors. Before analysis by machine classifiers, spike timing data were counted and discretely resampled into individual spike times (liquidstate machine [LSM]) or nonoverlapping 5, 50, or 500 ms bins, depending on whether the data were to be analyzed as a time series or a single trial. Only spike timing data from electrodes reporting at least 1 spike per minute were included for analysis.

Machine classification of unit responses. Biologically plausible classification algorithms were implemented in MATLAB to analyze single trial data. The simplest and most plausible anatomy for such a model is a readout neuron with tunable weights that allow the neuron to make



Figure 2. Stimulus presentation protocol. *A*, Stimulus-specific memory trials. Thirty static patterns of randomly generated dots were presented to the neuronal network for 100 – 200 ms. Neuronal activity is recorded for 500 ms before stimulus onset and for 1–2 s after stimulus onset. *B*, Cue-Probe Short-term memory trials. Two stimuli were presented sequentially to the neuronal network: a cue stimulus followed by a probe stimulus after a short delay. The cue stimulus could be one of two (or four) random dot stimuli, whereas the probe stimulus was fixed across trials. Each stimulus was presented for 100–200 ms. Cue and probe stimuli are separated by a variable delay, ranging from 200 ms to 2 s.

binary category membership decisions. Two algorithms were chosen to tune these weights: a support vector machine (SVM) or LSM algorithm. All algorithms have distinct training and testing phases (Haykin, 1999). During the training phase, a sample of data is taken to tune the weights using the appropriate algorithm. In the testing phase, no new learning takes place and the trained neural network is used to classify the remaining testing data. This assays the generalizability of the categories learned during the training phase. Data were generally divided into 70–30% training and testing trials. Classifications were retrained and retested 5–50 times choosing random subsets. The accuracy of the classifier is taken as the mean percentage correctly classified test trials. In multiclass classification, each readout neuron is trained to make a "one versus all" category membership decision. Multiclass classifications require multiple readout neurons arranged in a winner-take-all competition to determine category membership.

The SVM was implemented using the online resource libsvm (Chang and Lin, 2011; http://www.csie.ntu.edu.tw/~cjlin/libsvm). Data were input as the mean activity in a 50, 100, or 250 ms time window. The dimensionality of the inputs and weights is equal to the number of active electrodes. As category membership decisions are binary outputs, the readout neuron was simplified and represented by just its synaptic weight matrix.

The LSM was used because living neural networks are known to use temporal signals to encode stimulus information (Kermany et al., 2010) and the LSM is sensitive to temporal patterns (Maass et al., 2002). The LSM was simulated using CSIM software (Natschlager and Markram, 2003). The LSM is a form of filtered classifier that nonlinearly separates patterns, increasing the probability that any two different inputs to the LSM will be linearly separable (Maass et al., 2002). The LSM consists of an input layer, a liquid filter or reservoir of randomly connected neurons, and a linear classifier that separates the signals that propagate through the neural reservoir. MEA data were used to drive the input layer to the liquid filter. To preserve spatial information of each electrode, the input layer consists of 8 \times 8 neurons arranged equivalent to the MEA layout. Each input neuron receives recorded spikes from one electrode and propagates the input into the reservoir of simulated neurons that contains $8 \times 8 \times 8$ neurons in a 3D space, with 80% of the neurons being excitatory and the remaining 20% being inhibitory. The Multiple-timescale Adaptive Threshold neuron model was used, with a firing threshold of 14 mV, resting membrane potential 0 mV, 30 ms time constant, and a refractory



Figure 3. ChR2 transfection. *A*, Light microscopic image of primary neuronal culture at DIV9 on a 60 electrode MEA transfected with plasmid DNA encoding Channel Rhodopsin-2 fused to Yellow Fluorescent Protein (ChR2-YFP). *B*, Image taken using a 4× objective and 510 nm excitation light to visualize ChR2-YFP expression. Cell count indicates approximately 2% transfection efficiency.

period of 2 ms (Kobayashi et al., 2009). Parameters for the threshold adaptation are set to: $\alpha_1 = 20 \text{ mV}$, $\alpha_2 = 5 \text{ mV}$, $\tau_1 = 2 \text{ ms}$, and $\tau_2 = 40 \text{ ms}$. Connection probability between two neurons was based on the Euclidean distance:

$$P(D) = C^* \exp\left(\frac{-D^2(a,b)}{\lambda^2}\right),$$

where D(a,b) represents the Euclidean distance between neurons *a* and *b*, $\lambda = 5$ and C = 0.5 for the input connections, and $\lambda = 7$ and C = 0.6 for connections in the reservoir. Values were chosen after parametric simulations to optimize classification on a test set of data. All the synaptic weights were randomly initialized following a γ distribution. Input connections are static, whereas the connections in the reservoir incorporate short-term plasticity with CSIM default parameters. A readout neuron is connected to all the neurons in the reservoir to make classification decisions. Its synaptic weights are optimized by Fisher's linear discriminant. For multiclass classification, multiple readout neurons were created and



Figure 4. Unit responses encode spatial stimuli. *A*, Left, Square stimuli were projected onto a DIV14 network for 200 ms using a 488 nm laser to excite ChR2 at different locations. Right, Images of the culture taken while illuminated by the two laser stimuli. Graphs in Figure 4*B* are based on an analysis of data from a single electrode at position (6,2) shown circled. *B*, *C*, Raster plots of action potentials recorded extracellularly from electrode position (6,2). Responses from 45 trials are shown. Stimulus onset and offset are indicated by green and red vertical lines, respectively. *C*, *D*, For Stimulus 1 (*C*) and Stimulus 2 (*D*), the average latencies were 92 and 156 ms.

a winner-take-all rule was applied using the average activity of the readout neuron over the trial. A new Liquid State Filter topology was generated each time the network was retrained and tested.

Classifier information time series. A time series was constructed to measure how stimulus information varies over time. The time series were constructed using multiple, independent LSMs or SVMs. Each LSM or SVM analyzed the spike times or mean firing rates from a small time window. For the LSM and SVM network, spike times were segmented into 100 or 50 ms bins. An independent LSM or SVM is dedicated to each bin and trained to classify the neuronal response on this interval. Using an array of independent classifiers allowed the neuronal response data to be replaced with a time series of data points, whose value represents the amount of information available at that moment for predicting which stimulus was presented at the start of the trial. The amount of available information was determined using the percentage of correctly classified testing trials. The classifiers were implemented independently, and this provides a lower limit on the amount of information (Nikolic et al., 2009). Unless otherwise stated, the performance measure is the classification accuracy on the testing set.

Sources of classification information. The SVM uses two basic sources of information to perform classification: the spatial pattern of activity and the magnitude of overall activity. The spatial pattern of activity consists of an activity vector that is the combination of units that are activated by stimulus presentation. The LSM performs classifications primarily relying on three sources of information: spatial pattern of activation (spatial information), the magnitude of overall activity (network spike rate), and the temporal pattern of spiking inputs (temporal information). The temporal encoding of a stimulus relates to sequences of spikes, spike latencies or variations in firing rate that are stimulus-specific; these properties are extracted by the reservoir functionality of the liquid filter. The contribu-

tion of each of these sources of information to classification accuracy was measured by filtering out the spatial and/or temporal information before training and testing. For the LSM, the role temporal information plays can be uncovered by comparing the rate of classification before and after randomizing the spike times. Spike times are randomly drawn from a uniform distribution over the same interval (50 or 100 ms bins). For both SVM and LSM, the role of spatial information can be uncovered by comparing the rate of classification before and after the pooling of active units. Pooling of active units eliminates all information around the spatial pattern of activation. After filtering out spatial information for the SVM and both spatial and temporal information for the LSM, the amount of information encoded by overall mean firing rate can be assessed.

Transfection assessment. These experiments were used to determine whether dishes had effective transfection. Four square stimuli from a 2×2 grid were presented sequentially, in a clockwise order. Each square stimulus was typically presented for between 150 and 250 ms and poststimulus response was monitored for 500 ms to 1.5 s. If a culture showed a reliable response to one or more squares, it was used in additional experiments. Dishes used in further experiments ranged from DIV8 to DIV20.

Experimental protocols. Two basic experimental protocols were followed as illustrated in Figure 2. Stimulus-specific memory trials sought to determine the capacity and duration of short-term memories in the MEA dish. Up to 30 different single random dot stimuli were presented for 100–200 ms and the poststimulus response was monitored for 1.5–2.5 s. The stimuli were presented in pseudorandom order, usually 800–1200 trials per dish. During

experiments where all 30 stimuli are presented, each stimulus is repeated \sim 30 times; whereas during other experiments where only four stimuli are presented, there were \sim 200 repetitions each. Spike times from electrophysiological recordings of single units were collected into 250 ms bins. Multiclass classification was used to analyze the network response to discover how many different stimuli can be discriminated using network responses. To determine the duration of stimulus-specific short-term memory, classification between 4 random dot stimuli was used. The specificity and duration of the memory were determined by a time series constructed using an array of LSMs or SVMs.

Short-term memory/delayed response trials sought to determine whether the dish can support short-term memories that could be used to solve a task. In these trials, two stimuli (a cue and a probe stimulus) are presented sequentially, separated by a pause. If memories of the cue stimulus influence the processing of the probe stimuli, then the network could use this memory to select actions or responses. The cue stimulus was either one of two or one of four random dot stimuli presented for 100–200 ms (Fig. 2*B*). The probe random dot stimulus was presented 200–2000 ms later for 100–200 ms. The amount of cue-specific information was measured using a time-series constructed by an array of LSMs or SVMs.

Statistical methods. Random label shuffling was used to measure the significance of the classification accuracies. Label shuffling was repeated 100–200 times, and the best and worst classification rates (95th percentile) were recorded for each classifier to establish the significance of the correct classification rates. Bias relating to the unfair sampling of training and testing vectors was controlled by retraining and testing the classifiers 50 times. For each repeated training and testing, a different set of training and testing data were randomly selected from the experimental data. The

mean and its SE are reported from these 50 repeated training and testing episodes. The false-positive rate was used in measuring the stimulus representation capacity of the cultured neuronal networks. The correlation between the distance of an electrode from the stimulus and the latency of its action potentials was computed. The distance of a neuron to the stimulus was measured as the distance of the recording electrode to the stimulus center. Distance was measured in electrode units. The latency was taken as the mean latency from stimulus onset until the first spike recorded from that electrode. Neurons were included in the analysis if they responded reliably to two of the square stimuli (20% of presentations) inside a 300 ms window beginning at stimulus onset. When comparing two conditions, the change in classification accuracy was calculated by recording the percentage change after the chance rate of correct classification was subtracted out from the signals under consideration.

Results

Primary cultures of dissociated E18 rat cortical neurons were grown on an MEA and transfected with a plasmid encoding ChR2-EYFP. The MEAs were plated with 10^5 cells, and the percentage of cells expressing ChR2-EYFP averaged ~2% (Fig. 3). Experiments were performed on cultures 7–14 DIV, when on average 60% of the electrodes revealed spontaneous neuronal activity and 10–15% of the electrodes were recording signals that could be reliably produced by presenting optical stimuli. The method of transfection influ-

enced the neuronal response; and while 28 MEAs had reliable responses to the presentation of monochrome square stimuli, the longest duration responses were recorded from MEAs transfected using electroporation (n = 8).

Representation of optically presented stimuli in living neuronal networks

First, we investigated how single units responded to simple spatial stimuli. Five MEA dishes were tested using square stimuli sequentially presented from a 5×5 grid. The square stimuli were $\sim 0.6 \times 0.6$ mm. Ninety-five units responded to the presentation of these squares; 13% of the units responded to only one stimulus, while the remainder responded to more than one stimulus. Only 49% of units responded when a stimulus was presented directly over the recording electrode. The remaining units (51%) could only be activated by illuminating regions more distal from the recording site. Distances from which stimuli were able to elicit responses ranged from directly overhead to as far as 2.5 mm away.

The latency at which single neurons responded to spatial stimuli did not increase with stimulus distance. Figure 4 displays data from experiments where single squares from a 2 × 2 grid were projected onto 10 MEA dishes (squares were 0.9×0.9 mm). No correlation was found between the latency to firing of a single unit and the distance of the stimulus from the site of the recording electrode (r = 0.10, p = 0.16, 119 neurons, 10 MEAs). When single squares from a higher resolution 5 × 5 grid were used, the



Figure 5. Spike Raster plots recorded from single units in response to optical stimulation. Optical stimuli were presented for 100 ms. Each row represents one trial. Spike times are represented as circles, which are color coded according to electrode number. Trials are aligned to the first spike time from the cyan unit and then sorted by the spike time recorded from the unit, which fired second on most trials. *A*, Stimulus 1 presented to a DIV14 MEA. An obvious alternating spiking activity is observed. The average latency of the cyan unit is 114 ms after stimulus onset (14 ms after stimulus offset). *x*-axis indicates time from earliest latency spike. *B*, Response of same network to a second random dot stimulus. Alternating activation in response to optical presentation of a random dot stimulus is seen. The average latency of the cyan unit is 126 ms after stimulus onset (26 ms after stimulus offset).

latency of units was found to be weakly but inversely correlated to the distance of the stimulus from the recording electrode (r = -0.11, p = 0.05, 95 neurons, 5 MEAs). Even when neurons that responded to multiple stimuli were examined individually, the average correlation between latency and distance for each unit was negligible (r = -0.01, 82 neurons, 5 MEAs). Although there is no systematic relationship between response latency and electrode distance, the latency of the responses of individual units was fairly deterministic and predictable (average coefficient of variation, 0.265). The average overall latency for these square stimuli was 137 \pm 63 ms.

Next, we investigated the dynamics of neuronal responses to more complex spatial stimuli. Random dot stimuli were projected onto the ChR2-expressing neurons to test the ability of the networks to discriminate complex stimuli. Optical stimulation with random dot patterns or squares led to an initial network response that is characterized by the selective activation of a few units with an average latency of 165 ms after stimulus onset (range, 53–239 ms; N = 115 neurons, 13 MEAs). The information carried by this initial response decays by 50% after 122 \pm 36 ms (n = 7 MEAs). On some trials, this initial response is followed by a synchronized burst of network activity. A typical network burst lasts between 200 and 600 ms, after which there is a return to baseline unit activity. If no burst occurs, the initial stimuluselicited response will continue to dissipate until baseline response levels are achieved after 200-300 ms. The occurrence of a network burst is dependent on past network activity and the identity

of the stimulus. Some stimuli elicit a network burst on 100% of trials, whereas others seemingly never elicit a network burst. During the initial 200 ms after optical stimulation, the response of single units does not follow a simple profile of activation followed by passive decay. Figure 5 plots the responses of single units from a DIV9 MEA. The units are color coded, and from Figure 5 it can be seen that most units start responding after the optical stimulus has terminated and the responses recorded at electrodes reveal a complex interplay between units that is typical of the activation observed in recurrent networks (Schrader et al., 2008).

Representational capacity for stimulus information

Network responses could reliably encode complex spatial stimuli consisting of \sim 20 randomly spaced dots. Thirty different random dot stimuli were presented to 8 cultured neuronal networks to measure how many stimuli can be reliably discriminated using electrophysiological recordings of network activity. The stimuli were presented in pseudorandom order using the protocol shown in Figure 2A. For individual dishes, the mean rate of correct classification measured on the testing set ranged from 0.14 to 0.5. Classifications were performed on a 250 ms bin of network activity that began at stimulus termination. Subsequent bins were excluded from the analysis as classification accuracy fell precipitously thereafter. Figure 6A shows that, on average, 28 of the 30 stimuli were classified at rates above chance when an SVM classifier was used, whereas 29 of 30 stimuli were classified at rates above chance when the LSM was used (n = 8 MEAs).

Machine classifiers can address the question of how stimulus identity information is encoded by the network. Stimuli can be encoded in the temporal patterns of spiking activity, overall spike counts, or the spatial patterns of active channels (Kermany et al., 2010). How stimulus identity information is encoded can be analyzed by filtering out temporal and/or spatial information prior to classification. We have employed a liquid state machine (LSM) classifier for this purpose (see Materials and Methods). Figure 6B shows that both temporal and spatial information plays a significant role in stimulus encoding. Eliminating temporal structure by randomizing the spike times reduces classification accuracy by 30%. When the spatial pattern of channel activation is eliminated by pooling the data, there is a 69% reduction in classification accuracy. When both spatial and temporal information is eliminated, only the overall spike count can be used for classification. This results in a 94% drop in classification accuracy (almost to chance levels). In this case, only 5 stimuli are discriminated at rates above chance. These results indicate that the spatial pattern of active units is the most important source of information for discriminating these spatial stimuli, although temporal information still makes a significant contribution.

Fading memory processes

To measure the duration of active stimulus memories, four random dot patterns were presented in pseudorandom order onto 18 networks, using the protocol shown in Figure 2A. The duration of active or fading memories was defined as the time that elapsed between the end of the optical stimulus and the time when stimulus-specific information was no longer detectable in the spiking activity of the cultured neuronal networks. When an SVM was used to classify these stimuli, the average stimulus memory persisted for 354 ± 70 ms (range, 14-1480 ms) after the optical stimulus was terminated. Analyzing the data from four dishes using the multiclass LSM revealed that stimulus-specific information persisted in these networks for an average of 675 ms after stimulus (range, 150-1200 ms). Figure 7 shows how classi-



Figure 6. Representational capacity and encoding of stimulus information. Classifiers were trained on MEA data to discriminate neuronal responses to 30 different random dot stimuli. A, Comparison of different methods for classifying stimulus-elicited responses: multiclass SVM classifier (hyphenated lines) versus multiclass LSM classifier (solid lines). Data from eight networks were analyzed using eight separate SVM and LSM classifiers. Stimuli are ranked in order of classification accuracy. Classifiers were trained using a 250 ms window. Blue lines indicate the mean correct classification rate across all 8 networks; and red lines indicate the rate of falsepositive classification associated with each stimulus. When the red line is above chance, this can indicate that the classifier has a measurable lack of specificity. Shaded areas represent the SEM. The chance rate of correct classification was 3.3% (hyphenated black line). Data points for individual networks were obtained by dividing the data into 70%/30% training/testing sets and then taking the mean of 50 epochs of retraining and retesting. **B**, Temporal and spatial information were filtered out to determine how much stimulus information is encoded by the spatial pattern of active units or by temporal variations in firing rate. A solid blue line indicates the control rate of correct classification of the LSM for 30 different random dot stimuli. When temporal information is eliminated by replacing actual spike times with an equal number of randomized spike times, a slight drop in classification accuracy is seen (hyphenated blue lines). When information about electrode position is eliminated by pooling the data across electrodes into a single channel, there is a very large drop in classification accuracy (dotted blue lines). Other conventions as in Figure 6A.

fication accuracy varies over time, for one of the networks analyzed using a multiclass LSM.

Synchronized network bursting is a universal property of cultured neuronal networks (Maheswaranathan et al., 2012). We therefore investigated whether the duration of active stimulus memories differs during trials with bursting and nonbursting network responses. The responses of the neuronal network to the same stimulus were not simple and homogeneous: some presentations of a stimulus would elicit a network burst, whereas others would not. Figure 8 displays data from the same neuronal network as shown in Figure 7, except that the trials are sorted into bursting and nonbursting trials. When these bursting and nonbursting trials are analyzed separately using a multiclass LSM, stimulus information is observed to persist for different durations. On nonbursting trials, the mean rate of correct classification was 28% above chance, but the duration of the signal was



Figure 7. Multiclass classification of neuronal responses to four random dot stimuli. *A*, Time series indicates the duration of stimulus-specific memories by charting how stimulus classification information varies in time. Data are from a DIV11 neuronal network. Time series is composed of points constructed from an array of independent multiclass LSM classifiers that analyze nonoverlapping 50 ms time windows of the spike timing data. The multiclass LSM classifiers were trained to discriminate 4 stimuli; 70% of the data was used for training and the remaining 30% for testing. The mean rate of correct classification during training and testing trials is plotted in red and blue, respectively. Shaded regions indicate the SEM across 50 epochs of retraining and retesting. Chance rate of classification for 4 stimuli is 25% (hyphenated black line). Dotted lines indicate the 95% confidence intervals for random label reassignment. *y*-axis indicates the proportion of correctly classified trials; and *x*-axis, time in seconds. *B*, Average spike rate per trial (Hz) across all channels in the network. A network burst tends to occur in the vicinity of 300 ms after stimulus presentation. *C*, The four random dot stimuli used in these experiments. *D*, Heatmap of average spike rates per electrode channel generated in response to presenting the four experimental stimuli. Spike rates were calculated in 20 ms bins; color values range from 0 to 2 Hz. Data are from 45 electrodes.

only 120 ms (n = 7 MEAs). On bursting trials, the mean rate of correct classification was 14% above chance, whereas stimulus-related information could be correctly classified for 400 ms (n = 7 MEAs). The multiclass LSM indicates that the initial network response to the optical stimuli contains the highest amount of stimulus-specific information for both trial types (Fig. 8*B*,*C*). However, network bursts contain a small but significant amount of stimulus-specific information. This indicates that network bursts are capable of encoding information about stimulus identity.

The mechanisms for encoding stimulus information vary during bursting and nonbursting trials. Filtering out spatiotemporal features of the data before classification reveals that temporal spike patterns play a more important role representing stimulus information on bursting trials than on nonbursting trials. In contrast, the spatial pattern of active channels is more important for discriminating stimuli during nonbursting than bursting trials. Although significant, overall spike count was the least informative source of information on stimulus identity. When spatial information was eliminated by pooling, the accuracy of classification was reduced by 45% during nonbursting trials and 24.8% during bursting trials. When temporal information was eliminated by randomizing spike times, the accuracy of classification was reduced by 9.7% during nonbursting trials and 19.7% during bursting trials (n = 7 MEAs). This sensitivity to temporal information on bursting trials seems to correspond to stimulusrelated differences in the time at which a network burst is initiated. Finally, when both spatial and temporal information was eliminated and only overall spike counts were available for classification, classification accuracy was reduced by similar degree for both nonbursting (66%) and bursting trials (62.5%). Although the spatial pattern of active units is important for both bursting and nonbursting trials, it is not static (e.g., Figure 5). To measure whether the subset of active units that is activated by a stimulus is preserved across the fading memory, an SVM was trained on a 50 ms window of network activity that began just after stimulus offset. This SVM was then used to classify network responses recorded from successive 50 ms time bins. Generalization was extremely weak and limited to adjacent bins (n = 7MEAs).

Hidden memory processes

We have investigated the presence of hidden memory processes using sequentially presented stimuli according to the cue-probe protocol shown in Figure 2*B*. In the cue-probe task, two stimuli are presented sequentially and a classifier is used



Figure 8. Encoding of stimulus information on bursting and nonbursting trials. *A*, Raster plot of action potentials. The times at which action potentials were recorded from different electrodes in response to the presentation of the four different experimental stimuli were pooled across the network and are shown plotted as black hash marks. Dish is the same as Figure 6. *y*-axis indicates individual trials. *B*, Encoding of stimulus information on nonbursting trials. Hyphenated black line indicates accuracy of classification of testing data by LSM. Thin solid blue line indicates classification accuracy when temporal information is filtered out by randomizing spike times ('jittering'), and classification is performed relying only on overall firing rate and the pattern of active electrodes. Solid green line indicates classification accuracy when spatial information is eliminated by pooling spikes across all electrodes into a single channel, leaving temporal information and overall firing rate to guide classification. Dotted red line indicates classification is 0.25 (solid gray line). *C*, Encoding of stimulus information is 0.25 (solid gray line). *C*, Encoding of stimulus information on bursting trials. Conventions as in *B*.



Figure 9. Response to probe stimulus is modulated by the identity of previous cue stimuli. *A–D*, Four time series are shown, revealing the amount of cue-related information in the network responses. For each experiment, one of two different random dot cue stimuli ("X/Y") was presented. On experimental trials (red), the cue stimulus was followed by a probe stimulus ("Z"). On control trials (blue), the probe stimulus was not presented. Independent binary LSM classifiers analyzed nonoverlapping 50 ms time windows of the spike timing data; 70% of the data was used for training and the remaining 30% for testing. The age of the cultures and pause between cue and probe were as follows: *A*, DIV11, 800 ms; *B*, DIV9, 400 ms; *C*, DIV8, 200 ms; *D*, DIV9, 1 s.

to analyze the response of the network to the probe stimulus and determine whether probe-evoked activity is the same on every trial or whether it changes based on the cue stimulus that preceded it. If the classifier can discriminate which cue preceded the probe at a rate above chance, this constitutes a "one-back" memory (Nikolić et al., 2009). Ten in vitro networks were tested using this protocol; and in 6 of 10 networks, such a one-back memory was seen: the response of the network to the probe stimulus varies according to the cue stimulus that precedes it. For the dishes with such oneback memory, analysis of a 200 ms time window after probe presentation shows that there was a recovery of cue-related information relative to control. On trials where the probe was presented, the identity of the cue could be correctly discriminated at a rate of 22 \pm 6% above chance. In comparison, on trials where no probe stimulus was presented, the identity of the cue was correctly classified at a rate of $4 \pm 2\%$ above chance. Presentation of a probe stimulus unrelated to the cue therefore increased classification accuracy in each case. Figure 9 shows several cases where cue related information peaks during cue presentation, dissipates, and then recovers when

a second, unrelated, probe stimulus is presented. The delay between cue and probe stimuli varied between 200 and 1500 ms in these experiments.



Figure 10. Response to probe stimulus is modulated by identity of cue stimuli. *A*, Time series charting cue-related information during an experiment with a DIV11 culture. Cues are presented for 100 ms, the delay is 1.2 s, and the probe is presented for 100 ms. All stimuli are random dot patterns. The amount of cue-related information is measured using a time series constructed from an array of binary LSM classifiers. Presentation of the probe stimulus causes a recovery of cue-related information, and that

The duration and capacity of shortterm memory during cue-probe experiments were much larger than what would be predicted from the analysis of active (fading) network responses. The duration of the one-back memory appears to terminate at some point after 1.5 s. Figure 10 shows data from an experiment where the delay between cue and probe stimuli was 1.2 s. Short-term memory could not be detected at delays of 2 s or longer. At a shorter delay up to 1400 ms, the response to the probe stimulus was significantly modulated by the identity of the cue stimulus that preceded it.

It is possible that this one-back memory does not encode stimulus-specific information and simply reflects the influence of habituation on an all-or-none network burst. To investigate this question, we performed four experiments where we used 4 different cue stimuli instead of the usual 2. A strong one-back memory was observed in 2 of 4 dishes. Figure 11A charts the amount of cuerelated information recorded from a network whose response to a probe stimulus was modulated by all 4 preceding cues, resulting in a classification accuracy rate of 90%. Filtering the classifier data indicates that, for this dish, cue-related information was primarily encoded by the spatial configuration of active units.

To determine whether additional cuerelated information could be revived when a second probe stimulus was presented, a cue-probe-probe stimulus sequence was presented to four dishes. Two of four dishes showed a one-back memory, where the probe presentation elicited an increase in the amount of cue-related information. However, in none of the dishes did the presentation of a second probe lead to a similar recovery of cuerelated information. Figure 11*B* shows a case where one-back but no two-back memory was seen.

information persists at levels above chance for a further 1.2 s. Classifiers were trained and tested on 492 trials (246 trials per stimulus). Chance rate of correct classification is 50% (hyphenated black line). Other conventions are as in Figure 6. **B**, Network activity in response to Stimulus 1 (blue) and Stimulus 2 (red). Lines indicate the mean spike rate per trial (Hz) across all channels in the network. A network burst tends to occur in the vicinity of 300 ms after stimulus presentation. **C**, Raster plot of action potentials. Responses to Stimulus 1 and Stimulus 2 are shown on the left and right, respectively. Vertical lines indicate stimulus onset and termination. **D**, Heatmap of mean spike rates for each electrode. Spike rates were calculated in 20 ms bins; color values range from 0 to 2 Hz. Other conventions as in Figure 6.



Figure 11. Encoding of short-term memory and test for two-back memory. *A*, Time series charting cue-related information during an experiment with a DIV11 culture. Cues are presented for 150 ms, the delay is 350 ms, and the probe is presented for 150 ms. All stimuli are random dot patterns. Hyphenated black line indicates accuracy of classification of testing data by LSM. Thin solid blue line indicates classification accuracy when temporal information is filtered out. Solid green line indicates classification accuracy when spatial information is eliminated by pooling spikes. Dotted red line indicates classification rate based only on overall spike count. Data indicate that most stimulus information is encoded in the pattern of active channels, not spike counts or timing. Chance rate of classification is 0.25 (solid gray line); other conventions as in Figure 7. *B*, Response of DIV7 dish to cue-probe trials. A "one-back" memory was observed, where the probe presentation elicited an increase in the amount of cue-related information. However, presentation of the second probe did not lead to a similar recovery of cue-related information. "A" and "B" are the two cue stimuli to be classified. Either "C" or "D" was randomly assigned to be the first probe stimulus. The graphs above present the average across these two cases, indicating that in these two cases a short-term memory could be detected. No "two-back" memory for cue identity was detected. Top, Time series showing available cue-related information. Bottom, Spike confidence intervals for a random label reassignment. Other conventions as in Figure 6.

Short-term memory relies on both active and hidden mechanisms

To investigate whether short-term memory relies on spiking activity or hidden memory processes to bridge the delay between cue and probe, responses from the previous experiment were sorted into bursting and nonbursting trials (Fig. 12). Two neuro-

nal networks showed that a short-term memory with lumped trials also showed a short-term memory on nonbursting trials. From the raster plots in Figure 12 it can be seen that no activity bridges the gap between cue and probe presentations. Nevertheless, the network responds differently to the probe depending on which stimulus preceded it. In Figure 12A, the network is presented with a random dot stimulus "X." When "X" is followed by the probe stimulus "Z," there is a short response to the stimulus but no network burst. In Figure 11B, the same probe stimulus has been preceded by the other random dot stimulus "Y." This time, however, the response of the network to the probe is very different: there is a network burst. For this dish, the cue stimulus appears to alter the probability that the probe will elicit a network burst. Figure 12C, D shows the responses of a second dish with very different dynamics. Although this dish had sparse activation on nonbursting trials, it also showed a short-term memory. In this case, the probe stimulus only activates 3-4 units in the multielectrode array and "shortterm" memory appears to be a result of cuedependent adaptation in two of the three units. Analysis of how the hidden memory was encoded in the network activity revealed that the information was primarily encoded in overall firing rate, whether the units engage in a network burst or not.

Hidden memory was disrupted during bursting trials. When the same dishes that displayed hidden memory during nonbursting trials were analyzed to determine whether there was evidence for shortterm memory during bursting trials, no evidence for short-term memory was observed. Figure 13 shows the bursting trial responses from the same neuronal networks as Figure 12. These results indicate that, on trials where a network burst occurs, the memory of the cue stimulus is lost and the response to the probe stimulus does not reflect cuerelated information. For both dishes, network bursts disrupted hidden memory for the cue.

Discussion

We have employed a combination of optogenetics, MEA recordings, and machine classifiers to investigate how stimulus in-

formation is encoded and stored in networks formed by dissociated cortical neurons in culture. The results from our experiments indicate that stimulus-specific information is maintained for durations longer than 1 s and that both active and hidden memory processes are present.

These experiments used a new optogenetic approach. Electrical stimulation elicits an artifact at the time of stimulation followed by a highly stereotyped, fixedorder response lasting 10-30 ms. The first 5-15 ms of this initial response is artifactual in that it appears to be mediated electrically, not synaptically (Kermany et al., 2010). A synchronized network burst often follows this initial response with a peak occurring between 20 and 200 ms, before a return to the prestimulus state (Jimbo et al., 1999; Eytan and Marom, 2006; Shahaf et al., 2008). With optical stimulation, there is no stimulus artifact, and the latency to the first recorded spike is much longer, between 50 and 200 ms. The initial response to optical stimulation typically involves fewer neurons and also is fairly stereotyped (Fig. 5). As with electrical stimulation, a network burst is usually induced by optical stimulation, but at a much later latency (typically in the range of 100-300 ms) and for a longer duration, typically lasting 300-500 ms before activity dissipates to prestimulation levels.

On both bursting and nonbursting trials, most actively encoded stimulus information is found in the first 200 ms. This is also a period where the network response is less variable. During this initial 200 ms, stimulus information is primarily encoded by the spatial configuration of active units. However, this is not a simple, static spatial pattern of units that resembles the image of the optical stimulus; only 50% of units respond to direct illumination and the units that encode the stimulus change over time, with the relationship between network activity and stimulus identity gradually deteriorating into randomness. On bursting trials, actively encoded stimulus information can persist as long as 1 s, but this information appears to be mainly encoded by differences in overall spike counts or temporal information, particularly differences in the onset and termination of network bursts.

When the network was challenged with a second "probe" stimulus, which contains no information regarding the first "cue" stimulus, the response to the probe was typical and primarily encoded in the spatial pattern of active units. However, there was evidence for a residual memory of the previously presented cue stimulus. The mechanisms that encode this short-term memory were not singular. In two dishes, the ability of the probe to recall cue information relied on the absence or presence of a network burst, and the short-term memory was primarily encoded by a change in overall spike count.



Figure 12. Hidden memory encodes stimulus identity in overall spike counts. *A*, Raster plot of sorted nonbursting cue-probe trials from same dish as Figure 6; 50 trials of data are shown. Conventions are as in Figure 7. *B*, Encoding of stimulus information. Both cue and hidden memory are primarily encoded in the overall spike rate (dotted red line). Filtering out of temporal information has no effect (solid blue line), and filtering out of single unit activity by pooling spikes has little effect (solid green line). Chance classification performance is 50% (gray line); classification of unfiltered nonbursting trials is shown as hyphenated black line. *C*, Raster plot of sorted nonbursting data from cue-probe trials recorded from a different DIV7 MEA. *D*, As with the other dish, the majority of stimulus information recorded in hidden memory is encoded in overall spike counts (dotted red line). Disrupting temporal (blue) or spatial (green) information fails to reduce classifier accuracy from control levels (hyphenated black line). Other conventions are as in Figure 10.



In another dish, short-term memory was mainly encoded by changes in the spatial pattern of active units.

Important questions concerning the information content of synchronous network bursts were also answered. Network bursts are capable of representing and extending stimulus information in active memory, but simultaneously it was shown that network bursts erased hidden memories of the cue. If hidden memory plays an important role in short-term memory, then these observations may have some implications for cognitive impairments that are associated with pathological discharges that can accompany seizure disorders. This is an area that requires more investigation.

The observation that a probe stimulus can elicit information about previously presented stimuli compares favorably with some theoretical models of short-term memory. Mongillo et al. (2008) proposed a model of working memory wherein a randomly organized neural network with adaptive synapses is presented with a stimulus and then periodically maintains the stimulus information by reactivating the network with a uniform or nonspecific probe input. Figure 9A shows one such case where we used a uniform white stimulus as the probe and a observed a probe-dependent extension of cue-related information However, whenever we administered a second probe stimulus, we were unable observe a "two-back" memory effect (Fig. 11B). The theory may, however, anticipate this result as one of its primary premises is that excitatory synapses in prefrontal cortex are facilitating. The excitatory neurons in our living neuronal networks are likely to mainly have de-

Figure 13. No short-term memory is observed on bursting trials. A, Raster plot of sorted bursting cue-probe trials from same dish as Figure 6; 50 trials of data are shown. Conventions are as in Figure 7. **B**, Encoding of stimulus information. No "short-term" memory is seen after probe presentation. Cue information is primarily encoded in the overall spike rate on bursting trials (dotted red line). Filtering out of temporal information has no effect (solid blue line), and filtering out of single unit activity by pooling spikes has little effect (solid green line). Chance classification performance is 50% (gray line); classification of unfiltered nonbursting trials is shown as hyphenated black line. C, Raster plot of sorted bursting data from cueprobe trials recorded from a different DIV9 MEA. **D**, As with the other dish, no short-term memory is seen on bursting trials. For this DIV9 MEA, cue information is mostly encoded temporally (blue and red). The loss of information is less when only spatial information is eliminated before classification (green). Chance classification rate is 50%. Other conventions are as in Figure 10.

pressing synapses, such as rat sensory cortex (Markram et al., 1998).

The results from our cue-probe experiment more closely match the results reported by Nikolić et al. (2009), who recorded from the visual cortex of anesthetized cats. Despite the large differences between our in vitro system and the in vivo system used by Nikolić et al. (2009) there are many similarities. The duration of memory for single stimuli is comparable with fading stimulus memories reported as lasting ~700 ms in vivo while in vitro the mean was 675 ms. The correlation between mean firing rate and correct classification rates is about the same in vivo (r = 0.65) and *in vitro* (r = 0.64), indicating that mean firing rate carries a similar amount of information in each system. The presence of a "one-back" memory but absence of a "two-back" memory was also reported in vivo. Nikolić et al. (2009) hypothesized that the absence of this "two-back" memory might be a mechanism associated with visual masking, but the fact that this trend exists in vitro suggests that it is a generic property of neural networks rather than an unusual property of visual cortex. The networks also have their differences. The in vitro network is prone to network bursts after \sim 300 ms. The ratio of inhibitory to excitatory neurons is similar for the two systems (Marom and Shahaf, 2002), so the added stability of the in vivo system likely is a function of its architecture. Another cause for this difference may be developmental rather than architectural, reflecting the absence of external inputs to the in vitro network during development. When networks growing on an MEA were continuously stimulated with random inputs, the development of synchronized network bursting was suppressed (Wagenaar et al., 2006b). Hence, external signals may be needed during development to ensure network stability, by helping to tune and balance the outputs of excitatory and inhibitory neurons.

The memory processes we demonstrate endow the network with the ability to perform state-dependent computations, such that the response to a stimulus varies with the identity of the preceding stimulus. This constitutes a form of "reservoir" computing (Buonomano and Maass, 2009): nonlinear interactions in time and space. The usual application for reservoir computing is pattern separation, and these nonlinear interactions allow similar signals to evolve into very different network states. The shortterm memory observed during the cue-probe presentation experiments demonstrates that networks of disassociated cortical neurons can have a nonlinear sensitivity to temporal order. At the same time, our time-series analysis of the network response to optical stimulation demonstrates that these networks have a fading memory that preserves stimulus information long enough to allow recurrent interactions. The only complication is that network bursts may limit some computations; although they can propagate stimulus information they appear to erase hidden memory.

In conclusion, optogenetic methods can allow researchers to explore neuronal mechanisms of memory *in vitro* using stimuli and experimental designs similar to those used *in vivo*. An *in vitro* platform provides the additional benefit of allowing researchers to rapidly screen for drugs that might have therapeutic effects, manipulate genes and gene expression, provide a test bed for stem cell research, or examine how computers might be used in conjunction with brain cells to enhance memory and information processing (Chiappalone et al., 2003; Bakkum et al., 2008a; Jain and Muthuswamy, 2008; Berger et al., 2011; Stephens et al., 2012). All these possibilities can have far reaching consequences

References

- Baddeley A (1996) The fractionation of working memory. Proc Natl Acad Sci U S A 93:13468–13472. CrossRef Medline
- Bakkum DJ, Chao ZC, Potter SM (2008a) Spatio-temporal electrical stimuli shape behavior of an embodied cortical network in a goal-directed learning task. J Neural Eng 5:310–323. CrossRef Medline
- Bakkum DJ, Chao ZC, Potter SM (2008b) Long-term activity-dependent plasticity of action potential propagation delay and amplitude in cortical networks. PloS One 3:e2088. CrossRef Medline
- Beggs JM, Plenz D (2003) Neuronal avalanches in neocortical circuits. J Neurosci 23:11167–11177. CrossRef Medline
- Berger TW, Hampson RE, Song D, Goonawardena A, Marmarelis VZ, Deadwyler SA (2011) A cortical neural prosthesis for restoring and enhancing memory. J Neural Eng 8:046017. CrossRef Medline
- Berndt A, Schoenenberger P, Mattis J, Tye KM, Deisseroth K, Hegemann P, Oertner TG (2011) High-efficiency channelrhodopsins for fast neuronal stimulation at low light levels. Proc Natl Acad Sci U S A 108:7595–7600. CrossRef Medline
- Buonomano DV, Maass W (2009) State-dependent computations: spatiotemporal processing in cortical networks. Nat Rev Neurosci 10: 113–125. CrossRef Medline
- Chang C-C, Lin C-J(2011) LIBSVM: a library for support vector machines. ACM Trans Intell Syst Technol 2:1–27.
- Chiappalone M, Vato A, Tedesco MB, Marcoli M, Davide F, Martinoia S (2003) Networks of neurons coupled to microelectrode arrays: a neuronal sensory system for pharmacological applications. Biosens Bioelectron 18:627–634. CrossRef Medline
- Dockendorf KP, Park I, He P, Príncipe JC, DeMarse TB (2009) Liquid state machines and cultured cortical networks: the separation property. Biosystems 95:90–97. CrossRef Medline
- Dranias MR, Ju H, Rajaram E, VanDongen AM (2011) Spatiotemporal and fading memory properties in an in vitro model of information processing. Soci Neurosci Abstr 37:239.206.
- Eytan D, Marom S (2006) Dynamics and effective topology underlying synchronization in networks of cortical neurons. J Neurosci 26:8465–8476. CrossRef Medline
- Eytan D, Brenner N, Marom S (2003) Selective adaptation in networks of cortical neurons. J Neurosci 23:9349–9356. Medline
- Gal A, Eytan D, Wallach A, Sandler M, Schiller J, Marom S (2010) Dynamics of excitability over extended timescales in cultured cortical neurons. J Neurosci 30:16332–16342. CrossRef Medline
- Haykin SS (1999) Neural networks: a comprehensive foundation, Ed 2. Upper Saddle River, NJ: Prentice Hall.
- Jain T, Muthuswamy J (2008) Microelectrode array (MEA) platform for targeted neuronal transfection and recording. IEEE Trans Biomed Eng 55:827–832. CrossRef Medline
- Jimbo Y, Tateno T, Robinson HP (1999) Simultaneous induction of pathway-specific potentiation and depression in networks of cortical neurons. Biophys J 76:670–678. CrossRef Medline
- Kermany E, Gal A, Lyakhov V, Meir R, Marom S, Eytan D (2010) Tradeoffs and constraints on neural representation in networks of cortical neurons. J Neurosci 30:9588–9596. CrossRef Medline
- Kobayashi R, Tsubo Y, Shinomoto S (2009) Made-to-order spiking neuron model equipped with a multi-timescale adaptive threshold. Front Comput Neurosci 3:9. CrossRef Medline
- Maass W, Natschläger T, Markram H (2002) Real-time computing without stable states: a new framework for neural computation based on perturbations. Neural Comput 14:2531–2560. CrossRef Medline
- Maheswaranathan N, Ferrari S, Vandongen AM, Henriquez CS (2012) Emergent bursting and synchrony in computer simulations of neuronal cultures. Front Comput Neurosci 6:15. CrossRef Medline
- Markram H, Wang Y, Tsodyks M (1998) Differential signaling via the same axon of neocortical pyramidal neurons. Proc Natl Acad Sci U S A 95: 5323–5328. CrossRef Medline
- Marom S, Shahaf G (2002) Development, learning and memory in large random networks of cortical neurons: lessons beyond anatomy. Q Rev Biophys 35:63–87. CrossRef Medline
- Mongillo G, Barak O, Tsodyks M (2008) Synaptic theory of working memory. Science 319:1543–1546. CrossRef Medline

- Nagel G, Brauner M, Liewald JF, Adeishvili N, Bamberg E, Gottschalk A (2005) Light activation of channelrhodopsin-2 in excitable cells of *Caenorhabditis elegans* triggers rapid behavioral responses. Curr Biol 15: 2279–2284. CrossRef Medline
- Natschlager T, Markram H (2003) Computer models and analysis tools for neural microcicuits. In: Neuoscience databases: a practical guide, 2003 edition (Kotter R, ed), pp 123–138. Norwell, MA: Kluwer Academic.
- Nikolić D, Häusler S, Singer W, Maass W (2009) Distributed fading memory for stimulus properties in the primary visual cortex. PLoS Biol 7:e1000260. CrossRef Medline
- Schrader S, Grün S, Diesmann M, Gerstein GL (2008) Detecting synfire chain activity using massively parallel spike train recording. J Neurophysiol 100:2165–2176. CrossRef Medline
- Shahaf G, Eytan D, Gal A, Kermany E, Lyakhov V, Zrenner C, Marom S (2008) Order-based representation in random networks of cortical neurons. PLoS Comput Biol 4:e1000228. CrossRef Medline

- Stephens CL, Toda H, Palmer TD, DeMarse TB, Ormerod BK (2012) Adult neural progenitor cells reactivate superbursting in mature neural networks. Exp Neurol 234:20–30. CrossRef Medline
- Van de Ven TJ, VanDongen HM, VanDongen AM (2005) The nonkinase phorbol ester receptor α 1-chimerin binds the NMDA receptor NR2A subunit and regulates dendritic spine density. J Neurosci 25:9488–9496. CrossRef Medline
- Wagenaar DA, Nadasdy Z, Potter SM (2006a) Persistent dynamic attractors in activity patterns of cultured neuronal networks. Phys Rev E Stat Nonlin Soft Matter Phys 73:051907. CrossRef Medline
- Wagenaar DA, Pine J, Potter SM (2006b) Searching for plasticity in dissociated cortical cultures on multi-electrode arrays. J Negat Results Biomed 5:16. CrossRef Medline
- Wang XJ (2001) Synaptic reverberation underlying mnemonic persistent activity. Trends Neurosci 24:455–463. CrossRef Medline