Liquid State Machines and Cultured Cortical Networks: The Separation Property

Karl P. Dockendorf [a], Il Park [a], Ping He [a], José C. Príncipe [b], Thomas B. DeMarse [a],*

[a]*Biomedical Engineering Department, University of Florida, Gainesville, FL, USA 32611*
[b]*Electrical Engineering Department, University of Florida, Gainesville, FL, USA 32611*

Abstract

In vitro neural networks of cortical neurons interfaced to a computer via multichannel microelectrode arrays (MEA) provide a unique paradigm to create a hybrid neural computer. Unfortunately, only rudimentary information about these *in vitro* network's computational properties or the extent of their abilities are known. To study those properties, a Liquid State Machine (LSM) approach was employed in which the liquid (typically an artificial neural network) was replaced with a living cortical network and the input and readout functions were replaced by the MEA-computer interface. A key requirement of the LSM architecture is that inputs into the liquid state must result in separable outputs based on the liquid's response (separation property). In this paper, high and low frequency multi-site stimulation patterns were applied to the living cortical networks. Two template-based classifiers, one based on Euclidean distance and a second based on a cross-correlation were then applied to measure the separation of the input-output relationship. The result was over a 95% (99.8% when non-stationarity is compensated) input reconstruction accuracy for the high and low frequency patterns, confirming the existence of the separation property in these biological networks.

Key words: liquid state machine, dissociated cortical tissue, separation property

1. Introduction

Dissociated neural cultures in conjuction with planar micro-electrode arrays (MEAs) have been used to study the computational properties of neural tissue (DeMarse and Dockendorf, 2005; DeMarse et al., 2001; Jimbo et al., 1999; Marom and Shahaf, 2002; Parodi et al., 1998; Ruaro et al., 2005; Shahaf and Marom, 2001). Although *in vitro* neural cultures differ appreciably from *in vivo* neuroanatomical organization, evidence of rudimentary learning (DeMarse et al., 2001; Marom and Shahaf, 2002; Shahaf and Marom, 2001), memory (Jimbo et al., 1998, 1999; Madhavan et al.,

2007; Maeda et al., 1998; Tateno and Jimbo, 1999) and some evidence for pattern recognition (Ruaro et al., 2005) have been reported. These living cortical networks (LCN) are spontaneously active producing complex patterns of activity whose dynamics are only partially understood (Madhavan et al., 2007; Parodi et al., 1998; Rolston et al., 2007; Segev et al., 2004; Wagenaar et al., 2006).

Here we explore the elicited activity of *in vitro* networks for computation using the Liquid State Machine (LSM) architecture recently proposed by Maass and colleagues (Maass et al., 2003). LSM is a new form of computing machine capable of universal computation in functional spaces. The model conceptually distinguishes two components of the computing machine: a linear or nonlinear dynamical system called the liquid, and a readout component (Figure. 1). The liquid state is the part of the internal state of the dynamical system

* Corresponding author

Email addresses: `karld@ufl.edu` (Karl P. Dockendorf), `memming@cnel.ufl.edu` (Il Park), `pennyh_01@yahoo.com` (Ping He), `principe@cnel.ufl.edu` (José C. Príncipe), `tdemarse@bme.ufl.edu` (Thomas B. DeMarse).

that is instantaneously accessible to the readout. *Unlike the Turing machine that is capable of universal computation in static discrete inputs, the LSM is universal for real time computing on sets of continuous time analog functions with finite memory (Maass et al., 2002), also called myopic maps (Sandberg, 1997).* Real world signals are continuous time analog signals carrying information in its local time structure, so LSMs are well matched to process sensory inputs.
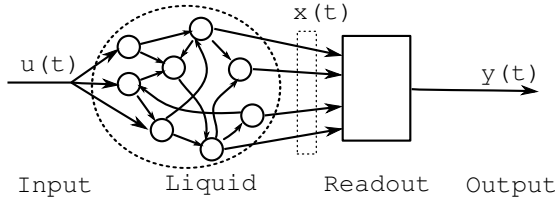

Fig. 1. Conceptual organization of a liquid state machine. The input is projected to the liquid (dynamical system consisting of spiking neurons), and the state of the liquid is projected through an instantaneous readout to an output.

One of the most surprising features of LSMs is that, when connected to a liquid consisting of a large number of artificial neurons, the readout is capable of achieving stable output even though the liquid state is a continuously varying high-dimensional pulse time series. This was unsuspected from previous work in neurocomputation, where dynamical systems with point attractors were postulated as necessary to read out stable information (**?**). Moreover, in spite of an instantaneous readout, enough information about the past of the input can be reconstructed (Maass and Markram, 2003; Maass et al., 2003). This also defies established procedures in optimal signal processing where storage of the input is commonly utilized to reconstruct the temporal structure of signals (e.g. the Wiener filter). The reason this is not so in LSMs is the recurrent and distributed nature of the state representation possessing different time scales. In a sense, the liquid states of individual neurons in the liquid $\{x_1(t), x_2(t), \cdots, x_i(t)\}$ can be thought of as a set of nonlinearly transformed traces of the input signal $u(t)$. As long as the set is complete, according to the projection theorem (Maass and Markram, 2004), the signal can be again recovered by a linear (or nonlinear) weighted sum implemented by the readout. This decomposition is rather different from the ones studied in approximation theory, where the bases are chosen *a priori* as simple time functions (polynomial basis or sine waves) and are independent of the input. Notice that the information from the liquid state is read out instantaneously, i.e. without any delay or transient response during the read operation, which makes LSMs

an ideal real-time processing machine. Maass and colleagues described two conditions that are necessary for the reconstruction to be exact: the separation property of the liquid state and the universal approximation property of the readout (Maass et al., 2002). However, there are so far no design principles to create an "ideal" liquid state for a special type of input (but see, Legenstein and Maass, 2007). The universality of the readout can be easily achieved and will not be addressed in this paper, see Maass et al. (2002) for discussion.
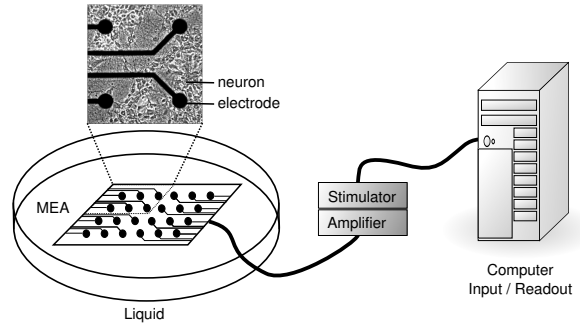

Fig. 2. Hybrid configuration of a biological system and a digital computer to implement liquid state machine (LSM). The liquid is implemented with real neurons grown on a micro-electrode array (MEA). A microscopic image (200x) snapshot of dissociated rat cortical neurons whose connectivity has developed over time and formed a highly interconnected network is shown. Extracellular electrodes are used to record the activity (liquid state), and the computer is used to perform the readout task. The extracellular electrodes also provide the input signal to the liquid. Spatially encoded voltage pulses are used to drive the local dynamics of the liquid in this paradigm.

In the conventional LSM approach, a simulated spiking neural network is used as the liquid. Our approach creates a hybrid biological-silicon computer by emplying a living network of *in vitro* rat cotrical neurons as the liquid. This network is interfaced to a silicon computer that serves as the readout stage based on spike activity (action potentials) recorded with a 60 electrode MEA. Inputs provided by electrically stimulating selected electrodes complete the LSM achieving the desired input-output mapping (Fig. 2).

Perhaps the most essential requirement of the biological subsystem to implement the liquid and achieve computation is the **separation property**, which ensures output decodability despite the stochastic nature of the biological neurons within the LCN. Unlike simulated neural networks where the initial state and network configuration can be controlled, an LCN exhibits complicated spontaneous dynamics (Rolston et al., 2007; Segev et al., 2004; Wagenaar et al., 2006) whose activity patterns evolve over the course of development

in vitro (van Pelt et al., 2004). These dynamics typically consist of episodes of increasing spontaneous asynchronous action potentials that culminate in a synchronous population burst of activity followed by a quiescent period. Simple inputs such as from a brief 200 us ±500 mV stimulation pulse delivered to a single electrode often evokes a burst of activity lasting over 100 ms and although the pattern of activity during the initial early phase of the network response is reliable (first 20 ms), the late phase (> 20 ms) is often more chaotic (Jimbo et al., 2000).

Here we advance the approach that separability of responses from site specific stimulation can be indirectly assessed in a classification context. Given different inputs to LCN, enough separation of the scattered output clusters guarantees that each output (network response) will contain unique information about the input. Since we encode the input to the liquid as spatial pattern of electrical stimulation, and use the response as the liquid state, a strategy that minimizes the chaotic late phase may provide a more accurate representation. In this paper, we demonstrate and compare two types of stimulation protocols to measure the seperability of the response; one inducing bursts (LFS; low frequency stimulation) and one that suppresses the burst response (HFS; high frequency stimulation), minimizing the chaotic late phase, and increasing potential input bandwidth. In both cases, we show the separability property is present by building a classifier that maps the liquid state to the input stimulation pattern. We begin by describing the LCN preparation and the interface with the computer. Then, we present and compare the performance of two classifier methods for separating spatio-temporal responses produced by the LFS and HFS protocols. Finally, we discuss the implications of this approach and its application to future LSM implementations with LCNs.

2. Methods

2.1. *Cell Culture*

Neuronal networks of rat cortical neurons were cultured using the methods similar to those discussed in prior work (e.g., Potter and DeMarse (2001)). Briefly, embryonic day 18 Sprague/Dawley rat cortical hemispheres obtained from a commercial supplier (BrainBits) were dissociated with Worthington Papain Dissociation System. About 20,000–50,000 cells were plated on each microelectrode array (MEA), which was pre-coated with 100 μL 0.1% polyethyleneimine (PEI, Sigma) and 10 μL laminin (Sigma) to enhance adhe-

sion and promote growth. The MEAs were covered with FEP lids (Potter and DeMarse, 2001), which reduce the culture media's evaporation, prevent infection, and allow gas exchange. Cell cultures were maintained in the 35.5 °C, 5% CO_2 for more than 1 month after which the pattern of spontaneous activity within these cultures becomes relatively stable (Kamioka et al., 1996; van Pelt et al., 2004). Half of the culture media, which consisted of Dulbecco's modified Eagle's medium (DMEM) (Gibco) containing 10% inactivated equine serum (HyClone), was replaced biweekly.

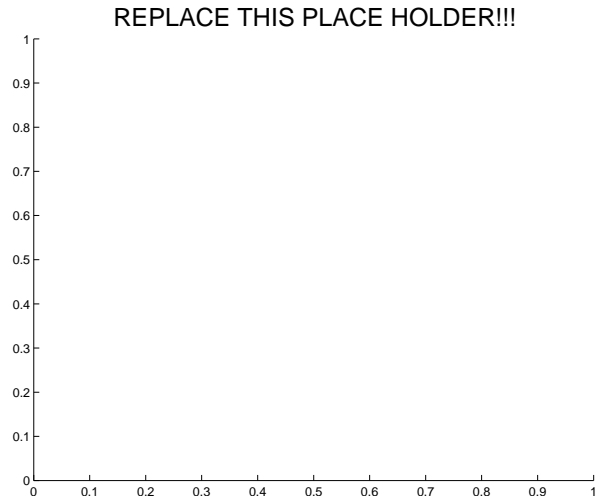### 2.2. *Acquisition*

REPLACE THIS PLACE HOLDER!!!

Fig. 3. Raw waveform trace before (upper) and after (lower) artifact filtering to reject noise from stimulation. Action potentials can be detected with a simple threshold method when the noise generated by stimulation is removed.

Neuronal activity was recorded extracellularly (sampling rate 25 kHz per electrode) using the MEA electrodes (8 by 8 grid), a 64-channel amplifier and acquisition board from MultiChannel Systems. Each MEA consists of sixty 30 μm electrodes spaced 200 μm apart. Data were collected using a Dell PC with Dual Intel Xeon 2.8 GHz processors with 3 GB RAM or an Apple Dual G5 2.0 GHz computer with 4 GB RAM. Spike detection was performed at a five times standard deviation threshold of estimated noise levels. Recordings were blanked (zeroed) for 2 ms following each stimulation pulse and any additional electrical artifact was removed using a local low-order polynomial fit (Wagenaar et al., 2005a), Figure 3.
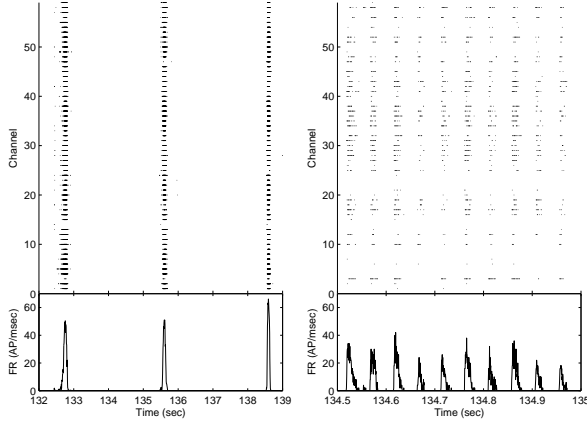
2.3. *Data and Stimulus Protocols*

Fig. 4. Stimulus Pattern. *Left*: Low frequency stimulation, $\frac{1}{3}$ Hz. *Right*: High frequency stimulation, 20 Hz. Note the differing timescales. In general, HFS results in individual responses with shorter duration than the burst responses of LFS that last longer than 100 ms. During stimulation, the average, culture-wide firing rate was 1.50 AP/ms (LFS) and 4.24 AP/ms (HFS).

Two distinct electrical stimulation protocols were used to provide inputs into the LCNs. The first protocol (LFS, n = 7 cultures) consisted of 10 repetitions of stimulation to each of the 60 MEA channels applied in random order at low frequency (1/3 Hz). Stimulations were biphasic 200 μsec pulses at 500 mV. Stimulation of most channels consistently induced population bursts lasting longer than 100 ms (see Fig. 4). The early phase of the response ($<$ 20 ms) likely represents activity from the direct activation of neurons nearby the stimulation electrode or neurons connected through short path lengths to neurons directly stimulated resulting in a reliable pattern of activity (Jimbo et al., 2000; **?**). In contrast, the response after this period (late phase) is often more variable and likely represents the reverberation of activity throughout the network. LFS also results in a refractory period over 100 ms in duration in which additional evoked bursts are not possible and therefore limits the potential input bandwidth for this protocol.

The second protocol (HFS, n = 3 cultures) sought to eliminate the late phase of the burst response and also increase the potential input-output rate of the system. In this protocol multichannel high frequency stimulation (20 Hz) was applied which has been shown to reduce the burst response to a tens of milliseconds in duration (Wagenaar et al., 2005b). The mechanisms under HFS that suppress bursting are not completely understood. However, stimulation at these frequencies likely elevates synaptic depression across the entire network

and therefore decreases the ability of the neurons to support longer bursts. Stimulation of 10 active channels (as determined by visual inspection) was repeated 600 times in randomized sequences at high frequency (20 Hz). Stimulations were biphasic 200 μsec pulses at 400 mV. Stimulation (after the first sequence) generated activity that typically ended in less than 30 ms (Fig. 4). Only data after the first second is used in the discrimination tasks for this protocol due to the burst of activity that is often elicited by the first stimulation(s).

3. Classification Methods

The input to the LCN was spatially coded via selection among the 8×8 grid of electrodes. Therefore, the goal of the classification task is to reconstruct the spatial stimulation pattern given the temporal structure of the spike trains recorded using the MEA. There are a wide variety of classification methods that could be used. For this paper, high similarity or small distance to a set of templates is used to classify windowed action potential trains. For similarity, a novel scheme based on cross-correlation of spike times is proposed and compared with a distance measure proposed by van Rossum (2001).

Since we are interested in proving the separation property, templates for the spatio-temporal pattern of action potentials with respect to each input are created, and later used to select the most similar template when an unknown pattern is given (nearest centroid classifier). The templates are built from simple smoothed peri-stimulus time histograms (PSTH, bin size equal to 1 sample) $S_{X \to Y}$ generated for every input-output mapping, i.e. between every stimulus channel X and every recording channel Y, as shown in equation (1) (Fig. 5).

$$S_{X \to Y}(t) = \frac{1}{N} \sum_k \sum_i \exp \left(\frac{-|t - t_{i,Y}^k|}{\tau} \right) \quad (1)$$

where $t_{i,Y}^k$ is the timing of the i-th action potential of k-th response contained in the training set of size N repeated stimulations corresponding to input X for output channel Y. Throughout this paper, the term training set is used to refer to the experimental data used to generate the templates for typifing stimulation of channel X. Training data is excluded from the descrimination task and is selected randomly from the entire set of experimental data unless otherwise noted. The smoothing function is a double exponential kernel with time constant $\tau = 0.25$ ms. τ is selected to be highly selective for differences in timing, but allow for spike jitter. Although the standard deviation of jitter in reliable spike

4

timing varies, most of the channels show precise timing rather than firing rate modulation (data not shown), thus we chose a fixed time constant for all action potentials.
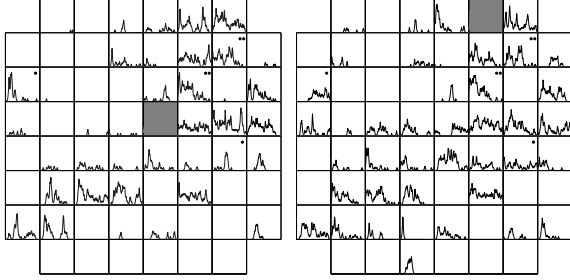

Fig. 5. Representative Example of Smoothed PSTHs for HFS protocol (N = 10 repeated stimulation of a single channel). Data shown represents the average response to stimulations on channel 54 (left) and 61 (right panel), indicated by grey. The horizontal axis corresponds to 3-18 ms interval after each stimulation input. The network's response to different stimulation sites is often unique producing different patterns dependent on stimulus location (response templates marked with a single dot at upper right corner). However, a few of the readout channels (marked with doubled dots) do show a remarkable similarity on a subset of stimulus locations, highlighting the importance of multichannel readout and classification techniques.

For the HFS protocol, we specifically excluded readout data from the stimulated channel. In order to do the same exclusion for the LFS protocol, we limited the number of stimulated channels to 10 where specified. Hence, only the nonstimulated channels were used for readout and classification in these cases. This is to ensure that the lack of detected action potentials due to the large isolated transient artifact on the stimulated electrode does not become a prominent feature used to match the template. For the LFS protocol with all 60 channels used as input stimulation, this exclusion is not possible and results in an artificial mechanism to identify the stimulated channel. However, incorporating 60 channels of data reduces the effect of this, as the significance of any one channel on average in the classifiers is approximately related to the inverse number of channels when the classifier equally weights the channels.

Spike trains recorded from many of the MEA channels are multi-unit activity (MUA). Spike sorting, to isolate individual neurons, is not used since it adds a source of error and it is not neccesary for the descrimination task. Futhermore, the correct method for combining (or weighting) the distances of action potential trains from multiple readout sources has not been determined, thus seperating each source would compound this issue further.

Template Normalization

Proper normalization of the PSTHs allows direct comparison to the template based on spike times regardless of the output spike train that is under classification. The normalization used is as follows,

$$\bar{S}_{X \rightarrow Y}(t) = \begin{cases} \dfrac{\mu_{X \rightarrow Y}}{\mu} & \text{if } \mu_{X \rightarrow Y} < \mu - \sigma \\ & \text{or} \quad \mu_{X \rightarrow Y} = 0, \\ \dfrac{S_{X \rightarrow Y}(t)}{\mu_{X \rightarrow Y}} & \text{if } \mu_{X \rightarrow Y} \geq \mu - \sigma. \end{cases} \quad (2)$$

where $\mu_{X \rightarrow Y}$ is the temporal average of $S_{X \rightarrow Y}(t)$ and μ and σ are the average and standard deviation, respectively, of $\mu_{X \rightarrow Y}$ for all X and Y.

Template matching can then be performed by the summation of the template values corresponding to the response (i.e., the timing of all the action potentials under classification).

$$Q_x = \sum_y \sum_j \bar{S}_{x,y}(t_{j,y}) \quad (3)$$

where t_j is the stimulation-locked time of the j-th action potential on output channel y. The response is classified as the channel x by having the maximum Q_x value.

L_2 Distance

As an alternate approach, L_2 distance between spike trains (van Rossum, 2001),

$$D_x = \sqrt{\sum_y \int \left(S_{x \rightarrow y}(t) - \lambda_y(t) \right)^2 dt} \quad (4)$$

is used where $\lambda_y(t)$ is the smoothed spike train with the same double exponential kernel similar to Eq. (1). The stimulation response is then classified by the minimum D_x value.

4. Results

We apply the template matching approaches from the previous section to LFS and HFS data with the goal of finding a one-to-one mapping of stimulus to response. Classification of inputs (stimulations) based on the spatio-temporal output pattern of action potentials is affected by a variety of factors. Here we note differences in the datasets and address the factors affecting input reconstruction including training data, stimulus protocols, windowing, and nonstationarity.

The spike trains produced by the two stimulation protocols differed in their composition, as seen by inspection of the activity (Fig. 4). Inspection of the spike trains under LFS showed unreliable responses as early as 13

Classification Method	LFS, 60 channel n = 7 cultures	LFS, 10 channel n = 7 cultures	HFS, 10 channel n = 3 cultures
Correlation	85.9%	95.4%	96.1%
L_2	64.5%	89.4%	98.1%

Table 1

Average (across cultures and repetitions) classifier performance in correctly identifying the input stimulation site. For LFS, 8 stimulations per input channel were classified using 2 responses as a template for each of the 60 channels (left column) or the 10 most active response channels (middle column) that improved classification performance. High frequency stimulation classifies 588 stimulations per input channel using 10 stimulations as a template. Overall, HFS resulted in more accurate seperability of the input-output relationship than the LFS protocol while also increasing the input bandwidth.

ms after stimulation and most reliable responses had terminated by 20 ms after stimulus. HFS produced the same reliable responses during the early phase; however, unreliable responses occurred less frequently. Occasionally, unreliable timings in responses were observed with HFS that was associated with the stimulation of particular channels. This response was typically isolated to a handful of readout channels with other channels simultaneously demonstrating precise timings. This contrasts with LFS which induced unreliable activity with varied firing rates as part of the population burst.

Input Reconstruction Performance

The performance of classifiers used in this work was dependent on the dataset, see Table 1. Sixty-channel classification is a more difficult task than classifying 10 channels due to increased input space. Further, LFS protocol occasionally delivers a stimulus mid-burst or does not evoke a (population) burst at all. However, this only accounted for 15 of the 600 stimulations on average in the cultures tested. Using data selection to circumvent these issues and choosing the ten most actively responding channels that remain compensates for both these issues and enables us to more fairly compare LFS data with HFS data. Regardless, both classifiers monotonically increased in performance with increases in the training data (N) for both LFS and HFS data (data not shown). Moreover, the normalization-based classifier outperformed the L_2 classifier when using fewer training data.

In contrast to the increase in performance associated with increases in training data, increasing the length of post-stimulus analysis window does not increase classification performance indefinitely. Indeed, the benefits of increases in window size on classification performance rapidly diminish following a peak at 5–10 ms and de-
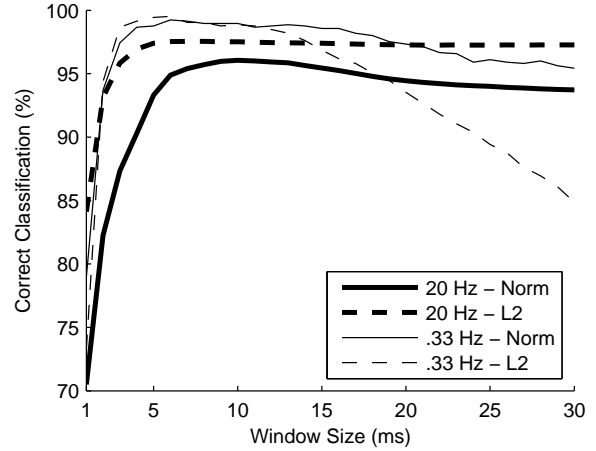

Fig. 6. Average Performance For Classifiers And Stimulation Protocols Dependent On Window Size. Analysis conditions are the same as in Figure 7 except the post-stimulus window always began at 3 ms and is variable. There is a sharper late phase decline of classification performance with the low frequency stimulation compared to high frequency stimulation reflecting the chaotic contribution of the late phase response.
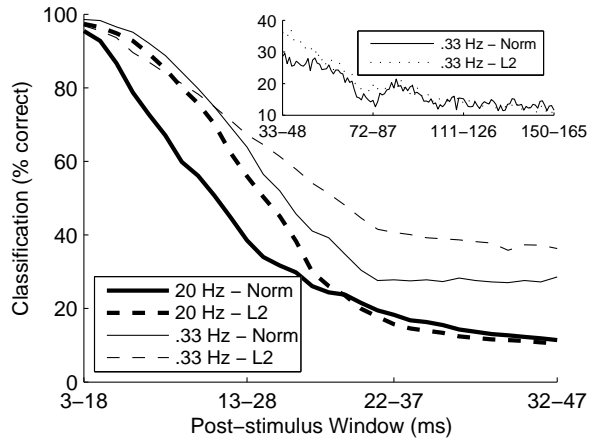

Fig. 7. Effect of Post-Stimulus Analysis Window on Classification Performance. Varying the post-stimulus window used for the template and classification demonstrates the reliability of timings relative to stimulation time. Five stimulations for each template were used to classify responses from 10 inputs (For LFS data only the most actively responding channels were used) during a 15 ms window. The degradation of the HFS performance is due to the lack of evoked late activity. The decline of low frequency stimulation performance is likely due to the small tau used in creating the smoothed PSTHs and the accumulation of spike jitter as larger taus have shown better similarity between trials in other work (Jimbo et al., 2000). Recurrent activity from stimulations inducing bursting (0.33 Hz) reduces the ability to perform correct classification.

cline gradually thereafter (see Figure 6). However, classification performance declined more slowly for HFS than for LFS. Shifting the template and comparison win-

dow toward later response data, we see a rapid decline in classification performance (Fig 7). We suggest that the deterioration of classification for the high frequency stimulation data with later post-stimulus windows is likely due to the lack of action potentials at longer delays (see Fig. 4). In contrast, the decrease in classification performance for the LFS protocol is likely due to accumulation of spike jitter and the resulting changes in action potential patterns across each stimulation. If we assume that recurrent activity especially during the late phase accumulates spike jitter (interpreted as error or noise in response output), then a decline would be expected in classification performance as recurrent activity increases (late phase) (Fig. 7). This would then result in an increase in ratio of highly recurrent activity to initial activity and reduce classification performance by effectively decreasing the signal to noise ratio (Fig. 6).

Surrogate Data

To insure the reconstruction of stimulated channels through output spikes is independent of artificial correlations, we shuffled the spike-times of each response while maintaining the number of action potentials for each channel for both training and test data. This surrogate data discards consistent stimulation locked correlations but maintains stimulation locked culture-wide average firing rates. The resulting classification performance of the 20 Hz stimulations fell to 15.8% and 45.3% for the normalization and L_2 methods, respectively.

Thus, the L_2 classifier's firing rate component was able to maintain marginal classification performance since the firing rate response varies by the channel stimulated. The normalization method performs near chance level, 10%, since firing rate information is removed in the formation of the template. This is confirmed in a second analysis by pairing random channels (instead of shuffling channels) with each spike time resulting in similar classification performance for both classifiers. However, assigning random spike-times to the original spike-channels resulted in classification rates for the normalization and L_2 methods of 10%. Not shuffling the data used to generate the template had little effect on the classification rates for the surrogate test data.

Training Set Sampling and Nonstationarity

Figure 8 shows a raster plot of each HFS trial for a single stimulation site and the template used for classification of a single output channel. Frequently, clear
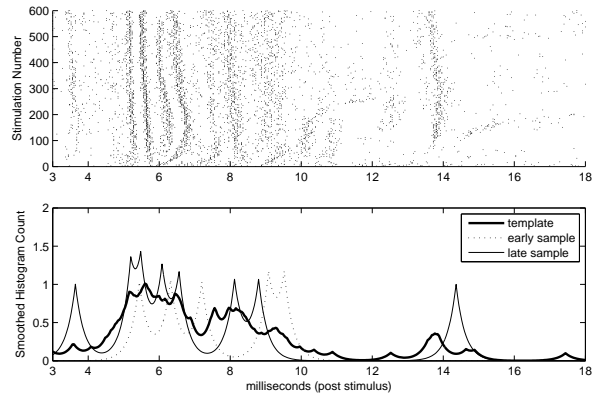

Fig. 8. Pattern of Evoked Activity and Associated Templates for One of the 60 Channels. Upper: Raster of activity from a single readout channel during repeated high frequency stimulation of an MEA electrode. Note the changes in action potential precise timings during the course of the 600 stimulations over 5 minutes. Lower: Comparison of the template and representative match from last 4 minutes relative to the first minute of stimulation. Note the blurring of the precise timings by randomly sampled template construction. (Thick line: smoothed template, Thin line: smoothed response from the last minute, Dotted line: smoothed response from the first minute).

temporally-dependent trends appear in the response over the course of the stimulation trials. The same trends were not noticed with the low frequency stimulus responses; although samples of the low frequency stimulation data are sparse and subtle changes would be indiscernible. Thus, only the stationarity of HFS data are mentioned and shown in Figure 8. A large portion (over 70% for each culture) of misclassifications occurred during the initial minute of high frequency stimulation; indeed, greater than 90% of misclassifications occurred in the first two minutes of data. As a result, the template is more representative of the latter 4 minutes as shown in Figure 8.

To illustrate the similarity of consecutive stimulus responses (and thus a gross measure of plastic changes on classification over long time periods), the ten immediately preceding responses of the same input stimulation are used to build the current template for classifying each succeeding response. The performance of the normalization classifier using this mutable template now becomes $99.77 \pm 0.07\%$ correct (n = 3). The changing template contains sharper peaks corresponding to reliable action potentials that shift over time and as a result are blurred in the static template (Fig. 8). Comparison with the numbers in Table 1 and considering the kernel size ($\tau = 0.25$ ms), the performance indicates that the changes in timings of action potentials are small between consecutive stimulations. Furthermore, this illustrates that nearly all the classification errors (in HFS

data) originate from the limitations of using the static template (or one that does not account for nonstationarity) and that only recent examples of stimulus response are required for correct classification while stimulation rate is maintained.

Changes in response to high frequency stimulation occur with many of the channels and are manifest as the origination or termination of precisely timed spikes or as a shift in the precise timing either earlier or later. Due to these changes, selecting the first stimulations for the generation of the template for comparisons results in more misclassifications than randomly selecting stimulations for the template. Hence, it is essential that the template incorporate stimulations from each period of experiments demonstrating differing timings and/or emphasize temporally relevant data extracted through HFS.

5. Discussion

The present study supports the notion of the hybrid computational device that consists of a biological system as a recurrent dynamical system, and a computer to explicitly interpret the biological activity based on the theory of LSMs. The ability to produce discriminable liquid states for different inputs, the separation property which is one of the key requirements of LSM, is demonstrated via the performance of a classifier. Despite the spontaneous non-stationary activity of the cultured network, we were able to show that the response is reproducible (reliable) for a short window (up to 15 ms) with high temporal precision ($\tau = 0.25$ ms) of the action potentials.

An MEA's electrode activates nearby tissue with a sufficiently low activation threshold – most of which are axons. Axons then propagate orthodromic and antidromic action potentials to the extent of the axonal arborization (Nowak and Bullier, 1998), thus, activating specific synapses with specific timing. In turn, a particular pattern is generated. However, spike jitter and unreliable vesicular release (among other reasons, Darbon et al., 2002; Reich et al., 1997; Shadlen and Newsome, 1998; Yvon et al., 2005) accumulates down stream error and is manifested as trial by trial variations in the response pattern. Without highly recurrent activity, these variations in timing propagate only for a short time. Thus, even a single early response is representative of the pattern seen from each repeated stimulation. This is particularly the case with respect to HFS as there is negligible late phase noise. Moreover, high input reconstruction accuracy (99.77%) is achievable by tracking

changes in the response spike trains.

High frequency input has improved reliability (minimal late phase with reduced recurrent activity), is better from an input data rate perspective (increases throughput), and is more likely to produce interacting responses (i.e., overlapping activity due to stimulation) than the more traditional slower stimulation rate. Rapid stimulation avoids a cascade of recurrent activation maintaining a large portion of action potentials due to stimulation relative to cascading recurrent activity (i.e., low noise versus high noise). However, high frequency input also reduces the activity length so much of the activity from consecutive stimulations may not directly interact through APs or be separated based on late portions since there is minimal activity after 30 ms. Yet, recovery time constants of synapses and neurons are slower than the period of HFS thus continuous activity may not be necessary to demonstrate fading memory.

However, the liquid is required to have the fading memory property which enables embedding of the temporal signal to the spatial representation (liquid states) and allow instantaneous readout for the implementation of a temporal signal filter (such as Wiener filter). Not only the stability of internal states and the dynamics of the liquid, but also the strength of the input stimulation limits the time period that the liquid is dependent on the past stimulus. In the HFS protocol, the recurrent dynamics is limited and activity is typically silent after 30 ms, therefore the state of the liquid can only be maintained through internal or unobserved states (Mongillo et al., 2008). We did observe some preliminary evidence of these states via conditional responses, however, the effects are currently unreliable.

To implement a full-fledged hybrid LSM, there are some important additional issues that needs to be resolved. One of the limitations of the current paradigm is the finite input space. Current stimulation spatially encodes each input to a single stimulation site, limiting the input space to 60 electrodes on the MEA. In the original LSM framework, the input is continuously fed into the network via a temporal pattern, however, technical limitations of the stimulation device prevents recording and stimulation at the same time. Therefore, a scheme that interleaves stimulation period and a response readout period is inevitable. Preliminary results suggest that using a combination of channels with a small time delay can still be separated. We are planning for extensive data collection in this direction.

References

Darbon, P., Scicluna, L., Tscherter, A., Streit, J., 2002. Mechanisms controlling bursting activity induced by disinhibition in spinal cord networks. Eur J Neurosci 15, 671–83.

DeMarse, T. B., Dockendorf, K. P., 2005. Adaptive flight control with living neuronal networks on microelectrode arrays. Vol. 3. pp. 1548–1551.

DeMarse, T. B., Wagenaar, D. A., Blau, A. W., Potter, S. M., 2001. The neurally controlled animat: biological brains acting with simulated bodies. Autonomous Robots 11 (5), 305–310.

Jimbo, Y., Kawana, A., Parodi, P., Torre, V., 2000. The dynamics of a neuronal culture of dissociated cortical neurons of neonatal rats. Biological Cybernetics 83, 1–20.

Jimbo, Y., Robinson, H. P. C., Kawana, A., 1998. Strengthening of synchronized activity by tetanic stimulation in cortical cultures: application of planar electrode arrays. Biomedical Engineering, IEEE Transactions on 45 (11), 1297–1304.

Jimbo, Y., Tateno, T., Robinson, H. P. C., 1999. Simultaneous induction of pathway-specific potentiation and depression in networks of cortical neurons. Biophysical Journal 76, 670–678.

Kamioka, H., Maeda, E., Jimbo, Y., Robinson, H. P. C., Kawana, A., 1996. Spontaneous periodic synchronized bursting during formation of mature patterns of connections in cortical cultures. Neuroscience Letters 206, 109–112.

Legenstein, R., Maass, W., 2007. Edge of chaos and prediction of computational performance for neural microcircuit models. Neural Networks, 323–333.

Maass, W., Markram, H., 2003. Temporal integration in recurrent microcircuits. In: Arbib, M. A. (Ed.), The Handbook of Brain Theory and Neural Networks, 2nd Edition. MIT Press (Cambridge), pp. 1159–1163.

Maass, W., Markram, H., 2004. On the computational power of recurrent circuits of spiking neurons. Journal of Computer and System Sciences 69 (4), 593–616.

Maass, W., Natschläger, T., Markram, H., 2002. Real-time computing without stable states: A new framework for neural computation based on perturbations. Neural Computation 14 (11), 2531–2560.

Maass, W., Natschläger, T., Markram, H., 2003. A model for real-time computation in generic neural microcircuits. In: Becker, S., Thrun, S., Obermayer, K. (Eds.), Proc. of NIPS 2002, Advances in Neural Information Processing Systems. Vol. 15. MIT Press, pp. 229–236.

Madhavan, R., Chao, Z. C., Potter, S. M., 2007. Plasticity of recurring spatiotemporal activity patterns in cortical networks. Physical Biology 4, 181–193.

Maeda, E., Kuroda, Y., Robinson, H. P. C., Kawana, A., 1998. Modification of parallel activity elicited by propagating bursts in developing networks of rat cortical neurones. European Journal of Neuroscience 10 (2), 488–496.

Marom, S., Shahaf, G., February 2002. Development, learning and memory in large random networks of cortical neurons: lessons beyond anatomy. Q Rev Biophys 35 (1), 63–87.

Mongillo, G., Barak, O., Tsodyks, M., March 2008. Synaptic theory of working memory. Science 319 (5869), 1543–1546.

Nowak, L. G., Bullier, J., 1998. Axons, but not cell bodies, are activated by electrical stimulation in cortical gray matter: I. Evidence from chronaxie measurements. Experiments in Brain Research 118, 477–488.

Parodi, P., Jimbo, Y., Kawana, A., Macri', D., Torre, V., November 1998. Segmentation of the response of a neuronal network into clusters with similar activity. Biosystems 48 (1-3), 171–178.

Potter, S. M., DeMarse, T. B., 2001. A new approach to neural cell culture for long-term studies. Journal of Neuroscience Methods 110, 17–24.

Reich, D. S., Victor, J. D., Knight, B. W., Ozaki, T., Kaplan, E., 1997. Response variability and timing precision of neuronal spike trains in vivo. J Neurophysiol 77, 2836–41.

Rolston, J., Wagenaar, D., Potter, S., 2007. Precisely timed spatiotemporal patterns of neural activity in dissociated cortical cultures. Neuroscience 148 (1), 294–303.

Ruaro, M. E., Bonifazi, P., Torre, V., March 2005. Toward the neurocomputer: Image processing and pattern recognition with neuronal cultures. IEEE Transactions on Biomedical Engineering 52 (3), 371–383.

Sandberg, I.W.; Xu, L., June 1997. Uniform approximation of multidimensional myopic maps. IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications 44 (6), 477–500.

Segev, R., Baruchi, I., Hulata, E., Ben-Jacob, E., 2004. Hidden neuronal correlations in cultured networks. Physical Review Letters 92 (11).

Shadlen, M. N., Newsome, W. T., May 1998. The variable discharge of cortical neurons: Implications for connectivity, computation, and information coding. Journal of Neuroscience 18 (10), 3870–3895.

Shahaf, G., Marom, S., 2001. Learning in networks of cortical neurons. J Neurosci 21, 8782–8.

Tateno, T., Jimbo, Y., January 1999. Activity-dependent enhancement in the reliability of correlated spike timings in cultured cortical neurons. Biological Cybernetics 80 (1), 45–55.

van Pelt, J., Wolters, P. S., Corner, M. A., Rutten, W. L. C., Ramakers, G. J. A., 2004. Long-term characterization of firing dynamics of spontaneous bursts in cultured neural networks. IEEE Trans Biomed Eng 51, 2051–62.

van Rossum, M. C. W., 2001. A novel spike distance. Neural Computation 13 (4), 751–764.

Wagenaar, D. A., DeMarse, T. B., Potter, S. M., 2005a. Meabench: A toolset for multi-electrode data acquisition and on-line analysis. In: Proc 2nd Intl IEEE EMBS Conf on Neural Eng. pp. 518–521.

Wagenaar, D. A., Madhavan, R., Pine, J., Potter, S. M., 2005b. Controlling bursting in cortical cultures with closed-loop multi-electrode stimulation. J Neurosci 25, 680–8.

Wagenaar, D. A., Nadasdy, Z., Potter, S. M., 2006. Persistent dynamic attractors in activity patterns of cultured neuronal networks. Physical Review E 73 (5).

Yvon, C., Rubli, R., Streit, J., 2005. Patterns of spontaneous activity in unstructured and minimally structured spinal networks in culture. Exp Brain Res 165, 139–51.