

9

Closing the Loop: Stimulation Feedback Systems for Embodied MEA Cultures

STEVE M. POTTER, DANIEL A. WAGENAAR, AND THOMAS B. DEMARSE

9.1 Introduction

In the vertebrate nervous system, all key functions are performed by many nerve cells working in concert. By contrast, single neurons are sometimes important for specific functions in invertebrates, such as the locust's giant motion detector neuron (LGMD), which integrates visual information and triggers jumping (Gabbiani et al., 1999). To understand more complex processing in vertebrates, we need to know how single-unit activity combines to form the network-level dynamics that take in sensory input, stores memories, and controls behavior. Cultured neuronal networks have provided us with much of our present understanding of ion channels, receptor molecules, and synaptic plasticity that may form the basis of learning and memory (Bi and Poo, 1998; Latham et al., 2000; Misgeld et al., 1998; Muller et al., 1992; Ramakers et al., 1991). To study the nervous system *in vitro* offers many advantages over *in vivo* approaches. *In vitro* systems are much more accessible to microscopic imaging and pharmacological manipulations than are intact animals. Recent developments in multi-electrode array (MEA) technology, including those described below, will enable researchers to answer questions not just at the single-neuron level, but at the network level. Most MEA research has involved recording activity that cultured networks produce spontaneously, via up to 64 extracellular electrodes. Although some studies also included electrical stimulation via the substrate electrodes, it was applied to only one or two of them at a time (Connolly et al., 1990; Fromherz and Stett, 1995; Gross et al., 1993; Jimbo and Kawana, 1992; Jimbo et al., 1998; Maeda et al., 1995; Oka et al., 1999; Regehr et al., 1989; Shahaf and Marom, 2001; Stoppini et al., 1997). We propose that in order to substantially advance our understanding of network dynamics, we need high-bandwidth (many neurons) communication in *both* directions, out of and into the network. This chapter describes technologies that allow recording and stimulation on every electrode of an MEA, and a new closed-loop paradigm that brings *in vitro* research into the behavioral realm.

9.1.1 *The Importance of Embodiment*

Nervous systems evolved to aid the survival of motile organisms, by directing their interactions with their surroundings. In the natural environment, sensory input to an animal's nervous system is largely a function of its recent output: as the animal moves and interacts within its environment, its sensory systems are actively oriented so as to provide the information that will be most useful in controlling subsequent behaviors (Nolfi and Parisi, 1999). Animals sense the consequences of an action for a few tens of milliseconds after the motor command is sent to muscles. Neural output is expressed continuously, while it is being modulated by a continuous stream of sensory inputs. This tight sensory–motor loop is likely to be important for learning to predict the consequences of actions and to create and fine-tune adaptive behaviors.

Contrast this with traditional *in vivo* neurophysiology in the lab: the animal is often anaesthetized and immobilized, unable to behave. Rarified, unimodal sensory input is provided in brief exposures, or “trials”, and neural responses are measured directly with electrodes. This open-loop approach has been helpful in understanding cortical maps, and neural receptive fields, among other things. But we must interpret such findings with care; it is likely that the dynamics of the nervous system under these unnatural circumstances are substantially different than in a freely behaving animal (Hartmann and Bower, 2001). We wish to bring *in vitro* networks closer to the kind of neural processing that nervous systems evolved to do, that is, to take in sensory information continuously, process it continuously, and express adaptive behavior continuously. Each of these processes interacts with the others, so they should not be studied separately.

9.1.2 *A New Closed-Loop Research Paradigm*

In vitro networks have always been incapable of expressing behavior, by the very fact of being removed from the donor's body. We have developed systems for re-embodiment of cultured networks, allowing them once again to express behavior. Our embodiments, or neurally controlled animats, are either simulated creatures on the computer (DeMarse et al., 2001), or actual robots (Bakkum et al., 2004). The whole system of MEA culture plus embodiment we call a “hybrot” because it is a hybrid robot with both living and artificial components. The greatest advantage of these systems is that their “brains” hold perfectly still on the microscope stage, amenable to detailed imaging for months at the submicron level, while controlling behavior and receiving sensory inputs (Figure 9.1). Such detailed, extended imaging is not presently possible in the brains of behaving animals. Two-photon time-lapse fluorescence optical microscopy can be used with rodent cortical networks on MEAs in order to find links between morphological and functional dynamics at the network level. This new approach of embodied cultured networks will help us and others address questions about how activity shapes network development, how neuromodulatory systems influence connectivity, and what the bases are of distributed neural activity for sensory processing, memory formation, and behavioral control.

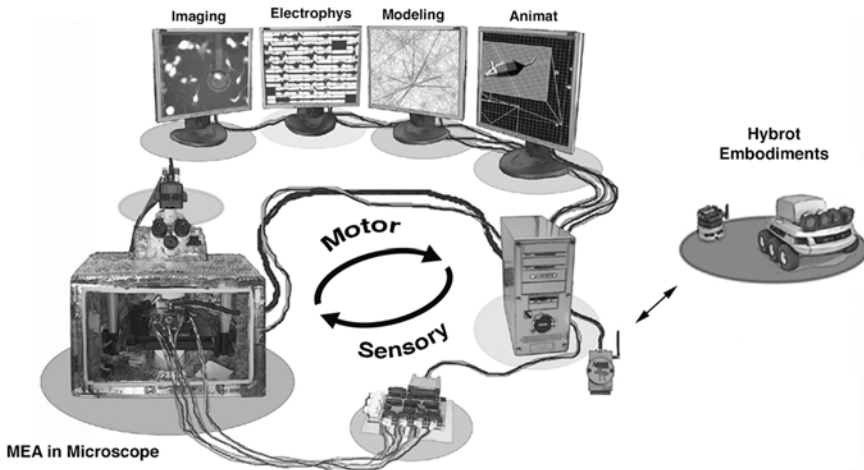


FIGURE 9.1. Embodying cultured neurons. We closed the sensory–motor loop between an MEA, serving as the “brain” of a Hybrot, and a simulated or robotic animat as its embodiment. While controlling behavior and receiving sensory stimulation, the cultured network can be imaged at the micron scale using time-lapse fluorescence microscopy. (From Zenas Chao)

9.2 Long-Term Cell Culture

As with any model system, dissociated primary neuronal cultures have advantages and disadvantages. Unlike cells from other tissues, mammalian neurons are (usually) terminally differentiated when obtained from late-term embryos or neonates, and thus cannot be multiplied by passaging. This means that a ready supply of brain tissue donors must be available, and that the cultures must be pampered more than those prepared from dividing cells, if long-term studies are to be carried out. This includes rigorous adherence to sterile techniques, careful choice and replenishment of media, and maintenance of pH, temperature, and osmolarity (Banker and Goslin, 1998). The reward for this extra effort is the ease of both observation and manipulation of neural circuits, allowing numerous types of inquiries not feasible in humans or even lab animals.

Both the blessing and the curse of neural cultures is that they are much simpler than living brains. Neurons, especially when cultured with glia (Ullian et al., 2001), spontaneously form functional synapses *in vitro*, and develop complex patterns of activity that closely resemble those recorded from developing brains of animals (Ben-Ari, 2001). Neurons retain their morphological and pharmacological identities in culture, but there are likely to be numerous subtle changes in their properties due to the unnatural environment in which they have been placed. A model system is only helpful if it is simpler than the thing being modeled. But *in vitro* researchers must always be wary of the limitations of their model systems. We are improving neural cell culture in an effort to remove some of those limitations, as described below.

9.2.1 Dense Monolayer Cultures

Of in vitro model systems, dissociated monolayer cultures provide the best access by electrodes (whether micro-pipettes or substrate-integrated electrodes), drugs, and microscopic imaging. If we are to make full use of multi-electrode array substrates, each extracellular electrode should be able to stimulate and record from at least one neuron. Efforts to trap neurons next to electrodes with micro-engineered structures such as silicon wells or pillars have been successful, at least for short-term cultures (Maher et al., 1999; Merz and Fromherz, 2002). Mammalian neurons survive longer and develop more synapses when in direct contact with glial cells (Ullian et al., 2001; personal observations), but these often provide the neurons with the tensile forces needed for escape. In working with silicon neurochips and neuroprobes in the Pine lab at Caltech, Potter and colleagues Mike Maher and Hannah Dvorak-Carbone even observed neurons escaping from their cages at the expense of leaving their nuclei behind! Needless to say, they died soon after. Keeping neurons near electrodes by treating the electrode regions with “neurophilic” chemicals or the surrounding regions with “neurophobic” chemicals, has also been reasonably successful. Branch et al. (1998) have maintained adherence to stamped patterns of poly-D-lysine for about one month in vitro. This approach shows much promise, not only for keeping neurons in close apposition to electrodes, but for creating well-defined simple neural circuits in vitro (Nam et al., 2003).

We chose a different, far simpler approach to ensure that each of the 60 electrodes in the MEAs we use (Multi Channel Systems, GmbH) is functionally interfaced to one or more neurons: plate them very densely (Figure 9.2). We usually plate 20,000 to 50,000 mouse cortical cells in a three millimeter diameter region, resulting in densities of 5,000 to 10,000 cells per square millimeter in the electrode region.




FIGURE 9.2. FEP-membrane sealed MEA (left) and dense monolayer culture of neurons and glia (right), after several months in vitro.

Thus, each 10 μm diameter electrode will have at least one and usually several neurons within recording and stimulation range.

Considering the three-dimensional structure of intact brains, to grow in a monolayer is clearly not what neurons and glia evolved to do. We and others who tried to produce dense monolayers often observed that cells would tend to clump up and form many-cell balls, following their tendencies to adhere to and migrate along each other during development in vivo. Although such clusters of neurons and glia (which Gross and Kowalski (1999) call “nacelles”) are interesting to study since they also form large fascicles of neurites between them that synchronize their activity (Segev et al., 2003), their formation necessarily causes most of the MEA’s electrodes to be wasted because the neurons move away from them as they form clusters.

To take full advantage of the multi-unit philosophy behind MEAs, we would like to have as many neurons in contact with electrodes as possible. In order to prevent clumping of neurons and glia in dense cultures, it is necessary to provide the cells with a very adherent substrate, which they prefer over adhering to each other. We found that treating the MEAs with polyethyleneimine (Lelong et al., 1992) and then laminin allows the cells to grow in a monolayer for months (Potter and DeMarse, 2001). Substituting polylysine for polyethyleneimine was not as effective on our MEAs, which have silicon nitride as the insulating layer. It is likely that each type of MEA insulation material will exhibit different cell adhesion properties, and a variety of treatments should be tried, such as exposure to polycations and extracellular matrix proteins, oxidation by flame or plasma, or covalent modification (Bohanon et al., 1996; Lucas et al., 1986; Stenger et al., 1993; Zeck and Fromherz, 2003).

Our dense monolayer cultures are not strictly monolayers: the glial cells, allowed to multiply and eventually become contact-inhibited, often form a very thin layer under, and sometimes over, the neurons and their processes. By scanning labeled cells with a two-photon fluorescence microscope, we have observed that our cultures are 15 to 20 μm thick, and the neuron somata form a monolayer. We make no effort to inhibit the growth or division of glial cells, because they contribute to neuron survival and synapse formation as mentioned above, and because they do not seem to impede the electrodes’ ability to stimulate or record neurons. Some electrodes on MEAs become damaged after repeated plating of cultures, from deterioration of the contact pads, titanium nitride electrode surface, or silicon nitride insulation. We routinely have neural activity on every MEA electrode that is physically intact.

9.2.2 *Sealed Dishes for Long-Term Cultures*

9.2.2.1 Why Neuron Cultures Die

Primary neuron cultures typically survive for less than a couple of months. One common cause of death is obvious: infection. Mold spores are ubiquitous. The air gap in most culture dishes that allows exchange of oxygen and carbon dioxide also

allows airborne pathogens to contaminate cultures. The warm humid environment of the incubator is itself often the source of mold, and when one dish gets infected, others are likely to suffer a similar fate. To bring cultures out of the incubator for imaging and manipulations in a typical nonsterile laboratory environment puts them at further risk of infection.

There is another common, but much less obvious, reason why primary neuron cultures die: changes in osmolarity. The humid environment in the incubator is supposed to prevent evaporation of cell culture media. However, unless the incubator door is never opened, the mean humidity inside is substantially less than 100%. In a busy lab where the incubator is opened often, it may be in the 80% range, and cultures suffer from hyperosmolarity due to evaporation. This problem is made worse by the tradition of feeding neural cultures by replacing only half the medium with fresh medium; hyperosmolarity persists even after feeding. This causes a gradual deterioration and death of neurons that is usually considered to be “normal” for neural cultures because checking medium osmolarity is not commonly done.

9.2.2.2 Gas-Permeable Membrane Dishes

We developed a new culturing method that greatly reduces the occurrence of both of these problems, and has allowed us to maintain several neural cultures for well over a year (Potter and DeMarse, 2001), and for over two years in one case. The culture dishes are hermetically sealed with a Teflon membrane (Figure 9.2), fluorinated ethylene-propylene, 12.5 μm (Dupont). Although this membrane has no pores (thus preventing infection), it is quite permeable to some small molecules, notably oxygen and carbon dioxide. It is hydrophobic, and thus relatively impermeable to water and water vapor. (Note the difference from Gore-Tex Teflon material used on rain gear, which has pores, and *is* permeable to water vapor.) This allows us to culture our cells in an incubator maintained at 65% relative humidity. A dry incubator full of sealed dishes never needs to be cleaned or sterilized. Of great interest to us and others using MEAs, the low humidity allows putting expensive electronics inside the incubator without the risk of damage by water condensation (Figure 4.6, Section 4.3). Sealed dishes can be repeatedly removed from the incubator for imaging or MEA electrophysiology without fear of contamination.

The membrane slows the shift in pH of carbonate-buffered media caused by removal from an incubator with 5% CO_2 atmosphere, by about a factor of two compared to a standard culture dish with an air gap (Potter and DeMarse, 2001). In practical terms, this allows 30 min to 1 h of experimentation outside the incubator before the medium must be exchanged or re-equilibrated with 5% CO_2 . The membrane is transparent and amenable to imaging on an upright microscope. The sealed MEA cultures are fed by removing the special Teflon lids (available from ALA Scientific) in the laminar flow hood, and replacing all (not half) of the medium, approximately once per week. Of course, proper sterile technique must

be utilized during feeding. If one culture does become infected, the use of sealed dishes prevents mold from spreading to others in the incubator.

The mean concentration of oxygen in the brain is far less than the atmospheric 20%, and usually between 1 and 5% depending on brain region (Studer et al., 2000). Presumably neurons in a 20% O₂ incubator are suffering from an unnatural level of oxidative damage. Therefore, we also routinely use an atmosphere brought to 9% oxygen by injecting pure nitrogen, because this has been shown to enhance survival of primary hippocampal cultures (Brewer and Cotman, 1989).

9.2.2.3 Transporting Live MEA Cultures

It was necessary to ship a number of young and elderly MEA cultures from Caltech to Georgia Tech when the Potter group moved to Atlanta in 2002. We created a system that resulted in viable firing cultures after a transcontinental FedEx journey. The most important consideration is that the cultures not be exposed to turbulence of the medium, which might tear them from the substrate. We made lids consisting of a 2 mm layer of Sylgard (silastic rubber, Dow Corning) on a glass microscope slide, and pressed the Sylgard layer against the glass ring of the MEA after overfilling the dish with degassed Hibernate medium (Brewer and Price, 1996), being very careful that no air remained in the dish. As long as there are no bubbles in this rigid vessel, neural cultures are extremely resistant to damage by shock (g-forces). The cultures were placed in styrofoam boxes with 4°C cold packs, to reduce metabolism. Hibernate medium is buffered for ambient carbon dioxide levels, so no special consideration for maintenance of pH was made. Upon arrival, the medium was replaced with standard serum-containing medium (Potter and DeMarse, 2001), and the glass/Sylgard lids were replaced with Teflon membrane lids. Neural activity was recorded as soon as the cultures warmed up and equilibrated to the culture medium.

9.3 Real-Time Data Processing

Most experimental scientists would prefer to have the results of their experiment as soon as possible, ideally, with intermediate results appearing on the computer screen even before the experiment is over. This allows the experimenter to make the best of unexpected contingencies that might warrant a redirection of efforts. For experiments lasting days, as ours often do, online data processing becomes crucial. Most electrophysiology and imaging systems incorporate this “online” philosophy. But online is not necessarily real-time, although the terms are often used interchangeably. By “real-time,” we mean systems in which results are available in milliseconds, and in which maximum delays are known and guaranteed. Real-time systems are necessary for closed-loop electrophysiology. Because real-time systems are not part of commercially available MEA setups, and we need real-time feedback for our closed-loop paradigm, we developed our own system.

9.3.1 Preventing Data Glut

Multi-electrode arrays are capable of generating large amounts of data in a short period; sampling each of 60 channels at 25 kHz creates a data stream of several megabytes per second, or tens of gigabytes in one afternoon. Even with the plummeting cost of disk storage, the largest affordable current storage systems are taxed by continuous recording for days. Clearly, some data reduction strategy is necessary, and this usually takes the form of extraction of spikes from the raw data stream.

It is assumed by most MEA users that neural signals smaller than action potentials, such as postsynaptic potentials, are hidden in the noise of an extracellular recording, so it makes sense only to record action potentials. This is not strictly true. We have observed that noise levels are higher when recording from a living culture than from a clean MEA with just medium in it. By blocking all sodium-channel-dependent activity in an MEA culture with tetrodotoxin, we verified that indeed, many of the tiny voltage peaks often called “noise” are actually biogenic, and may include both subthreshold depolarizations as well as spikes recorded from neurons at some distance from the electrode. This exercise is helpful in setting an appropriate threshold for detection of action potentials, which usually ranges from four to six times the standard deviation of the background signal, depending on our desire to include or exclude questionable peaks.

9.3.1.1 MeaBench Software

Wagenaar developed a suite of software modules to allow us to do MEA recording and stimulation in real-time. It was necessary to move from Windows to the Linux operating system, because Windows does not allow tight enough control of low-level interrupts that may disrupt real-time processing. Our philosophy is to make MeaBench open-source, extensible, and scriptable with a standard UNIX command-line interface. Data streams or files from MeaBench modules can also be sent to other software packages, such as MatLab (Table 9.1).

Spike-detection thresholds may change during the course of an experiment, especially long experiments where the physical relations between neurons and electrodes are changing due to cell growth. MeaBench has a module that adaptively adjusts the spike-detection threshold to track the RMS background signal, per electrode, on a time-constant of one second. To reduce the number of unusually large noise peaks detected as spikes, and to avoid counting a multiphasic spike as two or more spikes, a simple spike shape criterion (suggested by P. P. Mitra) is applied to all candidate spikes. MeaBench saves 3 ms of “context” centered on the peak of each identified spike, that is, a small segment of raw data, which may be used later for spike sorting. By saving just the spike time, electrode number, and context for each spike detected, and discarding the rest of the data, we usually reduce data storage requirements by a factor of a hundred, compared to saving raw data. Each spike saved with context is 164 bytes. Another order of magnitude of data reduction can be realized by forfeiting spike waveforms and only saving spike

TABLE 9.1. MeaBench features.

-
- Modular, open-source design
 - Direct streaming of raw electrode voltage traces to disk
 - Real-time artifact suppression
 - Real-time spike detection, with easily replaceable detection algorithms
 - Playback of raw electrode traces and recorded spikes at any speed
 - Online visualization of electrode traces, with:
 - Markers for spikes
 - Variable window size
 - Option to trigger off auxiliary channel pulses
 - Unique “scrollback” buffer, to allow closer inspection of interesting recent events
 - Online generation of raster plots, in three ways:
 - One raster showing all activity (accumulated over electrodes)
 - 8×8 geometric display of recording electrodes
 - One raster per stimulus type
 - Raster plot can be scrolled back during experiment to view any previous interval
 - Online sonification of spikes (stereo and tonal mapping)
 - A suite of MatLab functions to import data from MeaBench to MatLab
 - An easy interface that allows any program to hook on to any of MeaBench’s data streams
 - Linux OS, command-line based, modules are easily scripted
-

times and electrode numbers (16 bytes/event). By contrast, compression of raw data files by lossless algorithms such as the UNIX *gzip* (LZ77) function typically only reduces file size by a factor of two. Part of dealing with data glut is to decide carefully beforehand which features of the recordings are crucial for the questions being asked, and which are not.

9.3.1.2 Spike Sorting

In our dense cultures, each electrode records signals from three to ten neurons, and typically there are few inactive electrodes. Thus, spike sorting in real-time is still not feasible on a desktop computer without custom digital signal-processing hardware, such as Plexon’s Multichannel Acquisition Processor, designed for that purpose. For that reason, we usually combine the activity from several neurons recorded by any given electrode, in defining a relevant motor pattern. Spike sorting is still fraught with uncertainty due to spike overlap, especially during bursts, and variability due to cellular changes and noise. Recent theoretical advances such as noise modeling, clustering, and decomposition show promise of overcoming these problems (Shoham et al., 2003) but will be difficult to implement in real-time (millisecond latency) on desktop computers for more than a few channels.

9.3.1.3 Line-Noise Filter

Even with careful attention to grounding and shielding, it is usually not possible to eliminate all interference from devices running off the mains power. MeaBench has a module that can eliminate interference from the mains, including 60 Hz (or 50 Hz in Europe), and its harmonics. Note that it is a bad idea to use a 60 Hz notch filter for this purpose, because the interference is seldom, if ever, a pure sine

wave, and because much of the spectral content of interesting neural signals is near 60 Hz, such as gamma oscillations (Cunningham et al., 2003). We use a simple home-made mains thresholder consisting of two diodes and a voltage-divider, that plugs into the mains and sends out low-voltage transitions in phase with the 60 Hz high voltage. This is recorded on one of the analog input channels along with the neural data. MeaBench creates a template of one period of the mains-related interference on each electrode, by averaging many 60 Hz periods, triggered by the mains thresholder. This template is then subtracted from the incoming data stream from each electrode. This approach could be adapted to other periodic noise sources present in some lab environments, such as motors.

9.4 Stimulation Systems

Speaking in terms of embodied networks, the input or sensory side of MEA technology is not as technically well-developed as the output or motor side. But it is equally important in our closed-loop paradigm. The very same MEA electrodes used for recording can (and should) also be used to stimulate neurons.

9.4.1 *Optimization of Stimulation Parameters*

A wide variety of pulse shapes have been used by the groups mentioned in Section 9.1 to evoke neuronal activity through MEA electrodes, including voltage-controlled and current-controlled pulses. How does one decide what kind of stimulation to use? We have systematically investigated the efficacy of both types of stimuli as a function of amplitude and pulse width (Wagenaar et al., 2004). We found that in most cases negative current pulses are what excite neurons to fire action potentials. That does not mean that voltage-controlled pulses are not useful. In fact, positive-first biphasic voltage-controlled pulses were the most effective stimuli in our repertoire. This is easily understood if one realizes that the sharp downward voltage transient between the two phases corresponds to a strong negative current pulse.

Researchers have traditionally preferred current-controlled stimuli, because their efficacy appears more amenable to modeling (Buitengeweg et al., 2002), as the electric field and potential resulting from stimulation are directly proportional to the current passing through an electrode. There are, however, significant advantages to using voltage control: not only is the circuitry needed to control voltages simpler, but more important, with voltage control it is possible to avoid electrochemical reactions. Current-controlled stimuli can easily exceed voltages that can damage electrodes and harm neurons; this becomes significant when electrode voltages exceed one volt. Although the damage can be reduced by employing charge-balanced pulses, it is still desirable to avoid it altogether. Under current control, this is only possible if the impedance spectra of all the electrodes in the array are known. Moreover, the key advantage of current control—the ability to calculate the resulting electric field in the medium surrounding the




FIGURE 9.3. Short-latency action potentials in response to different stimuli. Each dot in a raster plot is one action potential induced by a biphasic pulse delivered to electrode 45 at time zero. As voltage was increased from zero to 1 volt, more responses were observed across the network. (Reprinted from Wagenaar et al. (2004) with permission from Elsevier.)

electrode—is compromised in MEAs because leakage currents through the insulation layer can reduce the current actually passing through the electrode by as much as 30% depending on the integrity and thickness of the insulation. Consequently, we use voltage-controlled positive-first biphasic pulses of less than 1 volt.

Pulse amplitude is the main determinant of stimulus efficacy (Wagenaar et al., 2004) (Figure 9.3). The number of cells directly stimulated by a given voltage-controlled pulse grows linearly with the amplitude of that pulse. Cultured networks are not as sensitive to the width of voltage pulses. Of course, the pulse must be wide enough to allow the cell membrane and all the parasitic capacitances in the system time to charge, (approximately 400 microseconds in our system). Increasing the width beyond 400 microseconds has little effect on neural response.

In Section 9.4.3 we describe two MEA stimulation systems designed with slightly different philosophies; the RACS by Wagenaar (Wagenaar and Potter, 2004), and the 64-CNS by DeMarse. But first we consider ways to deal with stimulation artifacts.

9.4.2 *Stimulation Artifacts and Software Solutions*

Anyone who uses MEA electrodes to stimulate and record neural activity knows the two are difficult to combine. Multi-electrode arrays typically record signals of 10 to 100 μV , whereas stimuli are on the order of one volt. The large disparity between these two voltage ranges means that artifacts resulting from stimulation can easily swamp out recorded action potentials, even on electrodes distant from the stimulation site.

Several factors contribute to stimulation artifacts (Grumet et al., 2000). A combination of capacitive crosstalk between electrode traces and conduction through the culture medium couples the stimulated electrode to all of the other recording electrodes. If the resulting transient is larger than the dynamic range of the amplification system—as is often the case—the nonlinear properties of saturated amplifiers and the connected filters greatly increase the size and duration of the artifact. We often observe cross-channel stimulus artifacts of several hundred microvolts lasting tens of milliseconds. Because many neurons respond to stimulation within this period, it is important to be able to record immediately after stimulation, so artifact suppression is essential.

Some hardware approaches based on track-and-hold (active suppression) circuits have been successful (Jimbo et al., 2003; Novak and Wheeler, 1988), but are not commercially available. One recently developed commercially available hardware solution (MEA1060-BC by Multi Channel Systems, which we have not tried) grounds the amplifiers during stimulation, and is claimed by the manufacturer to “completely remove any stimulus artifacts.” However, without a sample-and-hold circuit, the DC offset voltage often found on MEA electrodes can itself produce a substantial artifact when the amplifier is switched back into the circuit (Jimbo et al., 2003). For stimulation systems that don’t already have active artifact suppression in hardware, software solutions can be easily and cheaply applied.

We developed a software approach to artifact suppression that allows us to record within 1 to 2 msec after stimulation from all but the stimulated electrode itself (Wagenaar and Potter, 2002). Artifacts depend on the stimulation history and on the individual recording electrode, so each artifact tends to be different from any other artifact. This makes template-based algorithms (which subtract a fixed model of the artifact) perform poorly. We found that the shape of the artifacts is not well modeled even by a variable exponential decay. Instead, our SALPA* algorithm models each individual artifact, in real-time, by a curve constructed from locally fitting polynomials to the recorded signals. Subtraction of this model

* Suppression of Artifacts by Local Polynomial Approximation.




FIGURE 9.4. MEA recordings without (dotted traces and insets) and with SALPA (solid traces) applied, to practically eliminate stimulation artifacts on nonstimulated channels within one millisecond after amplifier becomes unpegged from its supply rails. (Reprinted from Wagenaar and Potter (2002) with permission from Elsevier.)

from the recording leaves an artifact-free signal in which spikes can be detected by voltage thresholding (Figure 9.4).

By leveling large artifacts upon which tiny action potentials ride, SALPA enables their detection an order of magnitude sooner, and can be used in online closed-loop systems in which the spikes detected are used to trigger stimuli. This opens up a new window for studying stimulus responses in culture, at latencies of only a few milliseconds. In particular, spikes from neurons directly stimulated by the electric pulse, without synaptic communication, became detectable for the first time (Figure 9.5). This revealed that a pulse delivered to a single electrode directly stimulates cells with arborizations covering the entire electrode array (Figure 9.3).

9.4.3 Stimulation Hardware

We developed two stimulation systems that can be plugged directly into the MEA1060 preamplifier (Multi Channel Systems) to enable one to stimulate any of the 60 electrodes of the MEA substrate. Commercial stimulation systems to date have been limited to small numbers of electrodes (usually fewer than ten), and required fully specifying a stimulation protocol ahead of time. We enabled short-latency stimulation feedback based on recorded signals. Our systems were designed with the idea of delivering a wide variety of spatiotemporal stimulus patterns to a cultured network, while recording continuously. Rapid switching between stimulation and recording on any electrode is crucial, especially if stimuli are to be triggered by the neural activity itself. Previous systems required manual plugging and unplugging of wires, to stimulate different electrodes. Ours use multiplexing of stimuli to isolation switches. If isolation switches are included in a

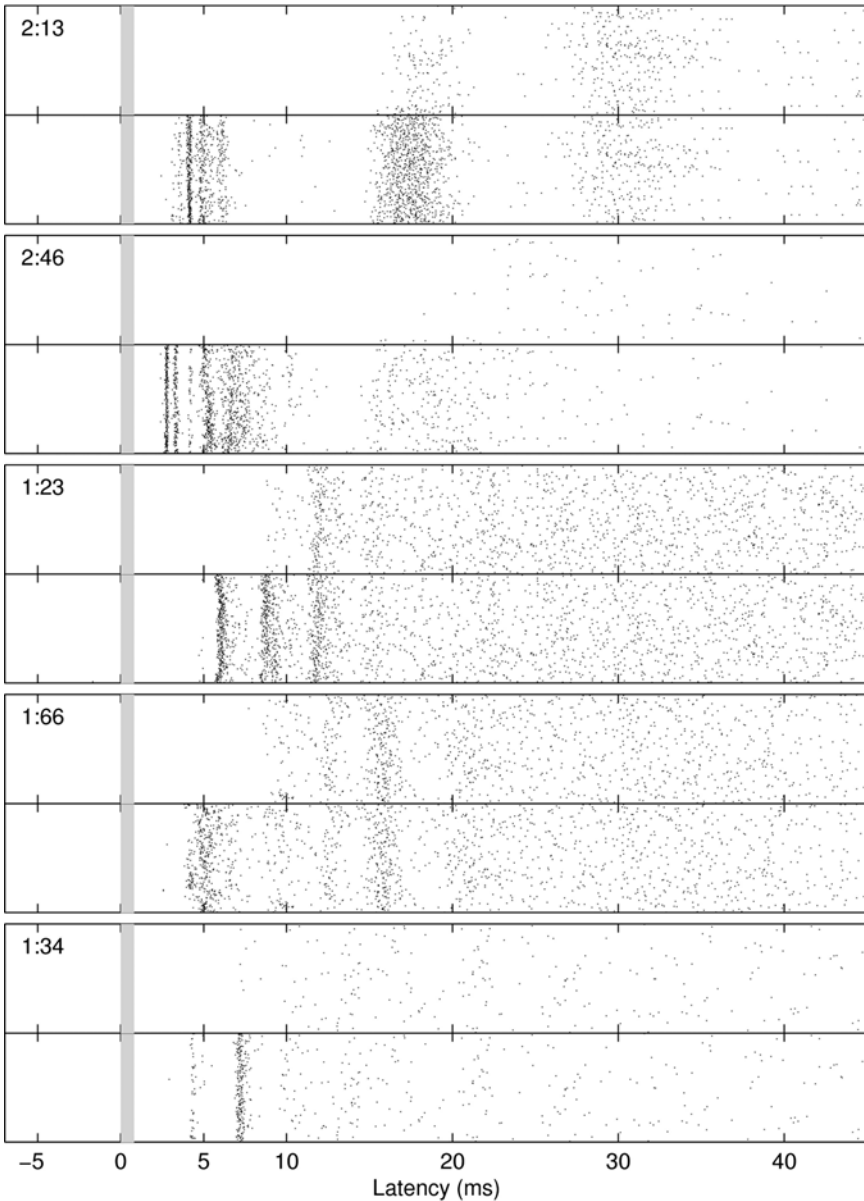


FIGURE 9.5. SALPA allows recording of action potentials (dots in raster plots) at shorter latencies than if no artifact subtraction is used. Two MEAs and five different electrodes' responses are shown (dish:electrode) without (top raster) or with (bottom raster) SALPA applied to raw voltage traces before spike detection by thresholding (at five times RMS noise, without stimuli).

stimulator design, it is important to locate these switches close to the MEA, because any low-impedance path away from the MEA can serve as an antenna, bringing unacceptable noise into recordings. Thus, we put the switches as close as possible to the MEA in both of our simulators, the RACS and the 64-CNS. Switches were carefully selected to have good isolation characteristics, low leakage current, and small charge injection.

9.4.3.1 The RACS: Real-Time All-Channel Stimulator

The philosophy behind the RACS (Wagenaar and Potter, 2004) was to make an inexpensive, flexible real-time stimulation system controlled by a dedicated low-end computer running the RT Linux real-time operating system. RACS is modular, with four banks of 16 stimulator lines that plug directly into each edge of the MEA1060 preamplifier (Figure 9.6, left). Adaptation to other recording systems (including in vivo probes) should be straightforward, provided they allow direct electrical access to the electrodes. The RACS main board (Figure 9.1) connects to the parallel port of a PC, and comprises a digital-to-analog converter and multiplexer. This routes stimuli and switching signals to the four modules that plug into the recording system. There are also auxiliary analog and digital outputs that may be used to trigger other lab equipment or to encode experimental parameters. The PCB layout uses standard 0.1 in. DIP components, and is publicly available (see Wagenaar and Potter (2004)). It can be replicated at an estimated cost of US\$ 250 and about one day of work.

Controlling the stimulator is extremely flexible: stimulation sequences can be generated using Perl scripts. By making such scripts read from a MeaBench spike stream, stimuli can be made contingent on the firing pattern of the culture. Because




FIGURE 9.6. All-channel stimulation systems for Multi Channel Systems preamplifiers. The RACS system (left) has a modular design, with one switching circuit plugged in to each side of the preamplifier. It features simple 0.1 in. DIP construction, with complex stimulus patterns being controlled by a separate computer running RT-Linux. The DAC and multiplexers are on a separate board (shown in Figure 9.1). The 64-CNS for “inverted microscope” amplifier (center) and “upright microscope” amplifier (right) has an onboard USB-interfaced microprocessor for control of complex stimulus patterns, and features surface-mount components for ease of access to the MEA by micro-pipettes.

the stimulator can switch between electrodes with microsecond timing, it is possible to stimulate using arbitrarily complex multi-channel patterns. Thus, for the first time, stimulating neuronal ensembles in culture with naturalistic patterns across all electrodes of the MEA is possible. Spatiotemporal analysis of multi-unit recordings is now common, thanks to MEAs and multi-wire probes. But spatiotemporal distributed stimulation is virtually nonexistent in the literature. One notable exception is from Heck (1995), who used distributed patterns of stimulation (across 11 wire electrodes) to test circuit hypotheses in cerebellar slices. This may be a far better way to probe the information-processing capacity of neuronal networks than sending in stimuli on single electrodes.

9.4.3.2 The 64-CNS: 64-Channel Neural Stimulator

Like the RACS, the 64-CNS was also designed to allow a host computer to dynamically stimulate electrodes on a MEA in patterns that can be modified in real-time in response to the ongoing neural activity. Unlike the RACS, it has everything on one low-profile PC board with surface-mount components (Figure 9.6, center and right). The system features an onboard 8 MHz micro-controller, obviating the need for the RT Linux box used by RACS. The micro-controller schedules the timing and delivery of biphasic voltage stimulation pulses to each of the 60 channels on the MEA. This relieves the host computer from the micro-second timing requirements needed to deliver patterned stimulations to the MEA, which is difficult to accomplish on common multi-tasking computers.

The board plugs directly into a version of the MEA1060 preamp built with headers in place of the grounding DIP switches. It has a programmable voltage source and low-noise programmable switches that control the delivery of stimulation pulses to any of the 60 channels on the MEA. A blanking system is used to ground the recording amplifiers during the delivery of stimulation pulses, reducing stimulus artifacts with the limitations described above. The micro-controller receives commands from a host computer via a high-speed USB serial interface. These commands are stored on the microprocessor's program stack and executed when triggered by the host. Stimuli of the same amplitude can be delivered to any number of electrodes simultaneously. Stimuli of different amplitudes can follow each other by as little as 500 μ s.

Our hardware and software systems allowed neural activity (analogous to motor commands) to be processed rapidly and used to trigger electrical stimuli via substrate electrodes (analogous to sensory inputs). Using MeaBench, and the RACS stimulation system, our 15 ms loop time includes 60-channel recording, spike extraction, pattern detection, and the triggering and delivery of stimuli to multiple electrodes (Wagenaar and Potter, 2004). This is at the fast end of sensory motor loops in mammals. Activity patterns to be used as triggers may include bursts on certain channels, a vector sum of activity across the array (Lukashin et al., 1996), or the occurrence of precisely timed spatiotemporal sequences of action potentials (Nadasdy, 2000). To help decide which neural activity patterns might serve as effective motor commands, we use a number of online visualization tools included

in the MeaBench suite, such as raw voltage traces, raster plots, burst analysis, and stereo sonification of spike data.

9.4.4 *The Future of MEA Stimulation*

The language that neurons use to communicate with each other is chemical, that is, neurotransmitters. Extracellular electrodes are able to stimulate neurons by the lucky fact that part of the communication process within a neuron is electrical, that is, the opening of voltage-sensitive ion channels. Although it is possible to induce action potentials by electrically depolarizing the excitable membrane of neurons, it would be more natural to induce them via the very chemicals that neurons use to induce them. Some advances have been made in this direction. One approach is to include into the MEA a micro-fluidic system for the localized delivery of neurotransmitters or other neuroactive compounds (Heuschkel et al., 1998). Another approach is to apply neuroactive compounds via a micro-manipulated puffer pipette (Liu and Tsien, 1995). One could potentially address more locations in a network with a scanning laser beam than a micro-pipette. A pulsed infrared laser can be used to “uncage” neurotransmitters (or agonists) that are only pharmacologically active after photolysis of an attached caging group. Denk et al. (1994) used this method to map acetylcholine receptor distribution on a single neuron, by recording the resulting whole-cell current when the agonist carbachol was photo-uncaged next to it. For any of these techniques to find routine applicability to many neurons in a MEA culture, we need advances in development of micro-fluidic structures and in photo-uncagable neuroactive molecules (Furuta et al., 1999). Our collaborators, Ari Glezer and Bruno Frazier at Georgia Tech, have begun building MEA systems with incorporated micro-fluidics, as part of an NIH Bioengineering Research Partnership. Similar projects are underway in Europe (Ziegler, 2000).

9.5 Closed-Loop Multi-Unit Electrophysiology

Closed-loop electrophysiology, where stimulation is contingent on what is recorded, is well established as a useful tool for studying single-neuron and small-circuit properties, using artificial conductance injection with a dynamic clamp (Nowotny et al., 2003; Raikov et al., 2004; Sharp et al., 1992, 1993; Suter and Jaeger, 2004). Glass micro-electrodes are used for both recording and injecting of currents in one or more nearby cells, under the control of a model for some dynamic property of the cell or network. This model may be implemented in software (Kullmann et al., 2004) or hardware (Raikov et al., 2004). This technique was a natural extension of the idea of voltage- (or current-) clamp recording, which dates back to 1948 (Huxley, 2002), except that instead of keeping some cellular parameter constant, it is varied dynamically according to the cell’s behavior and the model. The speed of the feedback for these experiments is in the tens of kHz range, corresponding to loop times in the tens of μs range. The need for very fast

feedback, and the fact that the technique requires delicate micro-manipulation of electrodes onto neurons, makes it difficult to use for more than a few cells at a time.

Closed-loop multi-unit electrophysiology, by contrast, is only just now coming of age (Ananthaswami, 2002). Because MEAs have extracellular electrodes, it is not feasible to do the sort of detailed conductance-model feedback experiments for which the dynamic clamp has been used. Instead, the recordings and stimuli focus on action potentials, the presumed currency of information transfer in the brain. One type of closed-loop multi-unit electrophysiology is to record from an electrode array implanted in an animal's brain, and to feed back perceptual stimuli through the animal's natural senses. In this approach being pursued by several groups, a monkey or rat is implanted with cortical electrodes (Carmena et al., 2003; Chapin et al., 1999; Taylor et al., 2002). Recorded neural signals control a robotic arm while the animal moves its own arm, and information about the prosthetic arm's movement is fed back to the animal through its eyes: either by watching the prosthetic arm, or by watching moving shapes on a video screen. The animal eventually learns to move the shapes (and the robot arm) without even moving its own arm. In the first closed-loop study to use what we call *hybrots* (which they call "neurobots") (Kositsky et al., 2003; Reger et al., 2000), Mussa-Ivaldi and Miller (2003) used an acute slice from a lamprey brainstem to control a Khepera wheeled robot. They mapped a circuit that normally processes vestibular information to a phototropism task: the robot moved towards or away from a light. Sensory input (from the Khepera's light sensors) was delivered to the slice via two tungsten wire electrodes, and motor commands were recorded by two glass extracellular micro-pipette electrodes. All of these closed-loop electrophysiology experiments are important steps toward studying distributed processing in embodied, situated neural systems.

9.5.1 *Neurally Controlled Animats*

Animals are situated and embodied. We want to enable the study of sensory-motor learning in cultured networks. If learning is defined as a process by which experience or practice results in a relatively permanent change in behavior (Morris, 1973), then to learn, a system must have a body to behave with, and an environment in which to behave. We have used the software and hardware systems described above to re-embodiment cultured cortical networks. We created a virtual environment and a very simple embodiment on the computer, as the first neurally controlled animat (DeMarse et al., 2001; Potter et al., 1997). (An animat is any simulated or robotic animal (Meyer and Guillot, 1994).) It was a neuroethology experiment, in the sense that we did not set the animat to any particular task, but merely observed the effect of the feedback stimulation on its behavior. We used an artificial neural network to cluster firing rate data in a high-dimensional space, and classify recurring patterns (DeMarse et al., 2001). We found that the diversity of activity patterns expressed

by cultured cortical networks was enhanced by real-time feedback stimulation, at least while the sensory–motor loop was closed.

We did not observe any evidence of lasting (more than 30 min) changes in the open-loop behavior (driven by spontaneous activity) in the animat, as a result of closed-loop sessions. In a more disembodied closed-loop study by Shahaf and Marom (2001), cortical networks cultured on MEAs were stimulated until they satisfied a “learning” criterion of increased firing at 50 ± 10 ms latency after a probe stimulus. In this case, the turning off of a periodic stimulus (delivered via two substrate electrodes) was the only stimulus parameter contingent on the multi-unit recordings. They propose that the stimuli serve as an exploratory driving force, which is supported by our observation that with feedback stimuli, the cultures expressed more differentiable activity patterns. They hypothesize that by turning off the stimuli upon achievement of the learning criterion, the most recent pattern is “selected” by the network, this being an adaptive or desired response (Shahaf and Marom, 2001).

9.5.2 *Hybrots*

There are a number of reasons to use physical embodiments (robots) for animats (Holland and McFarland, 2001; Webb, 2002). Simulating the mechanics of the real world to any degree of realism is computationally difficult, but with robots, you get the physics “for free.” It is becoming more and more clear that in animals, the physics of their bodies and interactions with the environment do a large amount of the sensory–motor processing that might previously have been attributed to neural systems alone. This is the premise of the interdisciplinary field of embodied cognition (Clark, 1997). Using robots also forces researchers to apply themselves to real-world problems, even if at a simplistic level.

We have used MEA cultures to control several hybrots in a closed-loop paradigm (Bakkum et al., 2004). One of these was the Koala 6-wheeled rover (K-Team, Figure 9.1). Under control of the neuronal network (the hybrot’s “brain”), the Koala (the hybrot’s body) was commanded to approach and then follow another robot being controlled randomly by the computer. For the control, Alec Shkolnik used a reproducible network property, the reduction or enhancement in dishwide response to the second of two stimulus pulses, depending on the interpulse interval (IPI). At short IPIs (± 20 msec), the response is maximal, roughly the sum of both responses when each stimulus is delivered alone. But at 100 to 300 msec IPI, the network is still in a refractory state from the first stimulus, and the response is minimal (Darbon et al., 2002). Beyond 500 msec, the response is intermediate. The magnitude of the Koala’s movement in one feedback cycle is in proportion to this response. The distance to target is encoded in the subsequent IPI. This mapping enabled the Koala to approach a stationary robot, and to follow it at a certain distance when it began to move. Other closed-loop hybrot experiments using cultured networks are under way in Europe, as part of the multi-national EU-funded NeuroBIT project (Martinoia et al., 2004).

9.6 Combining Imaging With MEA Electrophysiology

One of the most important advantages of *in vitro* neural models is that they are readily imaged with the light microscope. MEAs are usually transparent, and can be fabricated with clear leads of indium-tin oxide (Gross et al., 1985; Jimbo and Kawana, 1992) on thin glass for imaging on inverted microscopes. The “body” of re-embodied cultured networks (Bakkum et al., 2004) can move and behave, whereas the “brain” holds still on the microscope stage. Although it is possible to do microscopic imaging of living neurons *in vivo* (Gan et al., 2003; Helmchen et al., 2001; Levene et al., 2004; Majewska and Sur, 2003; Trachtenberg et al., 2002), the animal must be immobilized by physical restraint, paralytics, and/or anaesthesia. In those cases, the animal is incapable of expressing normal behavior, and likely not processing sensory input normally either. Hints at the morphological correlates of learning and memory *in vivo* must be gleaned by imaging before and after the animal learns, or in pairs of animals undergoing different experiences. Historically, this imaging has been done in slices of aldehyde-fixed brain tissue (Burgess and Coss, 1983; Rollenhagen and Bischof, 1994; Weiler et al., 1995), where inferences about activity-dependent morphological dynamics are difficult to make.

By using MEAs in conjunction with microscopic imaging, it is now possible to observe activity-dependent morphological dynamics, at a variety of time scales, while they are happening. With Scott Fraser of the Biological Imaging Center at Caltech, Potter developed new imaging techniques for maintaining mammalian neurons on the microscope stage for days, to allow time-lapse movies of morphological dynamics. These include advances in two-photon microscopy and specimen life support (Potter, 2000; Potter et al., 1996a, 2001a, 1996b, c, 2001b).

9.6.1 *Why Two Photons Are Better Than One*

Two-photon microscopy (Denk et al., 1990) allows repeated imaging of fluorescently labeled neurons with little or no light-induced damage (Potter, 1996; Williams et al., 2001). One of the benefits of two-photon microscopy is wasted on monolayer cultures, namely, the ability to image deeper into thick specimens by using infrared illumination. It is important to realize, however, that the two-photon effect limits excitation of the label to a plane about one micron in thickness (Potter et al., 1996c), and that monolayer cultures of neurons and glia are 15 to 20 microns thick. Thus, even for monolayer cultures, the total light dose is greatly reduced compared to one-photon confocal or wide-field fluorescence microscopy.

Commercially available multi-photon microscopes that are also visible-light confocal microscopes are not as efficient as they could be, because the requirements for confocal imaging are different than for multi-photon imaging (Pawley, 1995; Potter, 2000). Notably, with two-photon excitation it is not necessary to focus the light emitted by the specimen to create an in-focus image. This is because, at

any moment, light is only emitted by a diffraction-limited focal volume of the scanning infrared laser beam. If emitted light is scattered on its way out of the specimen, those photons can be collected by a photomultiplier tube and referred back to the point of excitation to create the image. This makes two-photon imaging inherently more photon-efficient than one-photon imaging, especially when a direct (nondescanned) detector is used to collect as many scattered photons as possible (Wokosin et al., 1998). To overcome the shortcomings of commercially available multi-photon microscopes, we have built a custom two-photon microscope from the ground up, according to the design of Tsai et al. (2001). This design includes direct detection, and a flexible open architecture that can accommodate new equipment as it becomes available.

9.6.2 *Keeping Cells Happy on the Microscope*

Because it is difficult to fit a microscope inside a cell incubator, for long-term imaging it is better to bring the incubator to the microscope. We built inexpensive microscope enclosures out of Reflectix insulation (Figure 9.1, <http://www.reflectixinc.com/>) to warm the culture to rodent body temperature, to block ambient light, and to maintain an atmosphere with 5% carbon dioxide, for pH homeostasis (Potter, 2000). By using Teflon-sealed dishes with “baggy” lids, the objective of an upright microscope can be submerged into the culture medium without compromising sterility or osmolarity. The microscope enclosure serves the additional purpose of preventing focus drift due to changes in room temperature that cause expansion and contraction of microscope components. This fact alone makes it superior to systems that merely warm the MEA itself, when doing time-lapse imaging.

9.6.3 *Variegated Neurons*

Thanks to the use of fluorescent proteins in more and more transgenic constructs (Hadjantonakis et al., 2003), it is easy to prepare MEA cultures from animals that come pre-labeled for fluorescence microscopy. There are now several color variants (xFPs) of the green fluorescent protein from jellyfish (Heim and Tsien, 1996), and red emitting fluorescent proteins from coral (Campbell et al., 2002; Okita et al., 2004). We prepared dissociated cultures from the combined cortices of two or more transgenic and wild-type mice (Feng et al., 2000) to produce cultures in which a small random subset of neurons is fluorescently labeled (Figure 9.7). For dense cultures, a ratio of 1:20 labeled:unlabeled cells allows detailed imaging of neurites without much overlap of labeled cells. Unlike dye labeling, or transient transfection with viruses or plasmids, xFP labeling in neurons from transgenic animals is self-renewing and harmless, allowing individual neurons to be followed *in vitro* for months.




FIGURE 9.7. YFP-labeled neurons on MEAs. Fluorescence (left) and phase-contrast (right) microscopic images from mixed neuron/glia cultures after 3 days in vitro (top) and 50 days in vitro (bottom). In some cases, we mixed cells from transgenic mice in which every cell expressed YFP under an actin promoter with an excess of unlabeled wild-type cortex cells (top). In other cases, we used mice (e.g., Thy1-YFP-H from Feng et al., 2000) that only express YFP in layer 5 pyramidal neurons (bottom). Although the older cultures show vacuoles from cells that had died earlier, many neurons show healthy morphology including complex neurites with dendritic spines and axonal varicosities, as well as robust electrical activity many months after plating.

9.7 Conclusion

By combining MEA electrophysiology with long-term time-lapse imaging, it is possible to make correlations between changes in network function and changes in neuronal morphology. By re-embodiment dissociated cultured networks, network function can be mapped onto behavior, and in vitro research can now make use of a new kind of behavioral studies that include detailed (submicron) imaging not possible in vivo. By closing the sensory–motor loop around MEA cultures, they are more likely to shed light on the mechanisms of learning, memory, and information processing in animals.

Acknowledgments. The authors thank the following for their significant contributions to this work: Funding: Whitaker Foundation, Georgia Research Alliance, NSF Center for Behavioral Neuroscience, NIH National Institute of Neurological Disorders and Stroke, NIH National Institute of Biomedical Imaging and

Bioengineering. Postdocs: Axel Blau and Mark Booth. Graduate students: Radhika Madhavan, Peter Passaro, Zenas Chao, Douglas Bakkum, and Komal Rambani. Undergrads: Mike Atkin, Gray Rybka, Sam Thompson, William Findley, Matthew MacDougall, Miro Dudik, John Brumfield, Blythe Towal, and Alec Shkolnik. Collaborators: Guy Ben-Ary, Jerry Pine, Scott Fraser, Josh Sanes, and David Kleinfeld. Technicians: Pooja Bhatia, Eno Ekong, Sheri McKinney, Dustin Romine, and Yan Wu. Lab managers: Mary Flowers and Bryan Williams. Manufacturer: Multi Channel Systems for help with data acquisition hardware.

References

- Ananthaswami, A. (2002). Mind over metal. *New Sci.* Feb 23: 26–29.
- Bakkum, D.J., Shkolnik, A.C., Ben-Ary, G., Gamblen, P., DeMarse, T.B., and Potter, S.M. (2004). Removing some ‘A’ from AI: Embodied cultured networks. In: Iida, Pfeifer, R., Steels, L., and Kuniyoshi, Y. (Eds.) pp. 130–45. Springer-Verlag, New York.
- Banker, G. and Goslin, K. (1998). *Culturing Nerve Cells, 2nd Edition*. MIT Press, Cambridge, MA.
- Ben-Ari, Y. (2001). Developing networks play a similar melody. *Trends Neurosci.* 24: 353–360.
- Bi, G.Q. and Poo, M.M. (1998). Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci.* 18: 10464–10472.
- Bohanon, T., Elender, G., Knoll, W., Koberle, P., Lee, J.S., Offenhausser, A., Ringsdorf, H., Sackmann, E., Simon, J., Tovar, G., and Winnik, F. M. (1996). Neural cell pattern formation on glass and oxidized silicon surfaces modified with poly(N-isopropylacrylamide). *J. Biomat. Sci. Polymer Edition* 8: 19–39.
- Branch, D.W., Corey, J.M., Weyhenmeyer, J.A., Brewer, G. J., and Wheeler, B.C. (1998). Microstamp patterns of biomolecules for high-resolution neuronal networks. *Med. Biol. Eng. Comput.* 36: 135–141.
- Brewer, G.J. and Cotman, C.W. (1989). Survival and growth of hippocampal-neurons in defined medium at low-density—Advantages of a sandwich culture technique or low oxygen. *Brain Res.* 494: 65–74.
- Brewer, G.J. and Price, P.J. (1996). Viable cultured neurons in ambient carbon dioxide and hibernation storage for a month. *Neuroreport* 7: 1509–1512.
- Buitenweg, J.R., Rutten, W.L.C., and Marani, E. (2002). Extracellular stimulation window explained by a geometry-based model of the neuron-electrode contact. *IEEE Trans. Biomed. Eng.* 49: 1591–1599.
- Burgess, J.W. and Coss, R.G. (1983). Rapid effect of biologically relevant stimulation on tectal neurons—changes in dendritic spine morphology after 9 minutes are retained for 24 hours. *Brain Res.* 266: 217–223.
- Campbell, R.E., Tour, O., Palmer, A.E., Steinbach, P.A., Baird, G.S., Zacharias, D.A., and Tsien, R.Y. (2002). A monomeric red fluorescent protein. *Proc. Nat. Acad. Sci. U. S. A.* 99: 7877–7882.
- Carmenta, J.M., Lebedev, M.A., Crist, R.E., O’Doherty, J.E., Santucci, D.M., Dimitrov, D.F., Patil, P.G., Henriquez, C.S., and Nicolelis, M.A.L. (2003). Learning to control a brain-machine interface for reaching and grasping by primates. *PLoS Biol.* 1: 193–208.

- Chapin, J.K., Moxon, K.A., Markowitz, R.S., and Nicolelis, M.A.L. (1999). Real-time control of a robot arm using simultaneously recorded neurons in the motor cortex. *Nature Neurosci.* 2: 664–670.
- Clark, A. (1997). *Being There: Putting Brain, Body, and the World Together Again*. MIT Press, Cambridge, MA.
- Connolly, P., Clark, P., Curtis, A.S., Dow, J.A., and Wilkinson, C.D. (1990). An extracellular microelectrode array for monitoring electrogenic cells in culture. *Biosens. Bioelectron.* 5: 223–234.
- Cunningham, M.O., Davies, C.H., Buhl, E.H., Kopell, N., and Whittington, M.A. (2003). Gamma oscillations induced by kainate receptor activation in the entorhinal cortex in vitro. *J. Neurosci.* 23: 9761–9769.
- Darbon, P., Scicluna, L., Tschertner, A., and Streit, J. (2002). Mechanisms controlling bursting activity induced by disinhibition in spinal cord networks. *Euro. J. Neurosci.* 15: 671–683.
- DeMarse, T.B., Wagenaar, D.A., Blau, A.W., and Potter, S.M. (2001). The neurally controlled animat: Biological brains acting with simulated bodies. *Auton. Robots* 11: 305–310.
- Denk, W., Delaney, K.R., Gelperin, A., Kleinfeld, D., Strowbridge, D.W., Tank, D.W., and Yuste, R. (1994). Anatomical and functional imaging of neurons using 2-photon laser scanning microscopy. *J. Neurosci. Meth.* 54: 151–162.
- Denk, W., Strickler, J.H., and Webb, W. W. (1990). 2-photon laser scanning fluorescence microscopy. *Science* 248: 73–76.
- Feng, G.P., Mellor, R.H., Bernstein, M., Keller-Peck, C., Nguyen, Q.T., Wallace, M., Nerbonne, J.M., Lichtman, J.W., and Sanes, J.R. (2000). Imaging neuronal subsets in transgenic mice expressing multiple spectral variants of GFP. *Neuron* 28: 41–51.
- Fromherz, P. and Stett, A. (1995). Silicon-neuron junction—Capacitive stimulation of an individual neuron on a silicon chip. *Phys. Rev. Lett.* 75: 1670–1673.
- Furuta, T., Wang, S.S.H., Dantzker, J.L., Dore, T.M., Bybee, W.J., Callaway, E.M., Denk, W., and Tsien, R.Y. (1999). Brominated 7-hydroxycoumarin-4-ylmethyls: Photolabile protecting groups with biologically useful cross-sections for two photon photolysis. *Proc. Nat. Acad. Sci. U. S. A.* 96: 1193–1200.
- Gabbiani, F., Krapp, H.G., and Laurent, G. (1999). Computation of object approach by a wide-field, motion-sensitive neuron. *J. Neurosci.* 19: 1122–1141.
- Gan, W.B., Kwon, E., Feng, G.P., Sanes, J.R., and Lichtman, J.W. (2003). Synaptic dynamism measured over minutes to months: Age-dependent decline in an autonomic ganglion. *Nature Neurosci.* 6: 956–960.
- Gross, G.W. and Kowalski, J.M. (1999). Origins of activity patterns in self-organizing neuronal networks in vitro. *J. Intell. Mater. Syst. Struct.* 10: 558–564.
- Gross, G.W., Rhoades, B.K., Reust, D.L., and Schwalm, F.U. (1993). Stimulation of monolayer networks in culture through thin-film indium-tin oxide recording electrodes. *J. Neurosci. Meth.* 50: 131–143.
- Gross, G.W., Wen, W.Y., and Lin, J.W. (1985). Transparent indium-tin oxide electrode patterns for extracellular, multisite recording in neuronal cultures. *J. Neurosci. Meth.* 15: 243–252.
- Grumet, A.E., Wyatt, J.L., and Rizzo, J.F. (2000). Multi-electrode stimulation and recording in the isolated retina. *J. Neurosci. Meth.* 101: 31–42.
- Hadjantonakis, A.K., Dickinson, M.E., Fraser, S.E., and Papaioannou, V.E. (2003). Technicolour transgenics: Imaging tools for functional genomics in the mouse. *Nature Rev. Genet.* 4: 613–625.

- Hartmann, M.J. and Bower, J.M. (2001). Tactile responses in the granule cell layer of cerebellar folium Crus IIa of freely behaving rats. *J. Neurosci.* 21: 3549–3563.
- Heck, D. (1995). Investigating dynamic aspects of brain-function in slice preparations—spatiotemporal stimulus patterns generated with an easy-to-build multielectrode array. *J. Neurosci. Meth.* 58: 81–87.
- Heim, R. and Tsien, R.Y. (1996). Engineering green fluorescent protein for improved brightness, longer wavelengths and fluorescence resonance energy transfer. *Curr. Biol.* 6: 178–182.
- Helmchen, F., Fee, M.S., Tank, D.W., and Denk, W. (2001). A miniature head-mounted two-photon microscope: High-resolution brain imaging in freely moving animals. *Neuron* 31: 903–912.
- Heuschkel, M.O., Guerin, I., Buisson, B., Bertrand, D., and Renaud, P. (1998). Buried microchannels in photopolymer for delivering of solutions to neurons in a network. *Sens. Actuat. B-Chem.* 48: 356–361.
- Holland, O. and McFarland, D. (2001). *Artificial Ethology*. Oxford University Press, Oxford.
- Huxley, A. (2002). From overshoot to voltage clamp. *Trends Neurosci.* 25: 553–558.
- Jimbo, Y. and Kawana, A. (1992). Electrical stimulation and recording from cultured neurons using a planar electrode array. *Bioelectrochem. Bioenerget.* 29: 193–204.
- Jimbo, Y., Kasai, N., Torimitsu, K., Tateno, T., and Robinson, H.P.C. (2003). A system for MEA-based multisite stimulation. *IEEE Trans. Biomed. Eng.* 50: 241–248.
- Jimbo, Y., Robinson, H.P.C., and Kawana, A. (1998). Strengthening of synchronized activity by tetanic stimulation in cortical cultures: Application of planar electrode arrays. *IEEE Trans. Biomed. Eng.* 45: 1297–1304.
- Kositsky, M., Karniel, A., Alford, S., Fleming, K.M., and Mussa-Valdi, F.A. (2003). Dynamical dimension of a hybrid neurorobotic system. *Trans. Neural Syst. Rehab. Eng.* 11: 155–159.
- Kullmann, P.H.M., Wheeler, D.W., Beacom, J., and Horn, J.P. (2004). Implementation of a fast 16-bit dynamic clamp using LabVIEW-RT. *J. Neurophysiol.* 91: 542–554.
- Latham, P.E., Richmond, B.J., Nirenberg, S., and Nelson, P.G. (2000). Intrinsic dynamics in neuronal networks. II. Experiment. *J. Neurophysiol.* 83: 828–835.
- Lelong, I.H., Petegnief, V., and Rebel, G. (1992). Neuronal cells mature faster on polyethyleneimine coated plates than on polylysine coated plates. *J. Neurosci. Res.* 32: 562–568.
- Levene, M.J., Dombek, D.A., Kasischke, K.A., Molloy, R.P., and Webb, W.W. (2004). In vivo multiphoton microscopy of deep brain tissue. *J. Neurophys.* 91: 1908–1912.
- Liu, G and Tsien, R.W. (1995). Synaptic transmission at single visualized hippocampal boutons. *Neuropharmacol.* 34: 1407–1421.
- Lucas, J.H., Czisny, L.E., and Gross, G.W. (1986). Adhesion of cultured mammalian central nervous system neurons to flame-modified hydrophobic surfaces. *In Vitro Cell Dev. Biol.* 22: 37–43.
- Lukashin, A.V., Amirikian, B.R., and Georgopoulos, A.P. (1996). A simulated actuator driven by motor cortical signals. *Neuroreport* 7: 2597–2601.
- Maeda, E., Robinson, H.P., and Kawana, C.A. (1995). The mechanisms of generation and propagation of synchronized bursting in developing networks of cortical-neurons. *J. Neurosci.* 15: 6834–6845.
- Maher, M.P., Pine, J., Wright, J., and Tai, Y.C. (1999). The neurochip: A new multielectrode device for stimulating and recording from cultured neurons. *J. Neurosci. Meth.* 87: 45–56.

- Majewska, A. and Sur, M. (2003). Motility of dendritic spines in visual cortex in vivo: Changes during the critical period and effects of visual deprivation. *Proc. Nat. Acad. Sci. U. S. A.* 100: 16024–16029.
- Martinoia, S., Sanguineti, V., Cozzi, L., Berdondini, L., van Pelt, J., Tomas, J., Le Masson, G., and Davide, F. (2004). Towards an embodied in vitro electrophysiology: The NeuroBIT Project. *Neurocomput.* 58–60: 1065–1072.
- Merz, M. and Fromherz, P. (2002). Polyester microstructures for topographical control of outgrowth and synapse formation of snail neurons. *Adv. Mater.* 14: 141ff.
- Meyer, J.-A. and Guillot, A. (1994). From SAB90 to SAB94: Four years of animat research. In: Cliff, D.H., Meyer, P., Wilson, J.-A., Cambridge, S.W., eds., *From Animals to Animats 3. Proceedings of the Third International Conference on Simulation of Adaptive Behavior*. MIT Press, Cambridge, MA, pp. 2–11.
- Misgeld, U., Zeilhofer, H.U., and Swandulla, D. (1998). Synaptic modulation of oscillatory activity of hypothalamic neuronal networks in vitro. *Cell. Molec. Neurobiol.* 18: 29–43.
- Morris, C.G. (1973). *Psychology: An Introduction*. Appleton-Century-Crofts, New York.
- Muller, T.H., Misgeld, U., and Swandulla, D. (1992). Ionic currents in cultured rat hypothalamic neurones. *J. Physiol.* 450: 341–62.
- Mussa-Ivaldi, F.A. and Miller, L.E. (2003). Brain-machine interfaces: Computational demands and clinical needs meet basic neuroscience. *Trends Neurosci.* 26: 329–334.
- Nadasdy, Z. (2000). Spike sequences and their consequences. *J. Physiol.-Paris* 94: 505–524.
- Nam, Y., Khatami, D., Wheeler, B.C., and Brewer, G.J. (2003). Electrical stimulation of patterned neuronal networks in vitro. *IEEE Eng. Med. Biol.*, Cancun, Mexico.
- Nolfi, S. and Parisi, D. (1999). Exploiting the power of sensory-motor coordination. *Advances in Artificial Life, Proceedings*. 1674: 173–182.
- Novak, J.L. and Wheeler, B.C. (1988). Multisite hippocampal slice recording and stimulation using a 32 element microelectrode array. *J. Neurosci. Meth.* 23: 149–59.
- Nowotny, T., Zhigulin, V.P., Selverston, A.I., Abarbanel, H.D.I., and Rabinovich, M.I. (2003). Enhancement of synchronization in a hybrid neural circuit by spike-timing dependent plasticity. *J. Neurosci.* 23: 9776–9785.
- Oka, H., Shimono, K., Ogawa, R., Sugihara, H., and Taketani, M. (1999). A new planar multielectrode array for extracellular recording: application to hippocampal acute slice. *J. Neurosci. Meth.* 93: 61–67.
- Okita, C., Sato, M., and Schroeder, T. (2004). Generation of optimized yellow and red fluorescent proteins with distinct subcellular localization. *Biotechniques* 36: 418ff.
- Pawley, J.B. (Ed.) (1995). *Handbook of Biological Confocal Microscopy*, 2nd Edition. Plenum Press, New York.
- Potter, S.M. (1996). Vital imaging: Two photons are better than one. *Curr. Biol.* 6: 1595–1598.
- Potter, S.M. (2000). Two-photon microscopy for 4D imaging of living neurons. In: Yuste, R., Lanni, F., and Konnerth, A., eds., *Imaging Neurons: A Laboratory Manual*. CSHL Press, Cold Spring Harbor, NY, 20.1–20.16.
- Potter, S.M. and DeMarse, T.B. (2001). A new approach to neural cell culture for long-term studies. *J. Neurosci. Meth.* 110: 17–24.
- Potter, S.M., Fraser, S.E., and Pine, J. (1996a). The greatly reduced photodamage of 2-photon microscopy enables extended 3-dimensional time-lapse imaging of living neurons. *Scanning* 18: 147.

- Potter, S.M., Fraser, S.E., and Pine, J. (1997). Animat in a petri dish: Cultured neural networks for studying neural computation. *Proceedings of the 4th Joint Symposium on Neural Computation, UCSD*: 167–174.
- Potter, S.M., Lukina, N., Longmuir, K.J., and Wu, Y. (2001a). Multi-site two-photon imaging of neurons on multi-electrode arrays. *SPIE Proceedings* 4262: 104–110.
- Potter, S.M., Pine, J., and Fraser, S.E. (1996b). Neural transplant staining with DiI and vital imaging by 2-photon laser-scanning microscopy. *Scan. Microscopy Supplement* 10: 189–199.
- Potter, S.M., Wang, C.M., Garrity, P.A., and Fraser, S.E. (1996c). Intravital imaging of green fluorescent protein using 2-photon laser-scanning microscopy. *Gene* 173: 25–31.
- Potter, S.M., Zheng, C., Koos, D.S., Feinstein, P., Fraser, S.E., and Mombaerts, P. (2001b). Structure and emergence of specific olfactory glomeruli in the mouse. *J. Neurosci.* 21: 9713–9723.
- Raikov, I., Preyer, A., and Butera, R.J. (2004). MRCI: A flexible real-time dynamic clamp system for electrophysiology experiments. *J. Neurosci. Meth.* 132: 109–123.
- Ramakers, G.J.A., Raadsheer, F.C., Corner, M.A., Ramaekers, F.C.S., and Vanleeuwen, F.W. (1991). Development of neurons and glial-cells in cerebral-cortex, cultured in the presence or absence of bioelectric activity—morphological observations. *Euro. J. Neurosci.* 3: 140–153.
- Rambani, K., Booth, M.C., Brown, E.A., Raikov, I., and Potter, S.M. (2005). Custom-made multiphoton microscope for long-term imaging off neuronal cultures to explore structural and functional plasticity. *Proc. SPIE* 5700: 102–108.
- Regehr, W.G., Pine, J., Cohan, C.S., Mischke, M.D., and Tank, D.W. (1989). Sealing cultured invertebrate neurons to embedded dish electrodes facilitates long-term stimulation and recording. *J. Neurosci. Meth.* 30: 91–106.
- Reger, B.D., Fleming, K.M., Sanguineti, V., Alford, S., and Mussa-Ivaldi, F.A. (2000). Connecting brains to robots: An artificial body for studying the computational properties of neural tissues. *Artif. Life* 6: 307–324.
- Rollenhagen, A. and Bischof, H.J. (1994). Spine morphology of neurons in the avian fore-brain is affected by rearing conditions. *Behav. Neural Biol.* 62: 83–89.
- Segev, R., Benveniste, M., Shapira, Y., and Ben-Jacob, E. (2003). Formation of electrically active clusterized neural networks. *Phys. Rev. Lett.* 90: Art. No. 168101.
- Shahaf, G. and Marom, S. (2001). Learning in networks of cortical neurons. *J. Neurosci.* 21: 8782–8788.
- Sharp, A.A., Abbott, L.F., and Marder, E. (1992). Artificial electrical synapses in oscillatory networks. *J. Neurophysiol.* 67: 1691–1694.
- Sharp, A.A., O’Neil, M.B., Abbott, L.F., and Marder, E. (1993). The dynamic clamp - Artificial conductances in biological neurons. *Trends Neurosci.* 16: 389–394.
- Shoham, S., Fellows, M.R., and Normann, R.A. (2003). Robust, automatic spike sorting using mixtures of multivariate t-distributions. *J. Neurosci. Meth.* 127: 111–122.
- Stenger, D.A., Pike, C.J., Hickman, J.J., and Cotman, C.W. (1993). Surface determinants of neuronal survival and growth on self-assembled monolayers in culture. *Brain Res.* 630: 136–147.
- Stoppini, I., Duport, S., and Correges, P. (1997). A new extracellular multirecording system for electrophysiological studies: Application to hippocampal organotypic cultures. *J. Neurosci. Meth.* 72: 23–33.
- Studer, L., Csete, M., Lee, S. H., Kabbani, N., Walikonis, J., Wold, B., and McKay, R. (2000). Enhanced proliferation, survival, and dopaminergic differentiation of CNS precursors in lowered oxygen. *J. Neurosci.* 20: 7377–7383.

- Suter, K.J. and Jaeger, D. (2004). Reliable control of spike rate and spike timing by rapid input transients in cerebellar stellate cells. *Neurosci.* 124: 305–317.
- Taylor, D.M., Tillery, S.I.H., and Schwartz, A.B. (2002). Direct cortical control of 3D neuroprosthetic devices. *Science* 296: 1829–1832.
- Trachtenberg, J., Chen, B.E., Knott, G.W., Feng, G.P., Sanes, J.R., Welker, E., and Svoboda, K. (2002). Long-term in vivo imaging of experience-dependent synaptic plasticity in adult cortex. *Nature* 420: 788–794.
- Tsai, P.S., Nishimura, N., Yoder, E.J., Dolnick, E.M., White, G.A., and Kleinfeld, D. (2001). Principles, design, and construction of a two photon laser scanning microscope for in vitro and in vivo brain imaging. In: Frostig, R., ed., *Methods for In Vivo Optical Imaging*. CRC Press, Boca Raton, FL.
- Ullian, E.M., Sapperstein, S.K., Christopherson, K.S., and Barres, B.A. (2001). Control of synapse number by glia. *Science* 291: 657–661.
- Wagenaar, D.A. and Potter, S.M. (2002). Real-time multi-channel stimulus artifact suppression by local curve fitting. *J. Neurosci. Meth.* 120: 113–120.
- Wagenaar, D.A. and Potter, S.M. (2004). A versatile all-channel stimulator for electrode arrays, with real-time control. *J. Neural Eng.* 1: 39–45.
- Wagenaar, D.A., Pine, J., and Potter, S.M. (2004). Effective parameters for stimulation of dissociated cultures using multi-electrode arrays. *J. Neurosci. Meth.* 138: 27–37.
- Webb, B. (2002). Robots in invertebrate neuroscience. *Nature* 417: 359–363.
- Weiler, I.J., Hawrylak, N., and Greenough, W.T. (1995). Morphogenesis in memory formation—Synaptic and cellular mechanisms. *Behav. Brain Res.* 66: 1–6.
- Williams, R.M., Zipfel, W.R., and Webb, W.W. (2001). Multiphoton microscopy in biological research. *Curr. Opin. Chem. Biol.* 5: 603–8.
- Wokosin, D.L., Amos, B.G., and White, J.G. (1998). Detection sensitivity enhancements for fluorescence imaging with multi-photon excitation microscopy. *IEEE EMBS* 20: 1707–1714.
- Zeck, G. and Fromherz, P. (2003). Repulsion and attraction by extracellular matrix protein in cell adhesion studied with nerve cells and lipid vesicles on silicon chips. *Langmuir* 19: 1580–1585.
- Ziegler, C. (2000). Information processing by natural neural networks: INPRO project IST-2000–26463.