

Revealing ensemble state transition patterns in multi-electrode neuronal recordings using hidden Markov models

Article

Accepted Version

Xydas, D., Downes, J. H., Spencer, M. C., Hammond, M. W., Nasuto, S. J. ORCID: <https://orcid.org/0000-0001-9414-9049>, Whalley, B. J., Becerra, V. M. and Warwick, K. (2011) Revealing ensemble state transition patterns in multi-electrode neuronal recordings using hidden Markov models. IEEE Transactions on Neural Systems and Rehabilitation Engineering, 19 (4). pp. 345-355. ISSN 1534-4320 doi: 10.1109/TNSRE.2011.2157360 Available at <https://centaur.reading.ac.uk/23197/>

It is advisable to refer to the publisher's version if you intend to cite from the work. See [Guidance on citing](#).

Published version at: <http://dx.doi.org/10.1109/TNSRE.2011.2157360>

To link to this article DOI: <http://dx.doi.org/10.1109/TNSRE.2011.2157360>

Publisher: IEEE

All outputs in CentAUR are protected by Intellectual Property Rights law, including copyright law. Copyright and IPR is retained by the creators or other copyright holders. Terms and conditions for use of this material are defined in the [End User Agreement](#).

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

Revealing Ensemble State Transition Patterns in Multi-Electrode Neuronal Recordings Using Hidden Markov Models

Dimitris Xydas, Julia H. Downes, Matthew C. Spencer, Mark W. Hammond, Slawomir J. Nasuto, Benjamin J. Whalley, Victor M. Becerra, *Senior Member, IEEE* and Kevin Warwick

Abstract—In order to harness the computational capacity of dissociated cultured neuronal networks, it is necessary to understand neuronal dynamics and connectivity on a mesoscopic scale. To this end, this paper uncovers dynamic spatiotemporal patterns emerging from electrically stimulated neuronal cultures using hidden Markov models (HMMs) to characterize multi-channel spike trains as a progression of patterns of underlying states of neuronal activity. However, experimentation aimed at optimal choice of parameters for such models is essential and results are reported in detail. Results derived from ensemble neuronal data revealed highly repeatable patterns of state transitions in the order of milliseconds in response to probing stimuli.

Index Terms—Cultured neuronal networks, hidden Markov models, neuronal state transitions, multi-channel recordings.

I. INTRODUCTION

GAINING a deeper understanding of neuronal information processing is a requirement not only for developing applications and devices at the nervous system-hardware interface, but also for understanding mechanisms of learning and memory at the level of network connectivity. Prior evidence [1]–[5] suggests that distinct meta-stable states and state transitions can be identified within the activity of individual neurons and neuronal networks and that such states may be associated with higher cognitive processes [1], [4]. Such meta-stable states are characterized by relatively robust (but not completely invariant) repetition of similar discrete patterns of activity which persist for some time and are otherwise embedded in a continuous flow of activity. Although the underlying biological mechanisms leading to the development of such meta-stable states embedded in continuous dynamics may be different in different brain areas, nevertheless they seem to be a common organizational principle around which various neural system organize their information processing *in vivo*.

These studies primarily involve *in vivo* experiments (or *in vitro* hippocampal slice preparations [2]) in monkeys and

rats, hence analyzing state transitions in neuronal cultures is a new and developing field which further motivates this study. Understanding neuronal activity at a network-level will help bridge the gap between lower-level functions and whole-brain studies, a necessary prerequisite for the success of any future advanced BMI procedures. These cultured neuronal networks may not have the structural integrity of their *in vivo* counterparts, but they still share many common properties with *in vivo* networks [6]. Exploring the input/output relationships of such cultures in a controllable environment, with closed-loop sensory feedback, will provide useful insight into mechanisms which process sensory information and hence will be beneficial to advanced neuro-prosthetics.

Cultured neuronal networks provide an opportunity to analyze and probe neuronal function at many levels [7]–[12]. In engineering applications, systems designed to exploit cultured neuronal networks, such as biological-machine hybrids [13], [14], often assume that a culture acts as a fixed, albeit unknown, transfer function. However, whilst inducing a particular state in culture prior to stimulation has been shown to improve its responsiveness to input stimuli [15], such an enforced state change may have unforeseen consequences as the resulting state changes may persist for extended periods [8], [16]. Hence, in order to achieve reliable and predictable responses, culture states prior to and during stimulus delivery must be considered.

Given the importance of meta-stable states for information processing *in vivo* and the relative ubiquity of biological mechanisms supporting them, it is plausible that such mechanisms are still present in cultures and hence the latter have the capacity to organize their activity around the same principle. Thus, discovery of such meta-stable states in *in vitro* cultures would be very important for characterization of their activity and could help in probing their properties and more efficient manipulation.

Here, inspired by the characterization of the meta-stable states in *in vivo* studies, results are reported from the use of HMMs to parameterize and identify internal dynamics of neuronal cultures in terms of delineation of meta-stable states and state transition sequence patterns. However, such parameterization initially raises a set of questions which must first be answered. Firstly, a compact scheme must be chosen to represent the highly dimensional input data. Secondly, HMMs represent complex dynamic systems using few parameters, whose configuration relies initially on the chosen

Manuscript received April 15, 2010; revised October 29, 2010 and March 18, 2011; accepted May 10, 2011. This work was funded by the Engineering and Physical Sciences Research Council (EPSRC) under grant No. EP/D080134/1 and the Royal Society.

D. Xydas, J. H. Downes, M. C. Spencer, M. W. Hammond, S. J. Nasuto, V. M. Becerra and K. Warwick are with the Cybernetics Research Group, School of Systems Engineering, University of Reading, Reading, Whiteknights, RG6 6AY, UK (e-mail: {d.xydas, j.downes, matthew.spencer, m.w.hammond}@pgr.reading.ac.uk; {s.j.nasuto, v.m.becerra, k.warwick}@reading.ac.uk).

M. W. Hammond and B. J. Whalley are with the School of Pharmacy, University of Reading, Reading, Whiteknights, RG6 6AP, UK. (e-mail: m.w.hammond@pgr.reading.ac.uk; b.j.whalley@reading.ac.uk).

input representation scheme. Thirdly, HMMs may or may not be robust to certain initializations [17] depending on the chosen configuration. Related research with similar dynamic networks (*in vivo*) influenced the choice of observation scheme and parameter initialization [1], [4], [5] and our methods and results are discussed in detail.

A. Hidden Markov models

HMMs provide a natural framework for capturing, from observed quantities, the underlying hidden dynamics of a system on the assumption that it can be represented by a discrete number of states and transitions between them. These states are representations of combinations of multi-unit spiking activity. HMMs are capable of uncovering the evolution of the system dynamics in the form of the most likely state sequence consistent with a particular observation sequence.

The training procedures used in HMM parameter estimation are mathematically well founded [17], thus making HMMs particularly attractive for analysis of sequential data types such as spike trains which exhibit complex temporal variability. HMMs have been successfully applied to both single [18]–[20] and ensemble neuronal recordings [21]–[23] in addition to a wide range of research fields such as speech modeling [17], image processing [24] and DNA sequencing [25].

Several *in vivo* studies suggest that the meta-stable patterns of activity on a fixed population of neurons or recording units play important role in information processing in neural systems. For example, such meta-stable state sequences could directly correspond to baseline-planning-execution phases of a motor control task as characterized in [23], or the modes of neural oscillations (continuous-mixed-bursting) induced by application of receptor agonists in hippocampal slices as reported in [18]. In other cases, drawing direct relations to biological properties is a harder task, particularly if a cue-reward reference system is absent. The identification of a meta-stable state may be straightforward in cases such as the characterization of pacemaker *vs* bursty activity in a single neuron, but more intricate for other cases of mixed activity patterns [19]. The identification of putative meta-stable states in *in vitro* activity is intricate as they may be embedded in, otherwise patternless, overall decay of activity characteristic for a response of a dissipative system to a transient onset of activity (intrinsic or extrinsically induced). Model selection based on minimization of known information criteria [19], [26] is a way of reducing complexity in terms of overestimating the number of states. HMM analyses of *in vivo* ensemble neuronal activity [1], [4] have revealed that neuronal responses to stimuli can be represented by a robust transition pattern through state sequences and that, whilst the timing of state transitions may be variable, the sequence is stimulus specific and consistent from trial to trial. Moreover, HMMs have enabled analyses of neuronal activity that was not locked in time to external events [4]. We aim to discover whether such patterns could be evident *in vitro* as well, which motivated the investigation reported here: are appropriately configured HMMs capable of parameterizing the dynamics of neuronal cultures under external excitation, a desired characteristic for more robust culture control?

This paper presents work on the use of HMMs to accurately classify and predict network state sequences using 59 channel multi-electrode array (MEA) data from cortical cultures. The main unit of interest in such data streams is a 'spike', representative of a single or compound action potential elicited by a given neuron or neurons. Such spikes are typically detected using threshold crossing algorithms and the resulting spike timestamp and channel number stored (the use of 'electrode' and 'channel' will be interchangeable henceforth). More detail is provided in section II-A.

The remainder of this paper is structured as follows: Section II (Methods) introduces the HMM considered and describes essential parameterization tests which are required, whilst section III (Results) shows results of these tests as well as how the studied neuronal data follows certain robust sequences of hidden states. Sections IV and V (Discussion and Conclusions) interpret the results, summarize the main findings of the paper and consider potential extensions of the work. Some basic background knowledge on HMMs is assumed [17].

II. METHODS

A. Real data choice and analysis

Analysis of neuronal data was performed on publicly available (under the GNU General Public License v2) data sets [27] originally used to present and classify the rich developmental repertoire of burst characteristics exhibited by cortical cultures. The datasets were obtained from dense cultures (~50,000 neurons plated), using standard 8x8 planar MEAs (providing 59 recording/stimulating electrodes in total), at a sampling frequency of 25 kHz. Novel stimulus artifact suppression by local curve fitting meant that spike detection was possible within 2 ms after stimulation on most electrodes (the stimulated electrode usually remained saturated for 50-100 ms) [28]. Spikes were detected using an adaptive filter algorithm [29] which adapts rapidly to changing noise situations, without desensitizing during bursts of spikes. Data were acquired following a stimulus (biphasic ± 0.8 V; 0.4 ms per phase) delivered cyclically on each electrode in turn at 300 ms intervals, for a total of 50 stimulations at each electrode. For more detailed and specific information describing experimental methods see [27]. Stimulation on a given electrode will be referred to as a separate 'condition' and the 50 stimulations each condition comprises will be referred to as 'trials'. Each culture is seeded onto a separate MEA and each recording is referenced with an ID, e.g. ID '3-4-20' indicates the culture code was part of batch 3, seeded on MEA number 4 and recorded at 20 days *in vitro* (DIV). The distinction is necessary since rat cells from a single batch are seeded on multiple MEAs and recorded for several weeks. Each culture develops its own specific connectivity and develops over a period of weeks hence recordings at different DIV may exhibit different behavior. These data sets are well documented, comprise a large number of previously analyzed experiments and are highly comparable to data acquired from our own experimental system, allowing easy application of developed tools to our own data.

Pre-processing analysis of spike timings and spike waveforms (extracellular voltage traces sampled 0.96 ms before

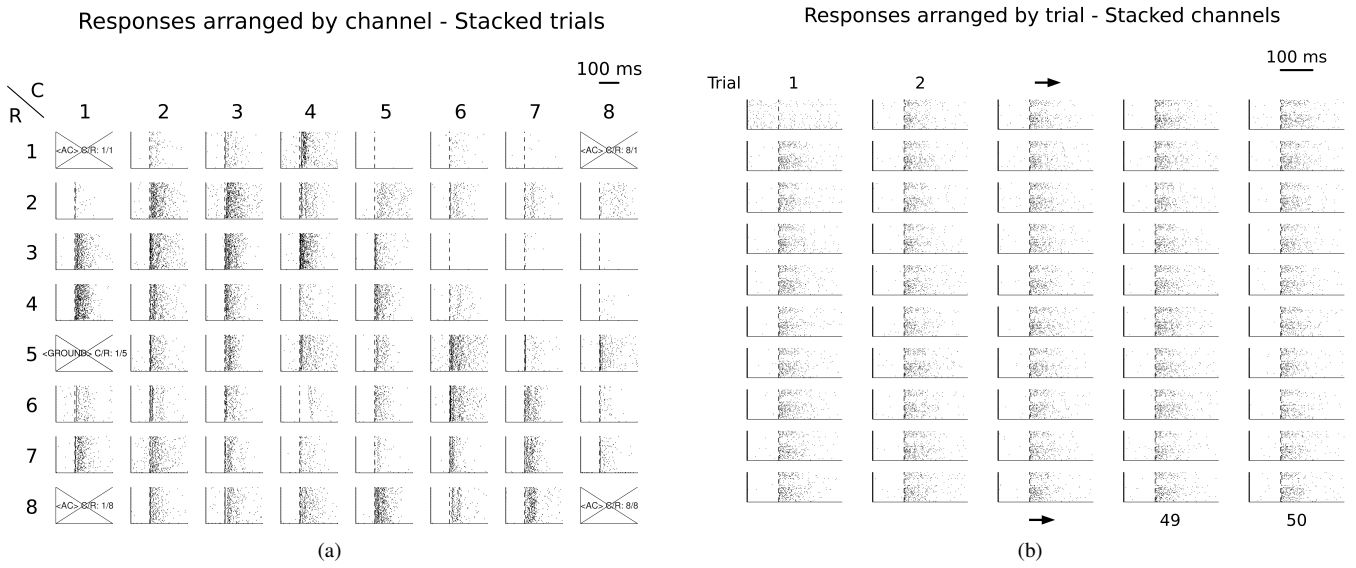


Fig. 1. Raster plots of example condition with 50 trials, in two separate groupings. **a)** Raster plots with channels arranged in 8x8 column/row format as per physical MEA layout. For each channel, individual trial response data is arranged around 100 ms pre-stimulus to 200 ms post-stimulus time duration (dashed vertical line after first 100 ms indicates the time of stimulation) and stacked (bottom to top) in chronological order for the 50 available trials. **b)** Raster plots arranged in a 5x10 format corresponding to the 50 available trials (left-to-right, top-to-bottom arrangement) based on chronological order. Channels are stacked (bottom to top) based on their hardware IDs (see section II-B).

and 1.96 ms after the spike's maximal height) revealed the existence of spike waveforms with both negative and positive polarities. Spikes with positive polarity constituted on average 39% ($STD = 6$) of all spikes (out of 40 cultures at DIV 19-21). Additionally, on average, 44% ($STD = 10$) of those positive spikes appeared in apposition to some negative spikes, appearing immediately after the latter (inter-spike interval < 2 ms). Since the electrodes record fluctuations of extracellular membrane potential from multi-unit sources (neuron groups), action potential activity typically manifests as a negative deflection of voltage. Considering detected but relatively small amplitude positive deflections as likely to be noise, artifacts or the after-hyperpolarization following negative going action potential events, it was deemed reasonable to dismiss them. It is highly unlikely that positive only spikes with a large enough amplitude would be detected, because in order for these to occur, the electrodes would have to create a weak perforated patch-clamp connection around the cell body itself. No spike-sorting technique was applied due to known drawbacks of standard techniques when dealing with bursts of neuronal activity, but novel techniques may be considered in the future.

Our results focus on activity changes induced by each condition. Consequently, the analysis uses data windows including small pre- and post-stimulation periods of 100 ms and 200 ms respectively (Fig. 1). We selected 8 cultures recorded at 19-21 DIV and chose several stimulated electrodes from each culture, which were capable of producing array-wide responses post-stimulus. A total of 60 HMMs were generated for these different conditions. Of particular interest are state 'segments', i.e. *sojourn state times* (time windows in which the state remained the same for a number of HMM steps) and segment 'progressions', i.e. *sequences of the most frequently occurring state segments throughout trials*.

B. Definition of observations

A crucial step prior to training an HMM is to determine how the system's observations should be expressed as input to HMM training procedures. Two principal options are to represent data as selections from either a finite set of discrete observables or from a continuous probability distribution. Discrete observables are assigned labels (usually from a finite set, such as integers) and continuous observables can be any real numbers. In the case of discrete HMMs as used here, each state has an associated conditional discrete probability (a single row in the observables probability matrix; see appendix for basic structure definitions).

The method for defining observations was inspired by Seidemann et al. [4]. In the original approach, data was time-binned into non-overlapping 1 ms windows and at each step a label was assigned corresponding to the neural unit which exhibited a spike, thus the number of observables corresponded to the number of recorded units. If spikes 'overlapped' in the same bin (e.g. two or more units showed activity, or a single unit showed activity more than once), a single label was assigned randomly. Seidemann and others [3] found this to be a rare event (usually $< 5\%$) for their *in vivo* activity. In contrast, the spike statistics in cultures are different because of epileptiform activity and lack of defined structure. This is illustrated by the following overlap statistics for a representative culture recording. For time bins representing 5, 2, 1, and 0.1 ms durations respectively, 25.5%, 16%, 10.8% and 0.8% of bins contained spike overlaps and only 20.3%, 13.8%, 10.6% and 4% of bins registered single spikes (with the remaining 54%, 70.1%, 78.6% and 95.3% of bins registering no spikes). Hence, the original technique was extended by explicitly handling spike 'overlaps'. This observation also motivates the choice of 0.1 ms rather than 1 ms binning in subsequent modeling. Note that the above data also indicates the tendency of responses to

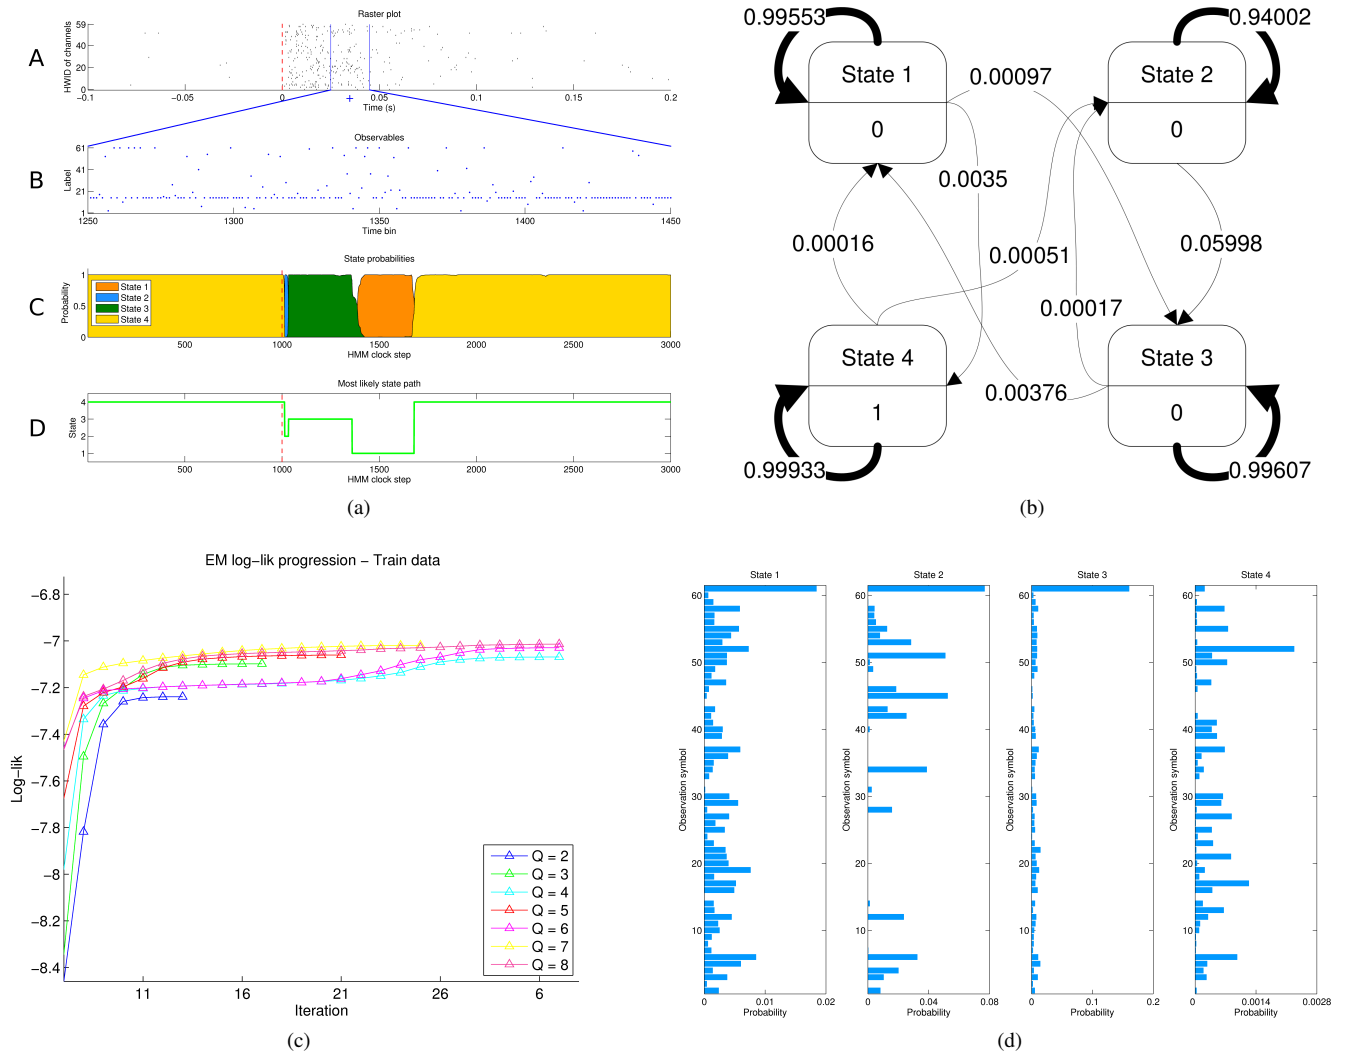


Fig. 2. Results from a trained HMM on array-wide responses to single-electrode stimulus. **a)** Results from a representative trial; Panel A displays the data sequence raster plot (300 ms total time), panel B illustrates results of the time binning and pre-processing of spike trains into observable labels (zoomed-in view of 200 HMM steps — each step corresponds to a 0.1 ms time window). Note that for every time step there is only one label assigned, also note the high occurrence of NSS (label 15) when firing rates are low (many 'empty' bins) and high occurrence of UCS (label 61) when firing rates are high (spike overlaps registered). Panel C is a visualization of state probabilities (forward-backward algorithm). Timings in panels are all vertically aligned. **b)** State-flow diagram representation of state transitions. The sojourn probabilities are clearly much higher than the rest. The HMM is fully-connected, but transitions with very low probabilities ($<10^{-5}$) are hidden for clarity. **c)** Progress of data log-likelihood throughout EM iterations for the best-performing models, with differing numbers of states (2-8). Although increasing the number of states generally increases the log-likelihood, the chosen information criteria will penalize the addition of more free parameters (see later section III-B). **d)** Illustration of the observable symbol probability distribution; NSS frequencies hidden for clarity because they are orders of magnitude higher than those of the other labels. Columns correspond to states, and rows of each subplot correspond to the observables.

include many 'empty' bins interspersed with bursts of activity.

Discrete observables were thus generated in a way similar to above, where a label corresponded to the channel on which a spike occurred in a given time bin (59 labels). A 'no-spike symbol' (NSS) label was also included if no spike was registered in a bin (by making use of the reference ground channel, label 15, which, by definition, acquires no neuronal signal). Extra labels were also initially introduced for special cases which included spike overlaps (simultaneous spikes on several channels). Rarely occurring spike overlaps were allocated a single 'undefined collision symbol' (UCS). The remainder of this paragraph describes how the extra labels for spike overlaps were produced. First, a new symbol was created from labels of active channels concatenated in ascending order.

For example, if channels 4, 6 and 17 showed activity during the same bin, the new symbol label would be '040617'. Labels re-occurring frequently enough, above a certain threshold (as a percentage of total spikes), would be retained, otherwise they would be replaced with the UCS label. Every time a new symbol was accepted, the number of observables, M , increased by 1. In practice however, identical spike overlaps were rarely repeated, resulting in all being assigned the UCS symbol unless the threshold was set to a very low value ($<0.05\%$ of all symbols). Therefore, the *final* observable set comprises 61 labels in total.

C. Model selection and training

Model selection involved three main considerations: initial conditions for the state transition matrix, choice of number of states and by consequence controlling model overcomplexity.

In previous HMM applications to neuronal data, it has been suggested that the state transition matrix (A matrix) should be initialized as an almost diagonal matrix since state transitions at millisecond intervals are highly unlikely to occur [1], [4]. Hence, in order to test the effects of initialization on the A matrix, a number of tests were performed on a specific condition. The A matrix was initialized either by setting the diagonal A_{0ii} equal to 0.99, 0.98, 0.8, 0.5, or by choosing elements in each row randomly from the unit interval ($A_0 = \text{Rnd}$, ascertaining that each row forms a proper conditional probability distribution, i.e. each row's sum equals to 1). The average value of duration d in state i is calculated as: $d_i = \frac{1}{1 - A_{ii}}$, hence the above values correspond to initializing mean sojourn times as either: 10 ms, 5 ms, 0.5 ms, 0.2 ms or random respectively. On the other hand, initial probability distributions (π vector) were randomly initialized. Also, no prior knowledge of channel firing frequencies within states was assumed, so the observable symbol distribution probabilities conditioned on states (B matrix) were also initialized randomly for every state [1]. Ultimately, the initialization of both π and B had little impact on the outcome of the training procedure.

In the case of HMMs and state-based modeling generally, the number of states has to be pre-determined. These states — combinations of multi-unit firing activity — are dependant on the definition of the observables and the timescales involved. Higher number of states lead to higher model dimensionality, so the number of states is typically determined by trial and error [1], [5]. Instead, we attempt to improve the model selection process by utilizing two standard information criteria for model selection, which strike a balance between capturing the regularities within data and model complexity. The results of these investigations are presented in section III.

For the training procedure, an EM algorithm (Baum-Welch method) was used to find the Maximum Likelihood Estimates of the HMM parameters [17]. The data sets consisted of 50 trials overall, which were partitioned into 5 sets of 10 trials. 5-fold cross-validation was performed by retaining one set as validation data and using the remaining sets as training data. The best model was considered the one producing the highest (validation data) log-likelihood out of the 5 folds, however a mean (validation data) log-likelihood was also obtained by averaging over the log-likelihoods of all folds. This value is a useful measure of goodness of fit. In order to characterize neuronal culture dynamics, the most likely state transition paths occurring over an entire observation sequence were estimated using the Viterbi algorithm (see appendix).

III. RESULTS

A. Training results

Before presenting results on parameter initialization (Section II-C), the results of a typical training procedure are shown

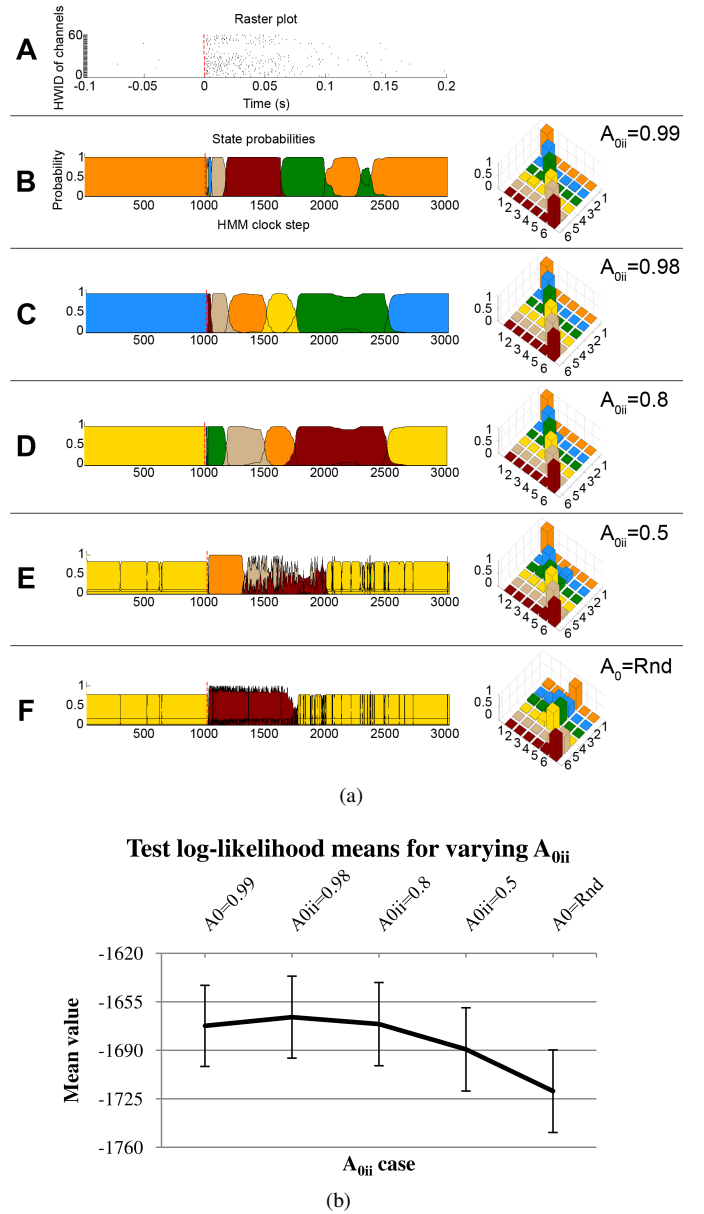


Fig. 3. Investigating the effects of A matrix initialization. **a)** Panel A depicts an example raster plot (same time windows as in previous figures), subsequent panels depict state and A matrix results after training, with A matrix diagonals initialized as 0.99, 0.98, 0.8, 0.5, or entire matrix initialized randomly, respectively. The first column shows state probabilities (forward-backward algorithm), while the second column shows 3D representations of the A matrix after training. Note that colors between panels are unrelated as they correspond to states of separate models. Log-likelihoods for multiple (training) trials converged to similar values after training for highly diagonal cases. Although in some cases state segmentation may appear different, in fact the underlying observable symbol distribution is very similar. For example, panel B state 6 and panel C states 1 and 4 share such similar distributions. Reducing the diagonal to $A_{0ii} = 0.5$ permitted maintenance of the highly diagonal structure after training to some degree, although state separation was unclear at various points, in a similar way to random initialization. **b)** Mean values for log-likelihoods of the 10 test trials. Values were around -1680 for the highly diagonally initialized cases and slightly lower (-1719) for the random case.

(Fig. 2) to serve as a guide of how subsequent data will be presented. For a typical trial, the sojourn probabilities are much higher than probabilities of transitions to different states (Fig. 2b), as governed by state matrix initialization (see later),

but the results also indicate that sharp transitions between certain states are possible (e.g. a few ms immediately after stimulation). In the illustrated example, the HMM revealed that the network response — following stimulation — visited four different states before settling back in its pre-stimulation state (Fig. 2a, state 4). Notably, in most of the analyses, the pre-stimulation state was almost always associated with very sparse neuronal activity (typically <10 spikes total) and was revisited at a period of time after stimulation.

B. Transition matrix comparisons

Interestingly, similar to methods reported by other groups [1], [4], with regards to the A matrix initialization, the HMMs settled to clearly defined state changes and high sojourn probabilities since, for the higher diagonal initialization cases ($A_{0ii} \geq 0.8$), all final diagonal values of the A matrix remained high ($A_{0ii} > 0.89$). It was found that starting with sojourn probabilities ≥ 0.8 always produced similar state transitions regardless of π and B initialization, although it may appear that similar initializations produce different state segmentations (Fig. 3a). Certain observable symbol distributions resulting from states between different models were found to be very similar; hence a single state segment in one model may appear split into two or more similar segments in another. This is not a large issue in models employing similar number of states, but may lead to ‘unnecessary’ states as will be seen next when increasing said number. Test trial log-likelihoods resulting from diagonal initialization were similar between cases and all higher than for the random case. However, as the diagonal was reduced further ($A_{0ii} = 0.5$) the log-likelihood also reduced and state separation became less clear. The reason for finalizing the choice of initialization of A to $A_{0ii} = 0.98$ over even higher values was due to the fact that the mean log-likelihood on unseen test data (Fig. 3b) reached higher values for this choice and the fact that initializations in the region of 0.8-0.98 typically converged to very similar state transition patterns.

The subsequent tests investigated the effects of varying the number of states (Fig. 4a), thus creating N -state models, with $N = 2, \dots, 8$. In initial tests, adding states had the effect of increasing the number of state transitions during the initial ~ 50 ms period of heightened activity following stimulation, while periods of quiescence prior to stimulation were rarely divided into more states. Furthermore, when increasing N , the extra added states appeared only for a small number of steps if at all. Generally, mean log-likelihood values increased when increasing number of states. However, care should be taken as increasing the number of states will eventually result in an over-fitted and unusable model. Reports in the literature [4], [5] suggest that recorded neuronal activity is describable by a relatively limited number of states. In order to robustly select the number of states which best described the fitted data (for each model) and avoided overfitting, two common criteria for model selection were used [30], Akaike’s information criterion (AIC) and the Bayesian information criterion (BIC):

$$AIC = -2\log L + 2f \quad (1)$$

$$BIC = -2\log L + f\log S \quad (2)$$

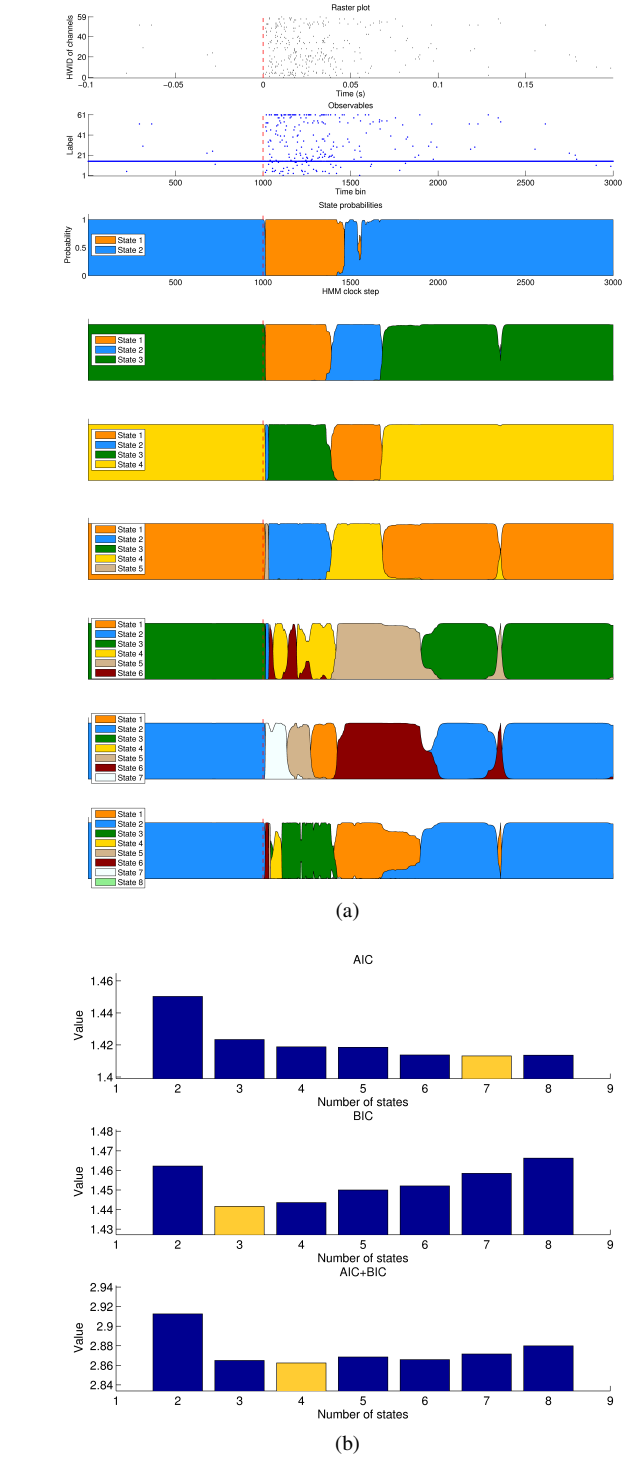


Fig. 4. Example models with varying number of states. **a)** A two-state model simply divided the periods of quiescence from heightened activity in the midst of the trial. Pre-stimulus quiescent activity periods were always represented by a single segment. More than 5 states never or rarely appeared, so addition of extra states did not provide informative segmentation, but simply added to model complexity. This is well reflected in the AIC and BIC values. **b)** Results from application of chosen information criteria for an example model. The combined values of AIC and BIC are shown in the third panel. AIC penalizes the number of free parameters less strongly than BIC, hence the AIC curve fluctuates as more states are added, but the BIC forms a basin at around 3-4 states (yellow bars indicates smallest/best values).

where $\log(L)$ is the maximized log-likelihood of the fitted model, f is the number of free parameters and S is the number




Fig. 5. Example state results from 8 different conditions (4 cultures, 2 conditions from each), with two sample trials of each condition; time windows are the same as in previous figures. Culture number and DIV information is represented in the ID (e.g. ID 2-2-19 is culture number 2-2, DIV 19), stimulated channel numbers are also displayed and trial number is indicated at the right of each resulting plot pair (raster and state probability plots). Again, color coding is arbitrary and different for each condition. Results show that channel activity is usually limited before stimulation, stimulation typically elicits 3-4 robust state segments of variable timing and network response returns to initial state after stimulation.

of observations used in fitting the model. In all cases $S = 40 \times 3000$ (trials \times observations) were used. The number of free parameters in each model was calculated as in Eq. 3.

$$f = N(N - 1) + N(M - 1) + (N - 1) \quad (3)$$

N is the number of states in the model and M is the number of observation symbols. The three terms in Eq. 3 correspond to the contribution from the state transition matrix, the observation symbol matrix and the initial state matrix respectively. Since the rows of each contributing matrix/vector add up to one (probability distributions), one value of each row is not freely estimated. It can be noticed that in both criteria the first term, which consists of the model likelihood, is the same. The second term, which reflects parsimony, is slightly different and penalizes the result when more free parameters are added to the model. The AIC penalizes parameters less strongly than the BIC. The model which produced the minimum value of AIC and BIC summed was chosen as the best representative model of the given stimulated electrode, similar to [26]. These criteria provide a robust method of model selection in our analysis. Our search space was not increased above 8 states, since other cultures also indicated comparative results for the information criteria so it was clear that including more than 8 states would be redundant.

C. Neuronal data and state progression analysis

Final results focus on analysis of the neuronal data and their segment progression statistics. For each trial in every condition, the progression was identified up to a depth of ten segments. Progression statistics were calculated for each condition as the frequencies of the most likely segments in the progression. A one-way analysis of variance (ANOVA) with Tukey's honestly significant difference (HSD) post-hoc test on

progression frequencies revealed a difference between initial and last parts of the progressions (Fig. 6). Specifically, the first four segments' frequencies were significantly higher (each $p < 0.05$ by multiple comparisons) from the last six, but not from one another. Hence, typically four state segments after stimulation well characterize the response of the majority of these network ensembles.

Furthermore, the following observations can be made (Fig.5): Responses will typically traverse through three or four (e.g. 2-2-19) state segments after stimulation. The first segment was typically of very short duration (1-10 ms) and corresponded to the initial burst of stimulus-induced activity. The response then moved through further longer duration sojourn states before returning to the same low activity, pre-stimulus segment. There were also notable cases where the HMM captured periodic activity from one or more channels (e.g. 6-1-19, channel 54).

IV. DISCUSSION

The model structure developed and tested was applied to analyze real neuronal network responses to probing stimuli. Using a discrete HMM utilizing active channels to construct observation symbols, individual HMMs were trained based on stimulated channel responses, while initialization of important parameters and their performance was assessed.

The reasonable assumption that neural units are unlikely to switch between firing states every millisecond, suggested the initialization of a highly diagonal transition matrix [1], [4], [5]. Varying initializations for the state transition probability matrix showed that all high diagonal initializations eventually lead to all estimated state transition probability matrices settling to high values on the diagonal, in contrast to uniformly random initialization, indicating the importance of




Fig. 6. Averaged segment progression frequencies up to a depth of ten. Analysis revealed a significant difference between segments 1 to 4 compared to segments 5 to 10, hence robustness of results is maintained for up to 4 state segments after stimulation.

informed initialization of this parameter. Data log-likelihood results for the high-diagonal cases were comparable, whereas the randomized case exhibited lower values. Finally, the non-straightforward task of choice of number of states was informed by model selection techniques based on commonly used information criteria, in addition to empirical analysis and consistent with related studies reported in literature.

The compact observable representation scheme was chosen for its relative simplicity and proven efficiency in a number of previous reports with discrete HMMs. An attempt to increase the observable alphabet in cases of repeated conflicting signals proved inefficient. Hence, a single extra observable was utilized to universally designate spike overlaps. This compromise may be biasing the HMM in describing periods of maximally heightened activity as a single state in few-state models (e.g. Fig. 4a) without revealing further information about these periods. So, a potential drawback of the scheme could be noted at times of heightened activity or if the dimensionality of the input (recording channels) were to be further increased. This is particularly relevant with recent MEA technologies such as CMOS-based arrays which are capable of achieving much higher spatial resolution [31]. Solutions based on vector quantization schemes have shown that this method too has limitations [32]. Alternative models using Poisson and binomial observables are particularly suited to modelling spike firing dynamics and have shown promise [21], [23], [32], hence such parameterized (observation) distribution models are being considered as alternatives to address the issue.

Analysis of real neuronal data revealed that, in general, culture activity prior to stimulus application was low, followed by heightened network activity starting few ms after stimulation and lasting typically for about 50-100 ms, before decaying

back to pre-stimulus levels. Hence, periods of inactivity were represented by one state. Modeling revealed segmentation of activity into various states, with each state reflecting a combination of channel activity patterns, while trial-based analysis revealed that the majority of ensemble responses reliably progress through certain state segment sequences after stimulus application. This important fact further reinforces the argument that trial-based approaches reveal information that otherwise would be lost by statistical averaging methods. However, further investigation is needed into defining whether the state segmentation is simply a consequence of a dissipative and decaying response to stimulation with no other interesting features, or whether it reveals discrete meta-stable states caused by the switching action of groups of neurons. Since the majority of stimulus response data indicates some form of decaying activity, such discrete meta-stable states, if existing, would be embedded within the overall decay. Tests with simulated data have shown that the HMMs presented here are indeed capable of revealing such patterns embedded within decaying activity, which are otherwise not discernible from either the culture-wide activity profiles (i.e. not differentiated by overall activity levels) or from raster plots (due to overall activity decay obscuring discrimination) (results not shown).

Pre-stimulus period times were considered in order to study whether the cultures would express certain dynamics prior to external stimuli, but the lack of spiking activity meant that these periods were modeled as single 'inactive' states. However, the probing stimulation used is likely activating different underlying neuronal pathways depending on the source (channel) and furthermore does not rule out the possibility that developmental or pharmacological changes in the cultures may alter the dynamic responses of such sources, a possibility that warrants further study (e.g. by comparison of separate HMMs over longer developmental time scales).

The biological mechanisms underlying the formation of the distinct states and their robust transition patterns remains a subject for further investigation, however a possible explanation may be related to the mechanisms of encoding information in the nervous system. An increasing body of work [33]–[35] suggests that information is encoded by the order in which neurons are recruited by the stimulus. Such population based rank encoding would imply repeatable patterns of ordered activation of neurons or neuronal groups shortly after stimulus. Such patterns would likely be progressively less discernible as the information propagates through more and more processing stages. Interestingly, this interpretation is in accordance with repeatable and stimulus dependent patterns of segment progressions as revealed by our HMM approach. Also, the fact that such progressions continue robustly only for a finite number of steps seems consistent with rank based encoding.

The cyclic trains of pulses studied here were used as a means of probing the network state during development without significantly affecting the cultures [27]. However, feedback-controlled training stimuli administered at low frequencies have shown success in inducing plasticity in a number of experiments [7], [11], [16], [36].

Finally, the electrically evoked neuronal activity studied here bears some resemblance (rise-decay patterns, number of

active channels) to spontaneous burst firing patterns that routinely manifest in cortical cultures [27], [37]. Such bursts may result from the absence of afferent sensory input and/or convey and store information [38], making their significance rely upon the context within which they are studied. Consequently, whilst abnormal burst suppression may be a necessary feature for the induction of persistent plastic changes within a closed-loop system [15], spontaneous burst analysis may also reveal important facets underlying biological mechanisms that drive neuronal activity and how these could pertain to information processing by *in vivo* networks. The fact that cultures undergo major developmental changes is also reflected upon spontaneous activity and cannot be overlooked. The suitability of the HMM-based analysis to reveal underlying state changes means that the proposed methods can work equally well in uncovering intrinsically-driven state changes occurring within spontaneous burst events. Further analysis and comparison of results from both spontaneous and evoked activity will uncover the similarities and differences between spontaneous versus induced changes in these neuronal cultures.

V. CONCLUSION

The principal goal of the presented neuronal activity modeling is to develop a robust model of dynamical activity, in order to be used in a closed-loop bi-directional interface between a neuronal culture and an artificial embodiment [39]. Including knowledge of the culture's dynamical state and how that changes will allow for tighter control of neuronal plasticity in such biological-machine hybrids, and eventually lead to enhanced applications and devices at the nervous system-hardware interface. HMM modeling is capable of uncovering hidden state sequences which can provide detailed knowledge of the culture's activity in a given time window and even predict with a degree of certainty its future behavior, whether spontaneous or evoked.

The dataset used for analysis was obtained from cultures grown in similar conditions to our own. Therefore, we expect that the results we obtained are representative for such cultures. Our HMM results showed that this type of stochastic modeling is suitable for the analysis of ensemble networks *in vitro* in addition to *in vivo* applications previously reported in the literature [1], [4], [19], [23]. Suitable models have been selected that allowed capture of subtle shifts in culture-wide activity; trials obtained by applying a probe stimulus to a fixed channel exhibited a robust neuronal culture response in terms of a characteristic sequence of state transition sequences, with observed variability of state transition timings. These highly conserved sequences corresponded to a fast response to stimuli, before subsiding into a low activity state similar to that seen prior to stimulation.

Interpretation of the identified HMM states requires further analysis in terms of the identification of the invariant activity contributing to their segmentation. This in turn may lead to better understanding of the information processing principles leading to the meta-stable state formation in the cultures and also pave the way to characterization of the biological mechanisms responsible for their formation. The link with

the rank based encoding of information in the nervous system seems promising and worth further investigation in this context. Future work will use the tools developed here to analyze larger MEA datasets from our own cultured neuronal networks under varying conditions. The biological underpinnings of the obtained HMM states, especially via order based encoding, will also be pursued. Further improvements to the modeling process will take into account prior knowledge of channel firing probabilities and their interdependence based on functional connectivity analysis, as well as implementation of parameterized observation distributions. This will ultimately lead to improvements in biological-machine hybrid control via neuronal cultures, and in the long-term, lead to better understanding of neuronal computation mechanisms with implications for neuro-prosthetics and brain-machine interfaces.

APPENDIX A

Basic HMM definitions: At any moment of time, t , the HMM is in a state q_t and it emits the corresponding observable o_t . The observables are drawn from a set ('alphabet' of size M) $O = \{O_1, O_2, \dots, O_M\}$ and a particular sequence of T observables is represented as $o = o_1 o_2 \dots o_T$. The HMM hidden states are drawn from a finite state space (N states) denoted as: $Q = \{Q_1, Q_2, \dots, Q_N\}$ and a particular state sequence is $q = q_1 q_2 \dots q_T$. A standard discrete HMM can be fully described by three parameters (Eqs 4-6):

1. The initial state vector π that defines the probability of starting in each state:

$$P(q_1 = Q_i) = \pi(i), \quad 1 \leq i \leq N \quad (4)$$

2. The state transition matrix, A , defines the probability of movement from one state at time step t to the next at time step $t + 1$ (first-order Markov property):

$$P(q_{t+1} = Q_j \mid q_t = Q_i) = A(i, j), \quad 1 \leq i, j \leq N \quad (5)$$

3. Finally, the observation symbol matrix, B , defines the probability of each observable being generated when in a given state:

$$P(o_t = O_k \mid q_t = Q_j) = B(j, k), \quad 1 \leq j \leq N, \quad 1 \leq k \leq M \quad (6)$$

Each row of the matrix corresponds to a conditional probability distribution of a symbol given the state. A HMM model is referenced with the grouped notation: $\lambda = (\pi, A, B)$.

Optimality criteria: The Forward-Backward algorithm uncovers the probabilities for each state at each time step (without taking into account state sequences), thus in order to find the most likely *individual* state at each time step, we find the $\arg \max_{1 \leq i \leq N} [\gamma_t(i)]$ at each step t for every state i , where $\gamma_t(i) = P(q_t = Q_i \mid O, \lambda)$. The Viterbi algorithm however returns the entire most likely state path sequence by taking into account the progressive path at every step of the *entire* sequence. Formally this maximizes $P(q \mid O, \lambda)$.

Programming Environment: All code was written in the MathWorks MATLAB® environment (v.7.3, Release 2006b) on a Linux-based experimental platform (Debian 4.0 Etch, kernel v.2.6.18) and the implementation uses the HMM toolbox

(last updated 8th June 2005) [40], a popular HMM coding tool utilized in the past by other groups [41].

ACKNOWLEDGMENT

The authors thank Prof. Steve M. Potter, Laboratory for Neuroengineering, Coulter Department of Biomedical Engineering, Georgia Institute of Technology, Atlanta, for valuable discussions and generous access to his laboratory's repository of data acquired from MEA recordings.

REFERENCES

- [1] L. M. Jones, A. Fontanini, B. F. Sadacca, P. Miller, and D. B. Katz, "Natural stimuli evoke dynamic sequences of states in sensory cortical ensembles," *Proc. Natl. Acad. Sci. USA*, vol. 104, no. 47, pp. 18 772–18 777, 2007.
- [2] S. Fujisawa, N. Matsuki, and Y. Ikegaya, "Single Neurons Can Induce Phase Transitions of Cortical Recurrent Networks with Multiple Internal States," *Cereb. Cortex*, vol. 16, no. 5, pp. 639–654, 2006.
- [3] G. Rainer and E. K. Miller, "Neural ensemble states in prefrontal cortex identified using a hidden Markov model with a modified EM algorithm," *Neurocomputing*, vol. 32–33, pp. 961–966, 2000.
- [4] E. Seidemann, I. Meilijson, M. Abeles, H. Bergman, and E. Vaadia, "Simultaneously Recorded Single Units in the Frontal Cortex Go Through Sequences of Discrete and Stable States in Monkeys Performing a Delayed Localization Task," *J. Neurosci.*, vol. 16, no. 2, pp. 752–768, 1996.
- [5] M. Abeles, H. Bergman, I. Gat, I. Meilijson, E. Seidemann, N. Tishby, and E. Vaadia, "Cortical activity flips among quasi-stationary states," *Proc. Natl. Acad. Sci. USA*, vol. 92, no. 19, pp. 8616–8620, 1995.
- [6] O. Shefi, I. Golding, R. Segev, E. Ben-Jacob, and A. Ayali, "Morphological characterization of *in vitro* neuronal networks," *Phys. Rev. E*, vol. 66, no. 2, p. 21905 (5pp), 2002.
- [7] J. Le Feber, J. Stegenga, and W. L. C. Rutten, "The Effect of Slow Electrical Stimuli to Achieve Learning in Cultured Networks of Rat Cortical Neurons," *PLoS One*, vol. 5, no. 1, p. e8871 (8pp), 2010.
- [8] M. Chiappalone, P. Massobrio, and S. Martinoia, "Network plasticity in cortical assemblies," *Eur. J. Neurosci.*, vol. 28, no. 1, pp. 221–237, 2008.
- [9] L. Bull and I. S. Uroukov, "Initial Results from the use of Learning Classifier Systems to Control *In Vitro* Neuronal Networks," in *Proc. GECCO '07*. ACM, 2007, pp. 369–376.
- [10] M. Chiappalone, M. Bove, A. Vato, M. Tedesco, and S. Martinoia, "Dissociated cortical networks show spontaneously correlated activity patterns during *in vitro* development," *Brain Res.*, vol. 1093, no. 1, pp. 41–53, 2006.
- [11] G. Shahaf and S. Marom, "Learning in Networks of Cortical Neurons," *J. Neurosci.*, vol. 21, no. 22, pp. 8782–8788, 2001.
- [12] Y. Jimbo, T. Tateno, and H. P. C. Robinson, "Simultaneous Induction of Pathway-Specific Potentiation and Depression in Networks of Cortical Neurons," *Biophys. J.*, vol. 76, no. 2, pp. 670–678, 1999.
- [13] T. B. DeMarse, D. A. Wagenaar, A. W. Blau, and S. M. Potter, "The Neurally Controlled Animat: Biological Brains Acting with Simulated Bodies," *Auton. Robot.*, vol. 11, no. 3, pp. 305–310, 2001.
- [14] A. Novellino, P. D'Angelo, L. Cozzi, M. Chiappalone, V. Sanguineti, and S. Martinoia, "Connecting Neurons to a Mobile Robot: An *In Vitro* Bidirectional Neural Interface," *Comput. Intell. Neurosci.*, vol. 2007, p. (13pp), 2007.
- [15] D. J. Bakkum, Z. C. Chao, and S. M. Potter, "Spatio-temporal electrical stimuli shape behavior of an embodied cortical network in a goal-directed learning task," *J. Neural Eng.*, vol. 5, pp. 310–323, 2008.
- [16] I. Vajda, J. van Pelt, P. Wolters, M. Chiappalone, S. Martinoia, E. van Someren, and A. van Ooyen, "Low-Frequency Stimulation Induces Stable Transitions in Stereotypical Activity in Cortical Networks," *Biophys. J.*, vol. 94, no. 12, pp. 5028–5039, 2008.
- [17] L. R. Rabiner, "A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition," *Proc. IEEE*, vol. 77, no. 2, pp. 257–286, 1989.
- [18] A. Dragomir, Y. M. Akay, and M. Akay, "Modeling carbachol-induced hippocampal network synchronization using hidden Markov models," *J. Neural Eng.*, vol. 7, no. 5, p. 056012 (10pp), 2010.
- [19] A. C. Camproux, F. Saunier, G. Chouvet, J. C. Thalabard, and G. Thomas, "A Hidden Markov Model Approach to Neuron Firing Patterns," *Biophys. J.*, vol. 71, no. 5, pp. 2404–2412, Nov 1996.
- [20] S. H. Chung, J. B. Moore, L. Xia, L. S. Premkumar, and P. W. Gage, "Characterization of single channel currents using digital signal processing techniques based on Hidden Markov Models," *Philos. Trans. R. Soc. B-Biol. Sci.*, pp. 265–285, 1990.
- [21] A. J. Brockmeier, I. Park, B. Mahmoudi, J. C. Sanchez, and J. C. Principe, "Spatio-Temporal Clustering of Firing Rates for Neural State Estimation," in *Proc. 32nd IEEE EMBS Conf. '10*, Buenos Aires, Argentina, 2010, pp. 6023–6026.
- [22] S. Darmanjian and J. Principe, "Spatial-Temporal Clustering of Neural Data Using Linked-Mixtures of Hidden Markov Models," *EURASIP J. Adv. Signal Process.*, vol. 2009, p. (16pp), 2009.
- [23] C. Kemere, G. Santhanam, B. M. Yu, A. Afshar, S. I. Ryu, T. H. Meng, and K. V. Shenoy, "Detecting Neural-State Transitions Using Hidden Markov Models for Motor Cortical Prostheses," *J. Neurophysiol.*, vol. 100, no. 4, pp. 2441–2452, 2008.
- [24] T. Starmer and A. Pentland, "Real-time American Sign Language recognition from Video Using Hidden Markov Models," in *Proc. Int. Symp. Comp. Vis.* '95. IEEE, 1995, pp. 265–270.
- [25] C. Bystroff, V. Thorsson, and D. Baker, "HMMSTR: a Hidden Markov Model for Local Sequence-Structure Correlations in Proteins," *J. Mol. Biol.*, vol. 301, no. 1, pp. 173–190, Aug 4 2000.
- [26] K. Orfanogiannaki, D. Karlis, and G. Papadopoulos, "Identifying Seismicity Levels via Poisson Hidden Markov Models," *Pure Appl. Geophys.*, vol. 167, no. 8–9, pp. 919–931, 2010.
- [27] D. A. Wagenaar, J. Pine, and S. M. Potter, "An extremely rich repertoire of bursting patterns during the development of cortical cultures," *BMC Neurosci.*, vol. 7, no. 1, p. (11pp), 2006.
- [28] D. A. Wagenaar and S. M. Potter, "Real-time multi-channel stimulus artifact suppression by local curve fitting," *J. Neurosci. Methods*, vol. 120, no. 2, pp. 113–120, 2002.
- [29] D. A. Wagenaar, T. B. DeMarse, and S. M. Potter, "MEABench: A toolset for multi-electrode data acquisition and on-line analysis," in *Proc. 2nd Int. IEEE EMBS Conf. Neural Eng. '05*. IEEE, 2005, pp. 518–521.
- [30] I. Visser, M. E. J. Raijmakers, and P. C. M. Molenaar, "Fitting hidden Markov models to psychological data," *Sci. Program.*, vol. 10, no. 3, pp. 185–199, 2002.
- [31] A. Maccione, M. Gandolfo, M. Tedesco, T. Nieuw, K. Imfeld, S. Martinoia, and L. Berdondini, "Experimental investigation on spontaneously active hippocampal cultures recorded by means of high-density MEAs: analysis of the spatial resolution effects," *Front. Neuroeng.*, vol. 3, no. 4, p. (12pp), 2010.
- [32] G. Radons, J. D. Becker, B. Dülfer, and J. Krüger, "Analysis, classification, and coding of multielectrode spike trains with hidden Markov models," *Biol. Cybern.*, vol. 71, no. 4, pp. 359–373, 1994.
- [33] G. Shahaf, D. Eytan, A. Gal, E. Kermany, V. Lyakhov, C. Zrenner, and S. Marom, "Order-based representation in random networks of cortical neurons," *PLoS Comput. Biol.*, vol. 4, no. 11, p. e1000228 (11pp), Nov 2008.
- [34] S. Thorpe, A. Delorme, and R. Van Rullen, "Spike-based strategies for rapid processing," *Neural Netw.*, vol. 14, no. 6–7, pp. 715–725, 2001.
- [35] S. J. Thorpe and J. Gautrais, "Rank Order Coding," *Lect. Notes Comput. Sci.*, vol. 13, pp. 113–119, 1998.
- [36] J. Stegenga, J. Le Feber, E. Marani, and W. L. C. Rutten, "The Effect of Learning on Bursting," *IEEE Trans. Biomed. Eng.*, vol. 56, no. 4, pp. 1220–1227, 2009.
- [37] —, "Analysis of Cultured Neuronal Networks Using Intraburst Firing Characteristics," *IEEE Trans. Biomed. Eng.*, vol. 55, no. 4, pp. 1382–1390, 2008.
- [38] S. M. Potter, "How Should We Think About Bursts?" in *MEA Meeting '08*. Reutlingen, Germany: BIOPRO Baden-Württemberg GmbH, 2008, pp. 22–25.
- [39] D. Xydias, D. J. Norcott, K. Warwick, B. J. Whalley, S. J. Nasuto, V. M. Becerra, M. W. Hammond, J. H. Downes, and S. Marshall, "Architecture for Neuronal Cell Control of a Mobile Robot," in *Proc. EUROS '08*, vol. 44. Prague, Czech Republic: Springer Berlin / Heidelberg, 2008, pp. 23–31.
- [40] K. P. Murphy, "The Bayes Net Toolbox for MATLAB," *Comput. Sci. Stat.*, vol. 33, no. 2, pp. 1024–1034, 2001.
- [41] S. Thorvaldsen, "A tutorial on Markov models based on Mendel's classical experiments," *J. Bioinform. Comput. Biol.*, vol. 3, no. 6, pp. 1441–1460, 2005.



Dimitris Xydas obtained his BEng in Electronic Engineering & Cybernetics in 2006 from the University of Reading (UoR), UK. His final year work won 'best project' and was based on sign language recognition gloves which utilised artificial neural networks. He continued his studies in the UoR to obtain his MSc in Cybernetics in 2007. He is currently in the final year of his PhD degree, involved in the 'Animat' project investigating the computational capabilities of cultured neuronal networks. Dimitris interests lie in robotics, AI, computational neuro-

science and neuroprosthetics.



Dr Slawomir J. Nasuto is a reader in Cybernetics at the UoR. His research interests include investigation of the swarm intelligence, machine learning and information processing in the brain from the perspective of complex systems and the relationship between structure and function. His current research projects include use of long range synchronisation for investigation of memory processes using EEG, EEG based Brain Computer Interfaces, classification of action potentials, automatic reconstruction of neuronal structure from optical microscope stacks

and robots controlled in real time by cultures of biological neurons. He has published over 80 research articles.




Julia H. Downes obtained her MSc in Computer Science at the University of North Florida, USA. Her research was on wireless sensor networks, working closely with the National Weather Service to enable online access to live sensor data. Subsequently, she developed surround sound technologies for mobile devices at Sonaptic Ltd. During her PhD and within the Animat team, she worked on system infrastructure, enabling delivery of electrical stimulus contingent on culture activity and robot sensor data. Her current focus is on neuronal networks, investigating

the interactions between neural units using complex network theory, ultimately to understand network properties of different brain conditions.



Dr Benjamin J. Whalley is a Senior Lecturer in Pharmacology and practising pharmacist at the Reading School of Pharmacy. Dr Whalleys research interests lie in the use of *in vitro* and *in vivo* electrophysiological approaches to investigate neuronal processes that underlie complex physiological functions and disorders, including epilepsy, ataxia, dystonias, and learning and memory. This work has most recently included investigation of the anticonvulsant effects of cannabis constituents, mammalian epileptogenesis and the development of experimental

animat platforms.



Matthew C. Spencer completed his degree in Computer Engineering at the University of Ottawa, Canada, in 2005. After graduation, he worked as a software engineer in the Ottawa area, writing web applications for eLearning and later joining Thales Canada to work on a distributed command and control software suite for the Canadian Forces. In 2008, he started his PhD at the UoR where he is researching evolving complex networks for understanding learning and cognition in neuronal systems.



Dr Victor M. Becerra obtained his PhD for his work on the development of nonlinear optimal control methods from City University, London, in 1994. Between 1994 and 1999 he was a Research Fellow at the Control Engineering Research Centre at City University, London. He is currently a Reader in Cybernetics at the UoR, where he has been an academic since 2000. He is a Senior Member of the IEEE and a Member of the IET. He has published over 100 research papers and one book.



Mark W. Hammond attained a Masters in Neuroscience & AI from the University of Manchester, UK, in 2005 in which research focused on the role of theta rhythm in information flow within Subiculum. He then began a role as an RA/PhD at the University of Reading investigating the computational power of neuronal networks embodied by a robot or Animat which has featured in New Scientist and allowed him to talk at such prestigious institutions as the National Science museum. Currently his work focuses on the role of the neurotransmitter acetylcholine in

learning, memory and attention within the Animat paradigm.



Prof. Kevin Warwick received the BSc degree from Aston University, Birmingham, UK and the PhD degree from Imperial College, London, UK. He has been awarded higher doctorates (DSc) by Imperial College, London and the Czech Academy of Sciences, Prague. He is Professor of Cybernetics at the UoR, where he is involved in research into artificial intelligence, control, robotics and cyborgs. He is the author or coauthor of more than 500 research papers and is perhaps best known for his experiments using implant technology. Prof. Warwick was the recipient

of The Future of Health Technology Award in MIT, was made an Honorary Member of the Academy of Sciences, St Petersburg and received the IEE Senior Achievement Medal in 2004, the Mountbatten Medal in 2008, the Golden Eurydice in 2009, the Ellison-Cliffe Medal in 2011. He has also received Honorary Doctorates from the Universities of Aston, Coventry, Bradford and Robert Gordon.