

Real-time multi-channel stimulus artifact suppression by local curve fitting

Daniel A. Wagenaar^{a,*}, Steve M. Potter^b

^a Department of Physics, California Institute of Technology, Caltech 103-33, Pasadena, CA 91125, USA

^b Department of Biomedical Engineering, Georgia Institute of Technology and Emory University, 315 Ferst Dr. NW, Atlanta, GA 30332-0535, USA

Received 12 December 2001; received in revised form 22 May 2002; accepted 22 May 2002

Abstract

We describe an algorithm for suppression of stimulation artifacts in extracellular micro-electrode array (MEA) recordings. A model of the artifact based on locally fitted cubic polynomials is subtracted from the recording, yielding a flat baseline amenable to spike detection by voltage thresholding. The algorithm, SALPA, reduces the period after stimulation during which action potentials cannot be detected by an order of magnitude, to less than 2 ms. Our implementation is fast enough to process 60-channel data sampled at 25 kHz in real-time on an inexpensive desktop PC. It performs well on a wide range of artifact shapes without re-tuning any parameters, because it accounts for amplifier saturation explicitly and uses a statistic to verify successful artifact suppression immediately after the amplifiers become operational. We demonstrate the algorithm's effectiveness on recordings from dense monolayer cultures of cortical neurons obtained from rat embryos. SALPA opens up a previously inaccessible window for studying transient neural oscillations and precisely timed dynamics in short-latency responses to electric stimulation. © 2002 Elsevier Science B.V. All rights reserved.

Keywords: Artifact suppression; Micro-electrode array; Stimulation; Real-time; Local regression; Multi-electrode array; MEA; Spikes

1. Introduction

Micro-electrode arrays (MEAs) (Thomas et al., 1972; Gross, 1979; Pine, 1980) and related technologies such as tetrode probes (Gray et al., 1995), silicon probes (Bai and Wise, 2001) and multi-wire probes (Nicolelis et al., 1998) offer great promise to record action potentials extracellularly from a large number of cells simultaneously (Meister et al., 1994; Potter, 2001), in cell culture, in slice or in vivo (Bragin et al., 2000). In addition, electrical stimulation through such arrays has been reported in a wide variety of preparations, such as murine spinal cord (Gross et al., 1993), rat cortex (Jimbo et al., 1999), cat sciatic nerve (Branner and Normann, 2000) and rabbit retina (Grumet et al., 2000). Simultaneously stimulating and recording through a

single MEA is attractive for the study of input-output relationships (Novak and Wheeler, 1988; DeAngelis et al., 1998), but poses technical difficulties because the stimuli employed are often four or five orders of magnitude greater than extracellularly recorded action potentials ('spikes'). These may be as low as 10 μ V (shown below), while stimuli are typically on the order of a volt (Pancrazio et al., 1998; Jimbo et al., 1999), causing substantial stimulation artifacts that corrupt the data or saturate the recording electronics.

Several factors contribute to these artifacts (Grumet, 1999). The stimulus induces pickup on other electrode channels by a combination of capacitive crosstalk between leads and conduction through the tissue or recording medium, saturating the amplification system. The non-linear behavior of saturated amplifiers, together with the properties of the filters used for noise reduction, make this artifact last much longer than the stimulus that caused it, sometimes up to 100 ms (Maeda et al., 1995), even on channels not used for stimulation. In some cases this problem can be reduced by physically

* Corresponding author. Tel.: +1-626-395-6787; fax: +1-626-795-2872

E-mail address: wagenaar@caltech.edu (D.A. Wagenaar).

separating the recording site from the stimulation site (Grumet et al., 2000), or circumvented by using non-electronic means for either stimulation or recording, such as photo-uncaged glutamate (Wang and Augustine, 1995), optical recording (Obaid et al., 1996; Maher et al., 1999) or muscle twitch response (Branner and Normann, 2000). In all other cases, careful design of the electronics is required to minimize pickup of stimulation artifacts.

One would like to stop large artifacts from entering the recording system in the first place. To do so, Jimbo and Kawana (1992) recorded differentially between pairs of electrodes spaced at 10 μm , while stimulating between a similar, distant pair of electrodes. Sample-and-hold circuitry has also been used to prevent amplifier saturation (Novak and Wheeler, 1988; Jimbo et al., 1998; Grumet, 1999), but with mixed results. Jimbo et al. (1999) were able to record 5 ms after stimulation, even from the stimulated electrode, but the implementation details are not described. In contrast, Grumet (1999) reports little or no reduction of artifacts with a sample-and-hold approach. Presently commercially available electrophysiology equipment for electrode arrays does not employ this strategy.

When artifacts cannot be entirely prevented in hardware, various forms of digital filtering can be used to reduce them. For example, if artifacts are the same across trials, template subtraction can be employed: Jimbo and Kawana (1992) estimated artifacts in stimulus responses by scaling the artifact recorded under sub-threshold conditions. Unfortunately, due to hysteresis at the electrode interface or in the electronics, artifact shapes often do vary between trials. Alternatively, Okajima et al. (1995) manually subtracted a linear baseline from recordings of muscle action potentials. This is too labor-intensive for multi-channel recordings, and certainly cannot be applied in real time as the data come in. As a last resort, blanking (digitally setting the signal to zero) can be used to eradicate any artifacts (O'Keefe et al., 2001). Any action potentials occurring within the duration of the artifact are lost, so if one is interested in the early part of stimulus response, this is not an option. Our solution, which works in real time, is SALPA, an algorithm for Subtraction of Artifacts by Local Polynomial Approximation. We show that its performance is superior to some possible alternatives: two simple high pass filters and one linear phase filter (see e.g. Jackson, 1996).

2. Methods

2.1. Cell culture and MEA preparation

Dense cultures of dissociated rat cortical cells were grown on MEA dishes and maintained for several

months. Culture methods have been detailed elsewhere (Potter and DeMarse, 2001). Very briefly, cortex from E18 Wistar rats was dissected under sterile conditions and dissociated using papain. MEA dishes from Multi-Channel Systems (Reutlingen, Germany, <http://www.multichannelsystems.com>) with sixty 10 μm diameter electrodes arranged in a rectangular array with 200 μm interelectrode spacing were coated with polyethylene imine (PEI) and laminin. A 15 μl droplet of neurobasal medium containing 50 000 cells was applied to the electrode area, and the dishes were sealed with FEP Teflon® lids (Potter and DeMarse, 2001). After 30 min, 1 ml of medium was added, and the dishes were transferred to an incubator (35 °C, 5% CO₂, 9% O₂ and 65% humidity). After 1 day, and thereafter every 4 days, the medium was replaced entirely by the medium adapted from Jimbo et al. (1999). Electrical activity of these cultures was recorded through the MEA, amplified and sampled with 12 bits resolution at 25 kHz (Multi-Channel Systems). One electrode was used for stimulation, while all the others were used for recording. The data presented below were obtained by stimulating 5-months-old cultures through one electrode with single biphasic voltage pulses of ± 600 mV, lasting 400 μs per phase, positive phase first.

2.2. Artifact suppression

SALPA works by locally fitting (Hastie and Loader, 1993) a function to the recorded trace that has enough degrees of freedom to accurately model the artifact, but not enough to represent individual action potentials. By subtracting this fitted function from the recording, the artifact-free signal remains, and action potentials can be detected by setting a voltage threshold. We observed considerable variability in artifact shapes between electrode channels, and even on individual channels between trials of the same stimulus. Therefore, we chose not to make assumptions about regularities in artifact shapes, and instead fit independent functions to each individual artifact.

For every timepoint n_c in the recording from a given electrode, a third degree polynomial is fitted to a segment of $2N+1$ samples centered around n_c . The segment half-length N is experimentally tuned for optimal performance, as discussed below (see *Results*). The fitted value at the central point n_c is subtracted from the raw recording at that point to yield a cleaned signal. The rest of the fitted curve is discarded; to estimate the cleaned signal at n_c+1 , a new polynomial is fitted to the segment of $2N+1$ samples centered around n_c+1 .

The fitting process is different for the first N points of the raw data, starting when the stimulus-induced saturation of the electronics ends ('depegging'; saturation of the electronics is determined by the digital signal having its minimum or maximum possible value).




Fig. 1. Illustration of the fitting method used by SALPA. The first fit after depegging that has acceptable deviation (see text) is used to model $N+1$ samples (bottom-most curve). Thereafter, each fitted polynomial is used to model one sample only (other dashed curves). The thin solid curve is the raw recording. For visual clarity, only one in ten fits is shown, and they have been vertically displaced. The thick solid curve is the resulting model of the artifact. Circles mark the centers of each fit. The dotted vertical lines indicate which parts of the fitted polynomials are used for modeling the artifact.

A polynomial is fitted to the data centered on the $(N+1)$ -th point after depegging, and the artifact up to the center of that window is estimated using that single third-degree polynomial, as illustrated in Fig. 1.

The raw electrode signal is represented as a sequence of sampled voltages, V_n , where n is the ordinal number of the sample (we sample at 25 kHz, so the unit of n is 40 μ s in real time). We decompose this raw signal into an estimated artifact, A_n , and a cleaned signal, v_n , by assuming that in the vicinity of some central point n_c , we can approximate the artifact by a cubic polynomial expressed in terms of the distance $(n-n_c)$ from the central point¹:

$$A_n^{(n_c)} = \alpha_0^{(n_c)} + \alpha_1^{(n_c)}(n-n_c) + \alpha_2^{(n_c)}(n-n_c)^2 + \alpha_3^{(n_c)}(n-n_c)^3.$$

The fit parameters α_k are found by minimizing the function

$$\chi_{(n_c)}^2 = \sum_{n=n_c-N}^{n_c+N} (V_n - A_n^{(n_c)})^2$$

¹ The notation $[\cdot]^{(n_c)}$ will be used throughout to represent the quantity $[\cdot]$ evaluated for the fit centered around n_c .

with respect to those parameters. We then estimate the corrected voltage at the central point n_c as:

$$v_{n_c} \equiv V_{n_c} - A_{n=n_c}^{(n_c)} = V_{n_c} - \alpha_0^{(n_c)}.$$

Next, we obtain a new fit centered around n_c+1 to estimate $v_{n_c+1} = V_{n_c+1} - \alpha_0^{(n_c+1)}$, and so on.

Fitting a new polynomial for every single datapoint might seem to be computationally very expensive, but it is not, because the fit parameters can be calculated recursively, as follows. Let us introduce the shorthands

$$T_k = \sum_{n=n_c-N}^{n_c+N} (n-n_c)^k, \quad \text{for } k=0 \dots 6,$$

and

$$W_k^{(n_c)} = \sum_{n=n_c-N}^{n_c+N} (n-n_c)^k V_n \quad \text{for } k=0 \dots 3;$$

and the (4×4) -matrix S with entries $S_{kl} = T_{k+l}$ (for $k, l=0 \dots 3$). The parameter values that minimize χ^2 can then be written as

$$\alpha_k^{(n_c)} = \sum_{l=0}^3 S_{kl}^{-1} W_l^{(n_c)}, \quad \text{for } k=0 \dots 3,$$

which can be computed cheaply once the $W_l^{(n_c)}$ are known, since the entries of S^{-1} are constants depending only on N . The complexity is further reduced because T_k is identically zero for odd values of k .

A recursion relation for $W_k^{(n_c)}$ is obtained by straightforward algebraic manipulation of the expressions for $W_k^{(n_c+1)}$ and $W_k^{(n_c)}$:

$$W_k^{(n_c+1)} = \sum_{l=0}^k \frac{(-1)^{k-l} k!}{l!(k-l)!} W_l^{(n_c)} + N^k V_{n_c+N+1} - (-N-1)^k V_{n_c-N}.$$

It is the existence of this closed-form expression that makes the method viable for real-time operation.

As noted above, at the beginning of the recording, just after a channel depegs, we are forced to use a fit based on a window centered N samples ahead in time. Such a non-central window is likely to give a less accurate fit to the artifact, so it is important to assess the quality of the fit before trusting it. This assessment can be based on the *deviation*:

$$\mathcal{D}^{(n_c)} = \sum_{n=n_c-N}^{(n_c-N)+(\delta-1)} (V_n - A_n^{(n_c)}),$$

where the width of the estimator, δ , can be chosen to be some fixed fraction of N , e.g. $\delta = N/10$. For good fits, $\mathcal{D}^{(n_c)}$ is normally distributed with zero mean, and variance $\sigma_{\mathcal{D}}^2 = \beta^2 \delta \sigma_V^2$, where σ_V^2 is the variance of the recording, and β is a correction factor larger than unity if the noise in the recording is not white. For our




Fig. 2. An example of an artifact in electrode recordings with the output of various filters (left) and the effect on spike waveforms (right). Notice the difference in scales between left and right: the artifact in the raw data is an order of magnitude larger than the spikes. Thin curves are raw data; thick curves are filter output. From top to bottom: SALPA, subtraction of low-pass Butterworth (BW-L), high-pass Butterworth (BW-H), and 39 pole linear phase filter (LPC). Notice SALPA's blanking of the output during saturation of the electronics. The Butterworth filters induce significant phase distortion while leaving much more lost time than SALPA. Even the linear phase filter leaves some echo of the artifact. The spike waveforms shown are from the same recording as the artifacts, but at longer latency to allow direct comparison with non-corrupted raw data.

equipment, $\beta^2 \approx 5$. We advance n_c until $\mathcal{D}^{(n_c)}$ attains an acceptably small (see below) absolute value relative to $\sigma_{\mathcal{D}}$, before declaring the artifact successfully suppressed.

The following results were obtained with $N = 75$ (corresponding to 3 ms at 25 kHz sampling rate), $\delta = 5$ and by rejecting fits after depegging until \mathcal{D}^2 no longer exceeds $3^2 \times 5 \times \sigma_V^2$. Spikes were identified by thresholding at five times RMS noise, and validated based on a test of their waveform shape: spikes were rejected if there were any peaks of either polarity larger than 90% of the main peak within ± 1 ms (P.P. Mitra, personal communication).

Our C++ implementation of SALPA is capable of processing 60 channels of MEA data at 25 kHz in real-time on an AMD Athlon 1.33 GHz processor, using just 75% of CPU time. This allows online spike detection entirely in software on an inexpensive PC system. The C++ source code is available upon request.

2.3. Analysis

To assess the quality of the algorithm, we compared its output on typical artifact-corrupted data with the output of three alternative filters: a three pole Butterworth high-pass filter with 400 Hz cutoff (BW-H), subtraction of the output of a three pole Butterworth low-pass filter with 600 Hz cutoff (BW-L), and a 39 pole linear phase high-pass filter with 500 Hz cutoff, designed using cosine expansion (LPC) (Jackson, 1996). The order of this filter was chosen such that we could compute it in real-time using a simple C++ program. The Butterworth filters were chosen because they are a computationally inexpensive simple alternative.

Two performance measures were used: *Lost time*, the latency after depegging of the electronics at which the artifact is successfully suppressed, and *PNR loss*, the reduction of the ratio of action potential peak amplitude




Fig. 3. *Lost time* (left, solid; results are shown for four dishes separately) and *PNR reduction* (right, dashed) can be traded off by varying SALPA's filter length, measured here in milliseconds. Note that the left-hand axis increases downwards, so 'up' means 'better' for both axes. Artifact sizes and duration vary by almost an order of magnitude between dishes, depending on electrode impedances (nominally 300 k Ω at 1 kHz). This is reflected in SALPA lost time, shown here for four different dishes. PNR drops dramatically when the filter half-width N approaches the duration of action potentials. The optimal choice of N must depend on the application, and on the PNR in the raw recordings. The results in the rest of this article were obtained with a filter half-length of 3 ms. PNR change is measured relative to single pole high-pass filtering at 150 Hz (see text). The sample period, τ_{sample} , was 40 μs .




Fig. 4. Comparison of various filter methods. PNR change is measured relative to a single pole high-pass filter at 150 Hz (see text). Notice that the reference 150 Hz filter also reduced spike amplitudes by a small fraction, so relative PNR *gain* resulted in some cases. Restricted SALPA is SALPA without the third degree term. Lost time does not include the duration of amplifier and ADC saturation (1.04 ± 0.02 ms). Charted values are mean and standard deviation of the data collected from 55 electrodes.

to noise (PNR) induced by the filter. Lost time was determined by computing 5 ms wide box-car averages of the signal, and rejecting data until the box-car average no longer exceeded the RMS noise. PNR loss was measured relative to raw data filtered through a single pole high-pass filter at 150 Hz, which, before developing SALPA, we used routinely to clean the data of DC drift and any low frequency local field potentials for the purpose of spike detection in recordings of spontaneous activity. In general, artifact suppression filters will reduce the ratio of spike amplitude to RMS noise, because there is substantial spectral overlap between artifacts and spike waveforms.

3. Results

Rat cortical cultures were stimulated with 600 mV biphasic pulses. Large dish-wide artifacts were observed in the resulting recordings. Fig. 2 shows how SALPA and the other filters act on these stimulation artifacts and on action potential waveforms. Aside from reducing spike amplitudes, filters may distort spike waveforms in more subtle ways, exemplified by the positive 'ghost' phases induced by BW-L and BW-H. These may hamper subsequent spike sorting, and may even lead to spurious detection of non-existent upgoing action potentials. Fortunately, SALPA—being a linear phase filter (except in the initial N samples after depegging)—is free of such phase distortions.

All filters used in this comparison have parameters that can be tuned to trade off lost time against PNR reduction. For SALPA, this is the segment half-width N ; for the other filters the cut-off frequency plays this role. Fig. 3 presents the trade-off for SALPA. The optimal choice of N depends on the kind of experiment one is doing. In Fig. 4, the performance of SALPA at $N = 75$ (3 ms) is compared with the other filters (at the frequencies specified above).

Features of SALPA that give it an edge over the alternatives, are that it explicitly recognizes saturation of recording electronics, outputting zeros whenever the digital values of the recording are at the extreme ends of their range, and that it incorporates a statistic to test goodness of fit for the earliest timepoints, as detailed in *Methods*. As a result, SALPA performs well on a wide range of artifact sizes and shapes (Fig. 5).

The ultimate test for an artifact suppressor is whether it allows reliable detection of spikes at short latencies. Fig. 6 shows detectable spikes on two electrodes in different MEA dishes containing mature cultures, using SALPA, using only a 150 Hz high-pass filter, and using each of the other filters mentioned above. SALPA reveals a structure of oscillations and very precisely timed spikes in the early response that would otherwise go undetected.




Fig. 5. For widely different artifact waveforms (shown in insets), SALPA yields usable output as early as 2 ms post stimulus (less than 1 ms after depegging). Faint traces are raw data with two different vertical offsets added to show details. Bold is SALPA output. Notice the spikes riding on the slopes of the artifacts which cannot be deleted by thresholding the raw data.

4. Discussion

We have presented an algorithm for stimulus artifact suppression that can be applied to 60-channel MEA recordings in real time on inexpensive PC hardware. SALPA does not cause phase distortion of spike waveforms unlike simple high-pass filters, and it is less computationally intensive than straightforward implementation of a generic linear phase filter of equivalent length. Perhaps more importantly, the algorithm covers the first few milliseconds of the artifact naturally, because it takes amplifier saturation into account explicitly. Other filters tend to suffer from ringing as a result of the sharp transient at the time of amplifier depegging.

In the bulk of the data, local regression of the form used in SALPA is just a special case of linear phase filtering. SALPA effectively functions as a high pass filter with -3 dB cut-off frequency $f_0 = 0.6 f_s/N$, i.e. $f_0 = 200$ Hz for $N = 75$ and sampling frequency $f_s = 25$ kHz. This is sufficient, since a spectrogram of the stimulation artifacts would reveal that high frequency power is mostly concentrated in the first few milliseconds—beyond that, the tail of the artifact is reasonably well spectrally separated from spike waveforms. Simple filters would have to find a difficult compromise between preventing ringing from the initial sharp transient of artifacts, and preserving signal shape. SALPA surmounts this problem by not using any samples from the sharp transient in its estimate of the shape of the tail, through the use of asymmetric fitting windows shortly after depegging (as illustrated in Fig. 1).

One could attempt to improve the performance of the algorithm by increasing the order of the polynomials used. The computational expense of the algorithm would increase, more dramatically than one might expect, because the higher powers of $(n - n_c)$ in the equations would make representation of these numbers as 32-bit integers impossible. The main effect of increasing polynomial order is an increase of cut-off frequency, which can be achieved more economically by decreasing N . In any event, it remains desirable to use odd order, because this gives the fit at the beginning of the trace one more degree of freedom than in the bulk, further improving the response to sharp transients without compromising signal shapes.

With current commercially available hardware, SALPA is less well suited for recordings from the stimulated electrode, because saturation on that channel lasts beyond the duration of the early phase of the response. Improvements in hardware, for example using sample-and-hold based artifact reduction, might bring the stimulated channel within SALPA's domain.

SALPA opens up a new window on very short latency multi-neuronal responses to electrical stimulation. The early post-stimulus neural dynamics comprise oscillations and action potentials timed with a precision not observed before. We are currently using SALPA to investigate the nature of these response components and their role in stimulus-induced plasticity. The precisely timed responses can be used to drive 'behaviors' in a neurally-controlled animat (DeMarse et al., 2001).




Fig. 6. Evoked activity on two electrodes (from different dishes). Each set of raster plots shows action potentials detected using the reference 150 Hz high-pass filter ('raw') and each of the filters tested. Each raster contains results from 500 consecutive trials, with each dot representing one action potential. Stimulation was on electrodes 280 and 450 μm away from the recording site. Grey bars represent the time and duration (0.8 ms) of the stimulus. Spikes were detected by thresholding at five times RMS noise and validated based on their waveform shape (see text). Notice that the last time was estimated independently for each individual trial. This explains the limited reliability with which, e.g., the spike at 5 ± 0.5 ms in the bottom panel is detected by BW-H. Electrodes were chosen to illustrate a variety of short latency response types.

Acknowledgements

The authors wish to thank Jerry Pine for useful comments on the manuscript. This work was supported by grant RO1-NS38628 from the NINDS and by the Burroughs-Wellcome Fund/Caltech *Computational Molecular Biology* program.

References

- Bai Q, Wise KD. Single-unit neural recording with active microelectrode arrays. *IEEE Trans Biomed Eng* 2001;48(8):911–20.
- Bragin A, Hetke J, Wilson CL, Anderson DJ, Engel J, Buzsaki G. Multiple site silicon-based probes for chronic recordings in freely moving rats: implantation, recording and histological verification. *J Neurosci Methods* 2000;98(1):77–82.
- Branner A, Normann RA. A multielectrode array for intrafascicular recording and stimulation in sciatic nerve of cats. *Brain Res Bull* 2000;51(4):293–306.
- DeAngelis GC, Cumming BG, Newsome WT. Cortical area MT and the perception of stereoscopic depth. *Nature* 1998;394(6694):677–80.
- DeMarse TB, Wagenaar DA, Blau AW, Potter SM. The neurally controlled animat: biological brains acting with simulated bodies. *Autonom Robots* 2001;11:305–10.
- Gray CM, Maldonado PE, Wilson M, McNaughton B. Tetrodes markedly improve the reliability and yield of multiple single-unit

- isolation from multi-unit recordings in cat striate cortex. *J Neurosci Methods* 1995;63(1–2):43–54.
- Gross GW. Simultaneous single unit recording in vitro with a photoetched laser deinsulated gold multi-micro-electrode surface. *IEEE Trans Biomed Eng* 1997;26(5):273–9.
- Gross GW, Rhoades BK, Reust DL, Schwalm FU. Stimulation of monolayer networks in culture through thin-film indium-tin oxide recording electrodes. *J Neurosci Methods* 1993;50(2):131–43.
- Grumet AE. Electrical stimulation parameters for an epi-retinal prosthesis, PhD thesis, Massachusetts Institute of Technology, 1999.
- Grumet AE, Wyatt JL, Jr, Rizzo JF, III. Multi-electrode stimulation and recording in the isolated retina. *J Neurosci Methods* 2000;101(1):31–42.
- Hastie T, Loader C. Local regression: automatic kernel carpentry. *Stat Sci* 1993;8(2):120–43.
- Jackson LB. Digital filters and signal processing, with Matlab® exercises, 3rd edition. Boston: Kluwer Academic Publishers, 1996.
- Jimbo Y, Kawana A. Electrical stimulation and recording from cultured neurons using a planar electrode array. *Bioelectrochem Bioenerg* 1992;29(2):193–204.
- Jimbo Y, Robinson HPC, Kawana A. Strengthening of synchronized activity by tetanic stimulation in cortical cultures: application of planar electrode arrays. *IEEE Trans Biomed Eng* 1998;45(11):1297–304.
- Jimbo Y, Tateno T, Robinson HPC. Simultaneous induction of pathway-specific potentiation and depression in networks of cortical neurons. *Biophys J* 1999;76(2):670–8.
- Maeda E, Robinson HPC, Kawana A. The mechanisms of generation and propagation of synchronized bursting in developing networks of cortical neurons. *J Neurosci* 1995;15(10):6834–45.
- Maher MP, Pine J, Wright J, Tai YC. The neurochip: a new multi-electrode device for stimulating and recording from cultured neurons. *J Neurosci Methods* 1999;87(1):45–56.
- Meister M, Pine J, Baylor DA. Multi-neuronal signals from the retina—acquisition and analysis. *J Neurosci Methods* 1994;51(1):95–106.
- Nicolelis MAL, Ghazanfar AA, Stambaugh CR, Oliveira LMO, Laubach M, Chapin JK, Nelson RJ, Kaas JH. Simultaneous encoding of tactile information by three primate cortical areas. *Nat Neurosci* 1998;1(7):621–30.
- Novak JL, Wheeler BC. Multisite hippocampal slice recording and stimulation using a 32 element microelectrode array. *J Neurosci Methods* 1988;23(2):149–59.
- Obaid AL, Farries MA, Kisley MA, Sakai T, Salzberg BM. Simultaneous optical recording of electrical activity with single cell resolution from several ganglia in a mammalian plexus. *J Gen Physiol* 1996;108(2):73.
- Okajima Y, Tsubahara A, Kondo K, Chino N, Noda Y, Tomita Y. A new method of estimating the distribution of muscle fiber conduction velocities. *Electromyogr Mot Control-Electroencephalogr Clin Neurophysiol* 1995;97(6):310–7.
- O’Keefe DT, Lyons GM, Donnelly AE, Byrne CA. Stimulus artifact removal using a software-based two-stage peak detection algorithm. *J Neurosci Methods* 2001;109(2):137–45.
- Pancrazio JJ, Bey PP, Jr, Loloee A, Manne SR, Chao HC, Howard LL, Gosney WM, Borkholder DA, Kovacs GTA, Manos P, Cuttino DS, Stenger DA. Description and demonstration of a CMOS amplifier-based-system with measurement and stimulation capability for bioelectrical signal transduction. *Biosens Bioelectron* 1998;13(9):971–9.
- Pine J. Recording action potentials from cultured neurons with extracellular microcircuit electrodes. *J Neurosci Methods* 1980;2:19–31.
- Potter SM. Distributed processing in cultured neuronal networks. In: Nicolelis MAL, editor. *Progress in Brain Research*, vol. 130. Elsevier Science, 2001:49–62.
- Potter SM, DeMarse TB. A new approach to neural cell culture for long-term studies. *J Neurosci Methods* 2001;110(1–2):17–24.
- Thomas CA, Jr, Springer PA, Loeb GE, Berwald-Netter Y, Okun LM. A miniature microelectrode array to monitor the bioelectric activity of cultured cells. *Exp Cell Res* 1972;74:61–6.
- Wang SSH, Augustine GJ. Confocal imaging and local photolysis of caged compounds—dual probes of synaptic function. *Neuron* 1995;15(4):755–60.