

Learning in Networks of Cortical Neurons

Goded Shahaf and Shimon Marom

Department of Physiology and Biophysics, Faculty of Medicine, Technion, Haifa, 31096 Israel

The results presented here demonstrate selective learning in a network of real cortical neurons. We focally stimulate the network at a low frequency (0.3–1 Hz) until a desired predefined response is observed 50 ± 10 msec after a stimulus, at which point the stimulus is stopped for 5 min. Repeated cycles of this procedure ultimately lead to the desired response being directly elicited by the stimulus. By plotting the number of stimuli required to achieve the target response in each cycle, we are

able to generate learning curves. Presumably, the repetitive stimulation is driving changes in the circuit, and we are selecting for changes consistent with the predefined desired response. To the best of our knowledge, this is the first time learning of arbitrarily chosen tasks, in networks composed of real cortical neurons, is demonstrated outside of the body.

Key words: learning; multielectrode array; cultured neurons; neural network; reward; drive reduction

Learning a new behavioral task is an exploration process that involves the formation and modulation of sets of associations between stimuli and responses. In an effort to understand the phenomenon of learning, two different questions are asked. (1) What are the neural mechanisms that underlie the formation and modulation of associations? (2) What are the principles that underlie the selection of “appropriate” associations over “inappropriate” ones? The nature of mechanisms underlying the formation and modulation of associations has been the topic of intense research. Although much is yet to be discovered, many mechanisms were described, at various levels of neural organization, that can support activity-dependent modification of associations between stimuli and responses. This study addresses the second question, the principles underlying the selection of an appropriate association during the learning process.

Our learning experiments were performed in networks containing 10,000–50,000 cortical neurons obtained from newborn rats (Baughman et al., 1991), under the assumption that the organizing principles operating at the level of neuronal populations are intrinsic to neurons and are therefore manifested *ex vivo*. Such cultured cortical networks were thoroughly studied by others (Ramakers et al., 1990; Murphy et al., 1992; Maeda et al., 1995; Canepari et al., 1997; Voigt et al., 1997; Turrigiano et al., 1998), and a substantial amount of data has been accumulated, showing that they are structurally rich, develop and adapt functionally and morphologically over a broad range of time scales, and are experimentally stable over weeks.

In what follows, we show that the large random cortical networks developing *ex vivo* display general properties required from neural systems capable of learning: namely, numerous connections, stability of connections, and modifiability by external stim-

uli. We then describe closed-loop experiments in which these biological networks interact with a computer-controlled environment and demonstrate a simple procedure for learning and memorizing arbitrarily chosen tasks defined in terms of neuronal firing patterns. Specifically, we show that, during regular low-frequency stimulation, the network explores a large space of possible connections and can be instructed to select and stabilize one or a subset of them by withdrawing the stimulus at the point that the connection is observed.

MATERIALS AND METHODS

Culture techniques. Cortical neurons are obtained from newborn rats within 24 hr from birth, following standard procedures. The cortex tissue is digested enzymatically and mechanically dissociated. The neurons are plated directly onto substrate-integrated multielectrode array (MEA) dishes (see below). The cultures are bathed in MEM supplemented with 5% heat-inactivated horse serum, 0.5 mM glutamine, 20 mM glucose, and 10 mg/ml gentamycin, and maintained in an atmosphere of 37°C, 5% CO₂ and 95% air in a tissue culture incubator and during the recording phases. Half of the medium is exchanged twice per week. Experiments are performed in the third week after plating, thus allowing complete maturation of the neurons. Networks that did not respond (in the third week after plating) to repeated low-frequency stimulation (1, 0.5, and 0.3 Hz) were not kept for additional experimentation.

The electrical activity of the cultured network is dependent on synaptic transmission; there are many published reports (Maeda et al., 1995; Turrigiano et al., 1998; and references therein) showing that the electrical activity in a cultured cortical network may be blocked by perfusion with the NMDA receptor antagonist D-2-amino-5-phosphonovalerate (APV) and the non-NMDA receptor antagonist CNQX. We repeated these experiments using intracellular recordings, as well as MEA recordings. We find that, in the presence of a mixture of synaptic blockers containing 5 μ M bicuculline, 10 μ M DNQX, and 20 μ M APV, spiking activity within the cultured network is completely abolished (Tal, 2000).

Electrophysiological methods. We used the substrate-embedded multi-electrode array technology (see Fig. 1A) (Gross, 1994; Meister et al., 1994). We used arrays of 60 Ti/Au/TiN electrodes, 30 μ m in diameter, and spaced 200 μ m from each other [MultiChannelSystems (MCS), Reutlingen, Germany]. The insulation layer (silicon nitride) was pre-treated with poly-L-lysine, forming a good surface for network development. A commercial 60 channel amplifier (B-MEA-1060; MCS) with frequency limits of 10–3000 Hz and a gain of 1024 \times was used. The B-MEA-1060 was connected to MCPPlus filter amplifiers (Alpha Omega, Nazareth, Israel) for additional amplification (10–20 \times). Stimulation through the MEA is performed using a dedicated eight channel stimulus generator (MCS). The microincubation environment was arranged to support long-term recordings from MEA dishes. This was

Received July 24, 2001; revised Sept. 4, 2001; accepted Sept. 5, 2001

This study is supported by grants from the European Commission, the Israel Science Foundation, and the Bernard Katz Minerva Foundation. We thank Vladimir Lyakhov and Ella Romanko-Lyakhov for their invaluable technical support and Dr. Noam Ziv for encouragement and helpful discussions. We also thank Drs. Shraga Hoeherman, Larry Manevitz, Daniel Dagan, Itzik Schiller, and Jackie Schiller for their suggestions in the preparation of this manuscript.

Correspondence should be addressed to Shimon Marom, Department of Physiology and Biophysics, Faculty of Medicine, Technion, Haifa, 31096 Israel. E-mail: marom@tx.technion.ac.il.

Copyright © 2001 Society for Neuroscience 0270-6474/01/218782-07\$15.00/0

achieved by streaming a filtered, heated, and 95% humidified air–5% CO₂ gas mixture and by electrically heating the MEA platform to 37°C. Data were digitized using two 5200a/526 analog-to-digital boards (Microstar Laboratories, Bellevue, WA). Each channel is sampled at a frequency of 24,000 samples/sec and prepared for analysis using the AlphaMap interface (Alpha Omega).

Spike detection. Thresholds ($8 \times$ root mean square units; typically in the range of 10–20 μ V) are separately defined for each of the recording channels before the beginning of the experiment. No additional spike sorting techniques are applied for the following reason. Much of our data and their interpretation correspond to the time scale of intraburst activities. Whole-cell recordings from single cortical neurons, both in our hands as well as in many published records, show that the shape of an action potential changes dramatically within bursts of activity because of the dynamics of membrane excitability. Consequently, any attempt to sort the spikes according to their shapes within this time scale is doomed a priori. We were therefore forced to resort to a stricter approach that defines elementary activities on the basis of their participation in statistically significant activity pairs as explained below, in which every threshold crossing is considered in the analysis. The major limitation of this approach is that it takes more occurrences of a particular pair to define statistical significance. This limitation was overcome by performing long experiments.

Definition of activity pairs. We operationally define pairs of neural connectivity in terms of an action potential A that is followed by another action potential B, with a precise time delay ($\tau \pm 0.5$ msec) between the two. A and B may be action potentials recorded in the same or in different measuring electrodes. Both events (A and B) are defined by threshold crossing as explained above. The number of measuring electrodes (N_e) dictates the maximal number of detectable pairs. Thus, for a $\tau > 0$, the maximal number of A→B pairs is N_e^2 . For $\tau = 0$, the maximal number of A→B pairs is $N_e(N_e - 1)$; an activity cannot pair with itself within a zero time delay.

Statistical significance of activity pairs. The statistical significance (p value) of a given A→B pair is calculated under binomial distribution assumptions given the number of times A occurred, the number of times A→B occurred with a time delay τ , and the probability of event B. Thus, if $p(k)$ is the probability of observing k or more A→B pairs out of n A events, and p_B is the probability of a B event, then:

$$p(k) = \sum_{i=k}^{\infty} \binom{n}{i} p_B^i (1 - p_B)^{n-i} = 1 - \sum_{j=0}^{k-1} \binom{n}{j} p_B^j (1 - p_B)^{n-j}$$

$p < 0.01$ was used as a significance measure.

Functional strength of activity pairs. Given an A→B activity pair, the forecasting of B by A, which is the strength of the functional connectivity between the two, is given in terms of a correlation coefficient, calculated from the number of times that the given pair appears within 1 hr, divided by the number of occurrences of A or B (see Fig. 1B, inset).

Stability of activity pairs. For each A→B pair, statistical significance of a change in pair co-occurrence counts was calculated under the assumptions of the binomial distribution (see above and Fig. 1C). For instance, suppose that A→B pair (e.g., with $\tau = 20$ msec) appeared n_1 times in the first 0.5 hr bin and n_4 times in the fourth 0.5 hr bin. To state that n_4 is significantly different from n_1 , we calculate the probability of finding n_4 or more events (for the case of $n_4 > n_1$) or n_4 or less events (for the case of $n_4 < n_1$). This is done using the frequency of A→B (at $\tau = 20$ msec) in the first 0.5 hr bin as the theoretical probability and the number of A events in the fourth 0.5 hr bin as the number of trials. If the calculated probability is $p < 0.01$, then n_4 is significantly different from n_1 .

Stimulation parameters. The pair of stimulating electrodes was chosen according to its ability to induce a reverberating electrical activity of the type shown in Figure 1, D and E, in response to a biphasic current pulse (± 50 μ A or smaller, lasting < 500 μ sec; 250 μ sec for each phase). At stimulation frequencies higher than 1 Hz, the networks usually inactivate after a few pulses. Therefore, in the learning experiments, the frequency of stimulation for a given network was set at either 1, 0.5, or 0.3 Hz, the highest that was possible for the particular network without inactivating its response (see Figs. 2–6). Stimulating electrodes were spatially near each other (~ 200 – 400 μ m apart).

Peristimulus time histogram construction. A series of 1200 stimuli (420 μ sec, 50 μ A, 0.3 Hz) was delivered through a pair of electrodes, and the responses in 10 randomly chosen active electrodes were recorded (see


Fig. 1E). The total number of responses (counted in 1 msec time bins) divided by 12000 is presented, time-locked to the stimulus event.

Stimulation protocol and analysis of activity-dependent change of activity pairs. Each network was exposed to nine stimulation sessions (see Fig. 1F). The pattern of stimulation in each of these nine sessions was one of the following: pattern 1, 10 min at 0.3 Hz; pattern 2, 2 min at 0.3 Hz, followed by 8 min of no stimulation; or pattern 3, 10 min of no stimulation. Each of these stimulation patterns was delivered three times to each network, in a random temporal order. Every stimulation session was preceded and succeeded by 100 test stimuli. For a given network, all the stimuli, including the test stimuli, were delivered through the same pair of electrodes. The test stimuli enabled us to define significantly occurring activity pairs as explained above, with the prefix of each pair (A in A→B) being the stimulus itself. Using the binomial theorem (see above and Fig. 1B,C), we identified activity pairs whose count changed during stimulation patterns 1 and 2 in a statistically significant manner and display the average number of such pairs normalized to the average spontaneous change (pattern 3). The data shown in Figure 1F was obtained from testing all A→B pairs (A being the stimulus itself) with pair time delay (τ) between 0 and 100 msec, in 1 msec bins.

RESULTS

The cultured neurons form numerous synaptic connections. This is apparent from the large number of statistically significant correlated activities between pairs of electrodes. We operationally define such pairs of neural connectivity in terms of an action potential A that is followed by another action potential B with a precise time delay ($\tau \pm 0.5$ msec) between the two (see Materials and Methods). Analysis of the spontaneous activity of the network, without any stimulation, suggests that the average number of such statistically significant A→B connections is a large percentage of the maximum that is possible at relatively small values of τ (Fig. 1B). As the time delay between the activities of the elements of the pair becomes longer, the realized number of pairs decreases. Of course, a significant occurrence of A→B connection might represent a causal relationship between the activity of A and that of B or a noncausal correlation resulting from coactivation by a common source. Furthermore, many of the observed connections are actually parts of larger groups of significantly connected activities. However, as we proceed, it will become clear that, for the purposes of this study, distinctions between the possibilities mentioned above are not crucial. Rather, the important thing is that the number of connections is large (Fig. 1B) and that many independent activity patterns exist. The latter is implied from the fact that, in these networks, single neurons seldom fire spontaneously without being activated by other neurons (Maeda et al., 1995; Canepari et al., 1997) (see Materials and Methods), whereas the average correlation between elements of pairs is rather weak (Fig. 1B, inset). The stability of connections in the network may be appreciated by comparing the number of times each of the significantly occurring pairs appeared in 10 consecutive time bins (30 min each) over 5 hr of continuous recording of spontaneous activity, without any stimulation. We used the number of times that a given A→B activity pair appeared in the first 30 min bin, divided by the number of occurrences of A or B as a measure for the occurrence probability of a pair. Using the binomial theorem, we identified pairs whose count did not change in a statistically significant manner in subsequent time bins. Figure 1C shows that $\sim 70\%$ of the pairs remained unchanged after 5 hr of spontaneous activity.

When stimulating currents are delivered through a pair of substrate-embedded electrodes at a constant frequency, the network responds by generating a rich repertoire of reverberating electrical activities, lasting 100 msec or more (Fig. 1D,E). Modifications in functional connectivity would be manifested as



changes in the coupling of such responses to the stimulus. Indeed, repeated stimulation induces changes in network responsiveness, as shown previously by others (Jimbo et al., 1999). Furthermore, Figure 1*F* shows that the magnitude of such modifications increases with stimulation time, reflecting the myriad activation pathways and activity-dependent mechanisms that operate in these networks. This “exploratory” nature (of the change in response to series of stimuli) is further demonstrated in the data presented below.

The analyses presented above imply that cortical networks display general properties expected from neural systems capable of learning: namely, numerous connections, stability of connections, and modifiability by external stimuli. We now turn to the novel aspect of this study, which is demonstrating learning in a cortical network without the involvement of a neural rewarding

entity. The idea is simply to stimulate the network until the required response is attained, and once this occurs, to remove the “driving” stimulus. We then ask how long it takes to attain the required response. Will the appropriate responsiveness remain stable after such a simple procedure? How selective can such a change in connectivity be? If after the procedure the required response to stimulus occurs reliably and selectively, this could be considered as a form of learning.

Each experiment starts by stimulating the network through a pair of electrodes and observing the responsiveness of all other (i.e., the nonstimulated) electrodes. A nonstimulated electrode that responds 50 ± 10 msec after a stimulus with a response-to-stimulus (R/S) ratio of 1/10 or less is selected. In other words, before training, it takes at least 10 stimuli to evoke one action potential in the selected electrode within the designated time




Figure 2. Example of learning in a cultured network of cortical neurons. Each trace within a panel shows recordings obtained 10 msec before the stimulus to 70 msec after the stimulus, before (*left*) and after (*right*) the training procedure. Note that the responsiveness of the electrode within the designated time window increased appreciably.

frame of 50 ± 10 msec after a stimulus. During the training phase, the learning task is to increase the R/S of the selected electrode to $2/10$ or greater at the designated time window of 50 ± 10 msec after a stimulus. The two stimulated electrodes are continuously stimulated at a constant frequency of $1/3$, $1/2$, or 1 stimulus per second. A computer constantly monitors the R/S of the selected electrode, and once the criterion of $R/S \geq 2/10$ is fulfilled, i.e., whenever two responses were seen in any 10 consecutive trials, the computer automatically stops the stimulation. After 5 min, the network is stimulated again (at the same low frequency) until the criterion $R/S \geq 2/10$ in the same selected electrode is fulfilled again. This stimulation cycle, which is composed of 5 min without stimulation followed by low-frequency (0.3, 0.5, or 1 Hz) stimulation until $R/S \geq 2/10$ criterion in the selected electrode is fulfilled, is repeated many times. As a rule, if the criterion is not fulfilled within 10 min of stimulation, the stimulation is stopped for 5 min. Hence, the maximal duration of one stimulation cycle is 15 min (i.e., 10 min of stimulation and 5 min of quiescence). The latency for reaching the predetermined criterion (referred to as response time) in each stimulation cycle is used as a measure for the strength of S–R connection and may be viewed as a measure of the degree to which the task was learned.

An example for the result of this learning procedure is shown in Figure 2. It includes the responses of a selected electrode before (*left column*) and after (*right column*) training. The 11 traces of each panel show the responses to 11 consecutive stimulation pulses. Note that the activity within the 50 ± 10 msec window (depicted) is markedly increased after the training phase.




Figure 3. *A*, Eight learning curves, differing in their learning kinetics. The response time (i.e., the duration of stimuli series until criterion is fulfilled) is plotted against the number of stimulation cycle. Each point depicts the time (in seconds) to accomplish the task in one cycle. *B*, Four control curves from a protocol in which each cycle consisted of 10 min of stimulation and 5 min of quiescence, regardless of response (see Results). Vertical broken lines are referred to in Figure 5.

Figure 3*A* shows eight learning curves, differing in the learning kinetics. In these curves, the response time (i.e., time required for the selected electrode to fulfill the $R/S \geq 2/10$ criterion) is plotted against the number of stimulation cycle. (Recall that each stimulation cycle is composed of 5 min without stimulation, followed




Figure 4. Eight selective learning curves, differing in their learning kinetics. The response time (i.e., the duration of stimuli series until criterion is fulfilled) is plotted against the number of stimulation cycle. Each *point* depicts the time (in seconds) to accomplish the task in one cycle. *Vertical broken lines* are referred to in Figure 5.

by low-frequency stimulation until $R/S \geq 2/10$ criterion is fulfilled.) The curves are characterized by response time decrement and stabilization at lower values compared with the initial values. Note that the time required to instruct a network to perform the task varies, reflecting the arbitrariness of the procedure by which the tasks are chosen and the idiosyncrasies of the networks.

The notion that driving stimulus removal is necessary for selecting appropriate network responses is further supported by a control experiment. In this experiment, the fulfillment of the R/S criterion in the selected electrode did not lead to stimulus removal (i.e., the attainment of the criterion was ignored). The stimulation was delivered for 10 min interrupted by 5 min of




Figure 5. Changes in R/S, depicted by f , of the selected electrode (*filled circles*) and 10 control electrodes (*stars*) for eight selective learning experiments (*left*) and four control experiments (*right*) whose curves are shown in Figure 3B. The *left eight columns* correspond to the eight selective learning curves of Figure 4 by a *left-to-right, top-to-bottom* order. (For instance, the selectivity data that relates to the *top right* curve of Fig. 4 is shown in the *second column from the left* in this figure). For the purpose of calculating f , the point in time that separates the period before training and that of after training is depicted by *broken lines* in the curves of Figures 3B and 4. f is normalized to the R/S change of the selected electrode.

quiescence, regardless of the responses recorded from the selected electrode. Figure 3B shows that, under these conditions, the response time (i.e., the time required for first appearance of $R/S \geq 2/10$ within each stimulation cycle) plotted against the stimulation cycle number shows large fluctuations and a tendency to a decreased responsiveness over time.

Thus far, our criterion for stopping the stimulus has simply been the appearance of a response on the selected electrode. We refer to the eight trials that used this criterion as “simple learning” trials. To ensure selectivity of the R/S increase in the selected electrode, we also conducted 16 trials using a second criterion. In these trials, we monitored a second electrode in the array, which serves as a measure for global network responsiveness. Our condition for removing the stimulus was that the R/S criterion be fulfilled in the selected electrode and not fulfilled in the second monitored electrode. We refer to these as “selective learning trials.” Of these 16 selective learning trials, eight showed learning (Fig. 4). In the remaining eight selective learning trials, the latency for reaching the predetermined R/S criterion did not relax; that is, the response times did not decrease and did not stabilize at lower values compared with the initial values. Such “nonrelaxing” experiments were stopped after 25 stimulation cycles.

Figure 5 (*left eight columns*) summarizes the selective learning data. Changes in R/S of the selected electrodes (*filled circles*) and 10 control electrodes (*stars*) are depicted for eight experiments




Figure 6. Average control (*curve 1*), average response over all noncontrol trials (*curve 2*), and average learning curve (*curve 3*). Each point depicts the average time (in seconds) to accomplish the task in one cycle. Filled gray circles (in *curve 2*) and filled black squares (in *curve 3*) depict points that are significantly different from averaged control (*curve 1*; F test, $p < 0.05$). Fitted power function lines ($y = ax^b$) are added to emphasize the different trends of the curves. Note that this figure shows averages of trials with different durations (curves from Figs. 3 and 4, as well as eight nonrelaxing curves that were interrupted after 25 cycles; see Results). Hence, the number of samples contributing to each point varies as follows: *curve 1*, $n = 4$ for cycles 1–40; *curve 2*, $n = 24$ for cycles 1–10, $n = 16$ for cycles 11–20, and $n = 8$ for cycles 21–25; *curve 3*, $n = 16$ for cycles 1–10, $n = 8$ for cycles 11–20, and $n = 4$ for cycles 21–40.

from eight different networks. For each network, the 10 control electrodes were chosen by analyzing the data, after the completion of the experiment, based on their similarity to the R/S of the selected electrode before the training; specifically, the control electrodes are the 10 electrodes whose R/S before training were the most similar to the R/S of the selected electrode before training. The change, depicted by f , is defined as the ratio between the responsiveness before training and responsiveness after training, normalized to the change in R/S of the selected electrode. Thus, $f = 1$ means a change in R/S that is identical to the change measured in the selected electrode. $f > 1$ and $f < 1$ mean that the relevant response of a control electrode increased or decreased, respectively, relative to the selected electrode. Note that the strengthening in the response-to-stimulus ratio (R/S) of the selected electrode is generally higher relative to the responsiveness

change in the control electrodes. Also note that, because the selected and control electrodes demonstrate low responsiveness before the training, a bias toward an average increase of R/S during training is introduced. The reported effect is selective because the increase in R/S of the selected electrode is more than the average increase for the control electrodes. The probability of the selected electrode to be ranked fourth or higher (of 11), as is the case in the eight experiments shown, is $\leq (4/11)^8$. Note that, in the control trials (*four right columns*), no preferred ranking of the selected electrode is observed.

Figure 6 summarizes the entire data set obtained in the above described experiments: the average control curve (*curve 1*), the average behavior of the entire set of trials (*curve 2*; including the nonrelaxing trials), and the average learning curve (*curve 3*; the combined set of simple and selective learning curves). Each point depicts the average time (in seconds) to accomplish the task in one cycle within a series of cycles. Figure 6 provides an indication for the robustness of the main phenomena shown in this study: when the loop is closed and the response is allowed to remove the stimulus, learning curves may be obtained; when the loop is open, i.e., the computer is instructed not to remove the stimulus when the selected electrode criterion is fulfilled, the curves “explore away.”

DISCUSSION

The experiments described above show that sufficient conditions for the realization of learning by a selection process, without the involvement of a neural rewarding entity, are embodied in large random networks of neurons maintained *ex vivo*. These networks form a large space of connectivity configurations that are stable over many hours. The connectivity can be modulated by external focal stimulation in an activity-dependent manner. Most importantly, the networks explore the space of possible responses and stabilize at configurations that remove the stimuli.

From the theoretical point of view, the above demonstration conveys an important message, supported by behavioral studies and psychological theories advocated over 50 years ago by eminent psychologists such as Hull (1943) and Guthrie (1946): it is not necessary to assume a separate mechanism for the biological realization of a reward in distinction from the process of exploration for solutions; the behavioral concept of reward might be considered as a change (removal) in the drive underlying the exploration in the space of possible modes of network response. A drive to explore that is removed when a desired state is achieved is an intentionless natural principle for adaptation to rich and unlabeled environment. Of course, the fact that learning by stimulus removal is plausible biologically does not mean that it is implemented in real brains; however, the simplicity of this principle makes it very likely that it does.

REFERENCES

- Baughman RW, Huettner JE, Jones KA, Khan AA (1991) Culturing nerve cells (Banker G, Goslin K, eds), pp 227–250. Cambridge, MA: MIT.
- Canepari M, Bove M, Maeda E, Cappello M, Kawana A (1997) Experimental analysis of neuronal dynamics in cultured cortical networks and transitions between different patterns of activity. *Biol Cybern* 77:153–162.
- Gross GW (1994) Enabling technologies for cultured neural networks (Strenger DA, Meister M, eds), pp 277–317. San Diego: Academic.
- Guthrie ER (1946) Psychological facts and psychological theory. *Psychological Bulletin* 43, Presidential address of the American Psychological Association, Evanston, IL, 1945.

- Hull C (1943) Principles of behavior. New York: Appleton-Century-Crofts.
- Jimbo Y, Tateno T, Robinson HPC (1999) Simultaneous induction of pathway-specific potentiation and depression in networks of cortical neurons. *Biophys J* 76:670–678.
- Maeda E, Robinson HPC, Kawana A (1995) The mechanism of generation and propagation of synchronized bursting in developing networks of cortical neurons. *J Neurosci* 15:6834–6845.
- Meister M, Pine J, Baylor DA (1994) Multi-neuronal signals from the retina: acquisition and analysis. *J Neurosci Methods* 51:95–106.
- Murphy TH, Blatter LA, Wier WG, Baraban JM (1992) Spontaneous synchronous synaptic calcium transients in cultured cortical neurons. *J Neurosci* 12:4834–4845.
- Ramakers GJ, Corner MA, Habets AM (1990) Development in the absence of spontaneous bioelectric activity results in increased stereotyped burst firing in cultures of dissociated cerebral cortex. *Exp Brain Res* 79:157–166.
- Tal D (2000) Analysis of the input–output function of cortical neurons, PhD thesis, Technion-Israel Institute of Technology.
- Turrigiano GG, Leslie KR, Desai NS, Rutherford LC, Nelson SB (1998) Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nature* 391:892–896.
- Voigt T, Baier H, Dolabela DL (1997) Synchronization of neuronal activity promotes survival of individual rat neocortical neurons in early development. *Eur J Neurosci* 9:990–999.