

# Techniques for Long-Term Multisite Neuronal Ensemble Recordings in Behaving Animals

Jerald D. Kralik,\* Dragan F. Dimitrov,\*,† David J. Krupa,\* Donald B. Katz,\* Dana Cohen,\* and Miguel A. L. Nicolelis<sup>\*,1</sup>

\*Department of Neurobiology and †Division of Neurosurgery, Duke University, Durham, North Carolina, 27710

Advances in our understanding of neural systems will go hand in hand with improvements in the experimental techniques used to study these systems. This article describes a series of methodological developments aimed at enhancing the power of the methods needed to record simultaneously from populations of neurons over broad regions of the brain in awake, behaving animals. First, our laboratory has made many advances in electrode design, including movable bundle and array electrodes and smaller electrode assemblies. Second, to perform longer and more complex multielectrode implantation surgeries in primates, we have modified our surgical procedures by employing comprehensive physiological monitoring akin to human neuroanesthesia. We have also developed surgical implantation techniques aimed at minimizing brain tissue damage and facilitating penetration of the cortical surface. Third, we have integrated new technologies into our neural ensemble, stimulus and behavioral recording experiments to provide more detailed measurements of experimental variables. Finally, new data analytical techniques are being used in the laboratory to analyze increasingly large quantities of data. © 2001 Elsevier Science

The relatively new paradigm of multichannel electrophysiology has clearly been having a profound impact in the neurosciences. The reason for this interest is probably not due so much to a change in research interests, since it has been appreciated for some time that recording from large numbers of neurons simultaneously in the primate cortex will help to uncover the mechanisms underlying such higher brain functions as learning, planning, and memory (see Lilly (1), for instance). Rather, the excitement seems to stem more from the methodological side, in that the technologies are now becoming available to record extracellularly from potentially hundreds of neurons distributed broadly over the cortex and other brain regions (2-10). Advances in the methodology, then, must continue to improve, to keep up with the exciting research and medical opportunities that these improvements offer. Although long-term multichannel single-unit recordings in multiple cortical and subcortical areas have been successfully obtained in many animal species, including birds, rats, and primates, it is clear that chronic electrode implantation in even more brain areas while recording from larger populations of neurons is a goal that could be attained in the near future. Further, these techniques may someday be applied to human subjects, to alleviate debilitating impairments (11). Taking these next steps will require the development of new techniques to address key issues in all areas of neuronal ensemble electrophysiology: multielectrode design, surgical implantation procedures, multichannel singleunit, stimulus, and behavioral recordings, and the analyses of the data sets that such work generates. In this article, we provide a detailed description of critical issues in each of these main areas and the advances we have made to address them.

# 1. MULTIELECTRODE DESIGN

Many laboratories, including our own, have had significant success in chronically implanting standard electrode bundles and arrays (NBLabs, Dallas, TX) into various brain structures in several animal species. However, to be able to implant areas where the spontaneous firing rate is low, especially under anesthesia,

<sup>&</sup>lt;sup>1</sup> To whom all correspondence should be addressed at Department of Neurobiology, Box 3209, Room 327E, Bryan Research Building, Duke University Medical Center, 101 Research Drive, Durham, NC 27710. Fax: (919) 684-5435. E-mail: nicoleli@neuro.duke.edu.

and to maximize overall neuronal yield, we have focused on developing bundle and array electrodes that are not permanently secured in place at the time of surgery, but rather can be moved to access more neurons while the animal is awake and behaving. In addition, we have been working on a design to reduce the overall size of the electrode assembly, to make multichannel recordings in smaller species, such as in mice. This latter development could provide the opportunity, for example, to conduct experiments that combine multielectrode physiology with molecular genetics.

## 1.1. "Mini-microdrive" Drivable Bundles

To facilitate reliable implantation of deeper structures, we have developed a variation on the classic drivable bundle design originally described by Kubie (12) and used in a number of laboratories (see 4, 13). Adjustments performed in our laboratory make it possible to place multiple bundles on a single rat's head, and to divide the wires of a single assembly into two or more independently movable cannulas. The construction and drivability of this assembly greatly increase its yield: cannulation of the bundle translates into lower variability of implant angle, and drivability removes guesswork in the dorsal–ventral plane.

In this design, 16 microwires (each 25  $\mu$ m in diameter, with an additional  $\sim 10$ - $\mu$ m Formvar jacket) are threaded through a 27-gauge thin-walled stainlesssteel tube (the "electrode cannula"), or are split among two or more narrower cannulas. The electrode cannula is soldered to a 21-gauge tube (the "guide cannula"), which slides into a microdrive. The microdrive, essentially a 3 imes 10 imes 2-mm tablet of dental acrylic with a hole running lengthwise through it, is made in a Teflon mold (Fig. 1): an 0-80 screw projects into the mold from one end, and is met halfway across by a 21-gauge stainless-steel tube; after dental acrylic has been poured into the mold and allowed to harden, both screw and tube can be cracked out with gentle torque. An 0-80 set screw is then screwed into the threaded half of the hole, and the guide/electrode assembly is slid in from the bottom, such that it meets the set screw (Fig. 1). When the set screw is turned further, the guide cannula is slowly ejected from the other end, until the set screw rests at the end of the tapped portion of the hole.

The microwires are glued both to the cannula and to the top end of the microdrive, leaving a short length of taut wire between the two glued points. The set screw is then turned back by  $\sim 2$  mm, and the cannula assembly pushed into meet it; this causes a small amount of slack to be introduced into the length of wire between the two glued points (see Fig. 1). Finally, the microdrive is glued to the side of a connector plug, and the microwires are individually glued onto the pins with a dot of silver print paint. Connection points are covered with a protective coating. The assembly presents a remarkably slim surface when looked down on from above, taking up only the space required by the plug itself (Microtech, Inc., Boothwyn, PA; in our case,  $2 \times 10$  mm) plus the  $2 \times 3$ -mm profile of the mini-microdrive. As many as four or five such bundles may fit on a single rat's head. Figure 2 shows signals culled from simultaneously implanted amygdalar and cortical bundles.

The bundle is trimmed such that only 2-3 mm of wire protrudes from the bottom of the electrode cannula, and the tips are gently spread just prior to implantation. For insertion, the cannula and length of slack wire are covered with petroleum jelly and lowered into an area dorsal to the intended recording site; as the cannula is rigid, deep structures can be targeted with relative

FIG. 1. Drivable microwire bundles with "mini-microdrives." The schematic details the mold used to produce the mini-microdrives. "The schematic details the mold used to produce the mini-microdrives, and shows the drive itself ejected from the mold following removal of the mold screw and tube (see text for details). An 0-80 set screw is then turned into the tapped hole, and a drive-cannula/electrode-cannula assembly is slid up to meet it. The inset shows the finished product, with the microdrive (C) inbetween the cannula (D) and microdrive; this slack is taken up as the cannula and electrodes are driven deeper into the brain in the weeks following surgery. Bar (in inset)-3 mm.





**FIG. 2.** Recordings from multiple simultaneously implanted movable bundles. A screen capture showing the spikes (725  $\mu$ s of signals that have crossed an amplitude threshold) from two wires in the same rat. The top panel is from a wire implanted into insular cortex, the bottom from the amygdala. In both, the waveforms of particular colors have been identified as similar in form, using the time–amplitude boxes shown; the green spikes are clearly single units.

ease. The entire assembly is carefully cemented to the skull, but the petroleum jelly ensures that the drivable components are free to move. Later advancement can be very precise (one entire turn of the set screw advances the tips by  $\sim 300 \ \mu$ m), and can be done while the animal is awake and plugged into the recording system (although it may take some time for recordings to stabilize following tip movement). Thus, proper placement can proceed while the animal is awake and performing—a great benefit when movement-related units are sought, or when regions that respond poorly under anesthesia are probed.

## 1.2. Miniature Eight-Channel Multielectrode Implant Device with Individually Movable Electrodes

In numerous experimental situations, it is desirable to record simultaneously the activity of individual neurons from many discrete locations within a particular structure. Several laboratories have reported good success by implanting fixed arrays of microwire electrodes into target structures (see, for instance, Nicolelis et al. (9)). However, in certain circumstances, it would be particularly beneficial to be able to move individual electrodes within a recording array after it has been implanted. Individually movable electrodes would increase the overall yield of recorded neurons simply because each electrode can sample different neurons as it is moved through the target structure. Also, individually movable electrodes would allow neurons to be sampled simultaneously at different depths within the structure. For example, one could record simultaneously from different layers in cerebral cortex during the same recording session. Finally, individually movable electrodes would allow electrodes to be precisely positioned within regions of a structure that are not uniformly distant from the surface. For example, cerebellar Purkinje cells are arranged in a monolayer in the cerebellar cortex. The highly foliated nature of this structure results in substantial differences in the depth of neighboring Purkinje cells relative to the pial surface. Thus, to record from neighboring populations of Purkinje cells, individual electrodes must be positioned at varying and substantially different depths.

Here, we describe the design of a multielectrode array that was developed to address these issues regarding the need for electrode mobility, while maintaining the spatial arrangement of the electrodes in the array. In addition to the ability to move individual electrodes within the array, the entire electrode assembly was designed to be as compact as possible to facilitate implantation of multiple arrays within an adult rat. Finally, the entire device is relatively simple in design and, thereby, easy and inexpensive to construct.

The device consists of eight individual microelectrodes. These are fashioned from tungsten rods (100  $\mu$ m in diameter) that are electrolytically etched in 2 M  $KNO_2$  to have a sharp tip (see, for instance, Freeman (14) and Hubel (15)). The taper in diameter of the electrodes is such that the diameter 1 mm from the tip is approximately 25  $\mu$ m and the diameter 0.5 mm from the tip is approximately 15  $\mu$ m. Thus, these electrodes easily penetrate the dura mater and other overlying membranes and cause minimal disruption of tissue near the recording tip. The electrodes are insulated by being dipped several times in epoxylite and subsequently baked at 110°C for several hours (see Snodderly (16)). The electrode tips are exposed by electrically "blasting" insulation from the tip with a controlled, high-voltage spark (17). Exposed tips of approximately 10–15  $\mu$ m are typically used.

The electrodes are arranged in a linear row with centers 250  $\mu$ m apart (Fig. 3A). A small, molded acrylic block with eight parallel holes (120  $\mu$ m diameter) running through the length of the block accurately positions each electrode in the array (see Fig. 3B). Individual electrodes can slide within the holes in the block. Two stainless-steel tubes on either side of the block accurately position the block relative to the individual microdrives. Each microdrive is constructed of a 16gauge internally threaded (0-80) stainless-steel hypo tube. Within each of these threaded tubes is a 0-80 set screw that pushes a short piece of Teflon tubing (24 gauge). Thus, for each full turn of the set screw, the Teflon tube is extended 318  $\mu$ m. Each tungsten microelectrode is passed through a hole in the acrylic block and the uninsulated end of each electrode is pushed about 1-2 mm into the lumen of one of the Teflon tubes in each microdrive. The end of a short length of Tefloninsulated silver wire (75  $\mu$ m diameter) is also inserted into the Teflon tube (see Fig. 3A). The silver wire is electrically connected to the tungsten electrode by allowing a small amount of conductive silver paint to flow into the Teflon tube. When the paint hardens, it mechanically holds the tungsten electrode and silver wire together within the Teflon tube and electrically connects them with a very low resistance connection. The other end of the short silver wire is soldered to a terminal on a connector used to connect the electrode to an external amplifier. Thus, when the set screw is rotated, the Teflon tube is extended out of the threaded hypo tube. As the Teflon tube moves out, the attached tungsten microelectrode is also extended, resulting in accurate movement away from the end of small acrylic block. As seen in Fig. 3B, the odd-numbered electrodes have been extended about 1 mm from the block while the even-numbered electrodes have been extended about 900  $\mu$ m further. Together, these components form

a very compact, lightweight eight-channel microdrive system. Multielectrode arrays have been used for neuronal recordings in somatosensory cortex in rats. The device is stereotaxically lowered in surgery so that the tips of the electrodes are positioned 100  $\mu$ m above the dura. The device is then cemented to the skull using dental acrylic. A viscous gel such as silicone grease is placed into the space between the microdrives and the acrylic block to prevent the dental acrylic from interfering with electrode movement. The rat is allowed to recover for at least 5 to 7 days. At any time after the surgery (hours to weeks), electrodes can be individually driven through the dura into underlying target neural structures. The sharp tips of the tungsten electrodes



FIG. 3. Drivable microwire array. (A) Example of a multielectrode device. Four stainless-steel microdrives can be seen attached to the electrode connector (lower portion of figure). The other four drives are attached to the opposite side of the connector. The electrodes in the four drives on the front of the connector have been extended approximately 1 mm. Note the Teflon tubes extending above the stainless-steel drive tubes and the electrode tips extending 1 mm above the acrylic block. The microdrives attached to the opposite side of the connector have been extended about 1.9 mm. Again, note the Teflon tubes extending above the drives and the tips of the electrodes above the acrylic block. The silver wires can be seen running from the Teflon tubes (where they connect with the tungsten electrodes by conductive silver paint) to the individual contacts on the connector. Bar = 5 mm. (B) Closeup of the eight electrodes passing through the acrylic block. Odd- and even-numbered electrodes are extended to different depths. Electrodes are spaced 250  $\mu$ m apart (center to center). Bar = 2 mm. (C) Example of a recorded signal from primary somatosensory cortex in an awake rat. Two separate single units are present. Signal was bandpass filtered between 0.5 and 5 kHz.

facilitate penetration into the brain with minimal dimpling of dura and subsequent compression of underlying tissue. After the electrodes have been driven to a desired depth, single-unit neuronal activity can be recorded. An example of a recorded signal is shown in Fig. 3C. There are two discriminable single units. This activity was recorded in the barrel region of the primary somatosensory cortex in an awake rat approximately 3 weeks after surgery.

#### 1.3. Reduced-Size Electrode Arrays

Reducing the size of the electrode assemblies is essential when implanting smaller animals like birds and mice. For such species, commonly used electrode arrays are typically larger than the animal's head, which even in the best case prevents the animal from maintaining its normal behavior while in its home cage, as well as during recording. Reducing the overall size of electrode arrays is valuable even when implanting larger animals like rats and monkeys. A smaller electrode assembly occupies less space on the animal's skull, providing an avenue for implanting and thus simultaneously recording from many more neurons and brain structures.

In our laboratory, a significant reduction in electrode array size was achieved by reducing the size of the connector on the electrode assembly. This new connector (Omnetics Connector Corp., Minneapolis, MN) is about a third of the size of regular connectors (such as those from Microtech, Inc., Boothwyn, PA). Additionally, we were able to decrease the distance between the connector and the tip of the wires by about one-half. The resulting decrease in the overall size of the arrays achieved with these two changes allows the head stage to be much smaller, thereby reducing considerably the total weight of the permanently mounted assembly on the animal's head.

The basic procedure of making such smaller electrode arrays requires several adjustments. A silver-impregnated conductive epoxy is applied on the connector's pins. Eight microwires, 1 inch in length, made of 45- $\mu$ m-diameter tungsten and coated with 5  $\mu$ m Isonel (California Fine Wire company, Grover Beach, CA), are placed gently on the connector's pins and are baked to harden the conductive epoxy. The wires are then passed through a guide tube arranged in a desired configuration. Once in place, dental acrylic cement is used to fix the position of the wires. The guide tube is made out of several stainless-steel tubes (Small Parts Inc., Miami Lakes, FL) organized according to the desired electrode configuration. The diameter of the stainless-steel tubes determines the interwire distance.

The configuration is chosen based on the number and geometry of the structures to be implanted. In the mouse, the brain structures are rather small and located too close to one another to allow the implant of two independent arrays. In such cases, the guide tubes are arranged to create two separated  $2 \times 2$  arrays (Fig. 4A) instead of a single array of  $1 \times 8$  (Fig. 4B). The guide tube arrangement can be tailored to match the distance between target structures.

Finally, it should be noted that the drivable bundles described above were successfully adjusted to work with the smaller connector as well.

# 2. SURGICAL IMPLANTATION PROCEDURES

Once the multielectrode bundles and arrays are assembled, they must be surgically implanted into the targeted brain structures. Major components of this surgery are the induction and maintenance of an adequate level of anesthesia and the surgical implantation of the electrodes. In this section, we describe ways to address two critical issues here: (1) the maintenance of homeostasis in New World monkeys over long, invasive brain surgeries, and (2) the reliable penetration of the pial layer by the electrodes.

# 2.1. Neuroanesthesia

We have performed implantation of up to 128 microwires in multiple cortical areas in New World primates (8, 9). Such surgeries may someday be performed in humans for therapeutic purposes (11). To be able to perform longer and more extensive surgical procedures, we have recently implemented a level of care that is similar to that used for humans undergoing neurosurgery (see Miller et al. (18)). In nonhuman primates, carefully monitoring the animal's physiology and performing appropriate interventions intraoperatively and postoperatively allow us to perform extremely long procedures, up to 14 h, with excellent rates of survival, even in older animals. Furthermore, several physiological manipulations can be made that have a significant impact on the actual implantation of the electrodes, such as managing the size of the brain.

The central problem faced when placing an animal under anesthesia is the animal's inability to monitor its own homeostasis. In essence, one must assume the responsibility of monitoring and responding to changes in physiological parameters that the animal can no longer control. Using the human operating room as the standard, we have assembled a variety of monitoring devices and techniques that allow us to assess continuously the physiological state of the animal throughout the surgical procedure (see Fig. 5A).

The animal is intubated prior to surgery and mechanical ventilation is provided throughout using a neonatal, pressure cycled ventilator. During surgery, heart rate, oxygen saturation, blood pressure, respiratory rate, end tidal carbon dioxide, temperature, blood glucose, and urine output are monitored. At least one person is dedicated to recording physiological data throughout the procedure. Following trends throughout the surgical procedure is often more useful than knowing the absolute values.

Heart rate is one of the most useful parameters, which we measure in two ways. The electrical activity is measured using neonatal limb leads connected to a standard electrocardiogram (ECG) machine (Horizon 2000 Patient Monitor, Mennen Medical, Inc., Clarence, NY). The pulsatile flow of blood produced by the heart is measured using a pulse oximeter (CO<sub>2</sub>SMO+ Respiratory Profile Monitor, Novametrix Medical Systems, Inc., Wallingford, CT) (Fig. 5B). In combination, these two methods of measuring heart rate provide a great deal of useful information regarding the functioning of the cardiovascular system. The redundancy in the values provides a useful way of cross-checking the accuracy of one particular value. If the ECG rate suddenly drops to zero, for example, the important distinction between device malfunction versus cardiac standstill

**FIG. 4.** Two examples of the reduced size electrode arrays made in our laboratory. (A) Eight wires (45- $\mu$ m-diameter tungsten coated with 5 $\mu$ m Isonel) arranged in two separated arrays of 2 × 2. The distance between the two arrays can be adjusted according to need. (B) The eight wires are arranged as a single array of 1 × 8.



and the need for immediate intervention can be ascertained by a quick glance at the pulse rate registered by the pulse oximeter.

Heart rate also provides an overall assessment of the depth of anesthesia (see Miller *et al.* (18)). We continuously compare preoperative baseline values with intraoperative values and aim to keep the heart rate within 30% above the baseline value. In addition to measuring the pulsatile forward flow of blood, the pulse oximeter





B

FIG. 5. (A) Neuroanesthesia and monitoring setup. In clockwise order starting at the top left: (1) Vetroson isoflurane vaporizer veterinary anesthesia machine (Summit Hills Laboratories, Navesink, NJ), (2) Dinamap vital signs monitor 1846SX (noninvasive neonatal blood pressure monitor; Critikon Inc., Tampa, FL), (3) CO<sub>2</sub>SMO+ respiratory profile monitor (pulse oximeter/ETCO<sub>2</sub> monitor; Novametrix Medical Systems, Inc., Wallingford, CT), (4) Infant Star Neonatal Ventilator (i.e., the bottom two pieces of equipment at the far right; theInfrasonics, Inc., San Diego, CA), (5) Horizon 2000 patient monitor (Mennen Medical, Inc., Clarence, NY). Data collected from these individual units can be downloaded to a PC for correlation with neuronal firing. Most devices are equipped with internal battery packs that can be used during neuronal recording to minimize noise. In addition, most of these devices are available as refurbished items from medical supply corporations. (B) Pulse Oximeter/End-Tidal CO2 Monitor closeup (CO<sub>2</sub>SMO+ respiratory profile monitor, Novametrix Medical Systems, Inc.). The large number top left is the CO<sub>2</sub> in mm Hg of the exhaled breath. The ETCO<sub>2</sub> is represented graphically as it varies with inspiration and expiration seen on the top waveform. As air is exhaled, ETCO2 rises, causing a peak in the waveform. The number of peaks generated per minute, the respiratory rate, is depicted as the smaller top number. The large bottom number represents the percentage saturation of arterial blood with oxygen. The pulsatile flow of the blood is depicted in the bottom waveform, providing information on the quality of the pulse as well as the pulse rate per minute, which is depicted numerically as the smaller bottom number. Again, this information can be downloaded and correlated with neuronal firing. The unit can operate temporarily on an internal backup battery.

gives an estimation of percentage saturation of arterial blood with oxygen (oxygen saturation). Oxygen saturation is particularly useful since it gives an overall estimation of the cardiopulmonary function. In essence, if the pulse oximeter registers values above 90%, the animal's lungs and heart are generally functioning well. Any major perturbation to the cardiopulmonary system will cause a rapid decline in the oxygen saturation. Further, because oxygen saturation is measured in the periphery (tail or distal extremity) it is very sensitive to any perturbations that occur before they affect the brain and vital organs. Oxygen saturation is therefore a very sensitive way of ascertaining overall cardiovascular stability.

It is important to note that with smaller primates such as owl (Aotus trivirgatus) and squirrel (Saimiri *sciureus*) monkeys, it is technically challenging to place the sensor and emitter in such a way that they provide accurate, stable values. There is clearly a steep learning curve to working with such devices. We have found the proximal tail and lower leg to be the best places to attach probes on squirrel and owl monkeys. We shave and remove hair chemically from all extremities, abdomen, pelvis, and most of the proximal tail to ascertain the best possible placement of the probe. We have tried various methods of attaching the probe to the animal and noted significant differences between various methods. It is much easier to place probes on larger animals such as rhesus monkeys (Macaca mulatta) due to their larger fingers, toes, and ears, which make ideal sites for probe attachment. Ambient light severely impedes functioning of pulse oximeters. A piece of tin foil placed over the probe overcomes this issue.

Blood pressure measurement is a good adjunct to pulse oximetry and heart rate. Blood pressure can be measured noninvasively using an appropriately sized neonatal blood pressure cuff and a noninvasive blood pressure monitor (Dinamap Vital Signs Monitor 1846SX, Critikon, Inc., Tampa, FL). Measuring blood pressure using direct cannulation of the arterial system in owl and squirrel monkeys is difficult due to the small caliber of their vessels and the risk of limb necrosis from arterial damage or vasospasm. In larger animals such as rhesus monkeys, invasive blood pressure monitoring may be more feasible.

Baseline blood pressures of squirrel monkeys after ketamine anesthesia vary significantly among animals. Average baseline values are in the range of 90 mm Hg systolic and 50 mm Hg diastolic. The effect of isoflurane on blood pressure also varies significantly among squirrel monkeys. Theoretically, blood pressure measurements are useful in titrating isoflurane since it is known that the effect of isoflurane is inversely proportional to blood pressure (19). Due to the high pulse rate and low blood pressure in squirrel and owl monkeys, measuring noninvasive blood pressure is difficult. It is more reliable in rhesus monkeys, and therefore more useful.

End-tidal carbon dioxide (ETCO<sub>2</sub>) can be used to measure respiratory rate and estimate blood carbon dioxide levels (Fig. 5B). ETCO<sub>2</sub> measurements are useful immediately after intubation since they unequivocally distinguish placement of the endotracheal tube into the bronchial tree versus the esophagus. The return of measurable levels of carbon dioxide after intubation guarantees bronchial intubation (18). ETCO<sub>2</sub> is measured throughout surgery. It registers as a waveform and absolute number that correlates with the amount of carbon dioxide in the exhaled air as it passes a small window. The waveform provides not only the respiratory rate, but also the quality of the respirations (i.e., shallow and short versus deep and long), which helps assess the depth of anesthesia.

Absolute ETCO<sub>2</sub> numbers are useful since they closely estimate the carbon dioxide level in the blood. The carbon dioxide level in the blood is inversely related to respiratory rate. In other words, hyperventilation causes the animal to "blow off" carbon dioxide. The blood level of carbon dioxide directly impacts the degree of vasoconstriction in cerebral vessels. High levels of carbon dioxide cause vasodilation, while low levels cause vasoconstriction. The degree of vascular tone of the cerebral vessels is proportional to cerebral blood flow and, thus, intracranial pressure. By this mechanism the ETCO<sub>2</sub> has significant consequences in regard to brain physiology, and can conveniently be manipulated by altering the rate of respiration (18). We take full advantage of these relationships by controlling respiratory rate throughout surgery. This allows us not only to maintain homeostasis but also to tailor brain size to suit our particular needs. For example, maintaining a lower ETCO<sub>2</sub> by mild hyperventilation makes drilling craniotomies and opening the dura easier by keeping the brain small, while maintaining a higher than normal ETCO<sub>2</sub> causes the brain to swell, theoretically making brain penetration and electrode placement easier.

Mechanical ventilation provides necessary control over parameters of the animal's breathing. To accommodate owl and squirrel monkeys, which have very small tidal volumes, we use a pressure-cycled ventilator intended for human neonates. Rather than delivering a specific volume of gas, pressure-cycled ventilators deliver gas until a preset pressure is reached, which in practice works better for very small lung volumes.

Depending on our particular needs at various times throughout the operation, and depending on the animal's own drive to breathe, intermittent mandatory ventilation (IMV), in which the ventilator provides breaths intermittently to the animal, or continuous positive airway pressure (CPAP), in which the ventilator assists the animal's own breathing, mode is selected. Using these two modes of ventilation, we are able to fully control ventilation and oxygenation.

We have noted that owl and squirrel monkeys have a tendency to hyperventilate after intubation using isoflurane anesthesia alone. This hyperventilation is evidenced by a shallow rapid breathing pattern and results in a rapid decline in the ETCO<sub>2</sub> to unacceptably low levels. Not only does this cause the brain to shrink, making electrode placement difficult, it can cause ischemia of the brain as well. Overcoming this requires a combination of ventilator settings that match the animal's own respiratory pattern and the use of respiratory depressing agents. We have successfully used fentanyl, an opioid, to achieve respiratory depression in squirrel monkeys (18). Much like morphine in adults, fentanyl causes respiratory depression in neonates. Its use is limited by its propensity to induce hypotension. Naloxone, a competitive inhibitor, can be used to quickly reverse its effects. By depressing the respiratory drive with fentanyl and varying the respiratory rate on the ventilator, brain size can be quickly increased or decreased to accommodate surgical needs.

In addition to fentanyl, several other pharmacologic agents have become a standard part of our surgical routine. Preoperatively, we administer ketamine for light anesthesia, combined with glycopyrolate to lessen oral secretions for intubation (note that food and water are removed from the animal's cage in the evening prior to the day of surgery). We administer an intramuscular dose of a broad-spectrum antibiotic with good coverage of skin flora. Dexamethasone can be given intravenously if there is concern regarding cerebral edema, especially if a large number of electrodes are implanted. Intraoperatively, we keep epinephrine, atropine, glucose, naloxone, and midazolam on hand. Intraoperatively, we have noted that owl and squirrel monkeys are particularly susceptible to hypoglycemia and thus we check blood glucose using a commercially available kit. We add glucose to the intravenous fluid to maintain adequate blood glucose levels as necessary. Postoperatively, a dose of midazolam remains available throughout recovery in the event of a seizure. We avoid intraoperative use of benzodiazepines such as midazolam due to the theoretical decrease in spontaneous neuronal firing that they may promote.

Our aggressive approach toward neuroanesthesia allows us to perform extremely long surgical procedures safely. This has opened the door to implanting more microwire arrays into more cortical and subcortical areas.

## 2.2. Neurosurgery

Two characteristics of the brain continue to provide particular challenges to reliable multielectrode penetration: the brain "floating" in the cerebrospinal fluid, and the relative toughness of the pia layer surrounding the cortex (see, e.g., Carpenter (20)). Here we describe techniques to minimize the effects of these two factors on successful multielectrode penetration of the brain.

A unique design feature of the brain is that it is bathed in cerebrospinal fluid, essentially tethered to the dura and rigid cranium only at certain key points. While this provides a built in shock-absorbing mechanism and allows the brain to enlarge and shrink within physiological limits, it makes the cortex a challenging target for chronic electrode implantation. Once an electrode array is in place, various factors including natural movement of the animal and physiological brain swelling and shrinking may cause the brain to move in relation to the electrodes if they are rigidly attached to the cranium. This effect is believed to be even more pronounced in larger animals with a better-developed subarachnoid space, such as rhesus monkeys as compared with rats.

The pia mater, or "pia," consists of an inner membraneous layer composed of fine elastic and reticular fibers that intimately invests the surface of the brain, draping over cortical sulci and under cortical vessels and an outer epipial layer composed of fiber bundles intermeshed with the arachnoid trabeculae (20-23). The fiber bundles of the pia contain the structural molecule collagen, unlike the brain parenchyma itself, which does not appear to contain collagen, except in the walls of the larger blood vessels coursing through it (24; see Alberts et al. (25)). The internal resistance and structure of the brain parenchyma appear to be provided solely by myelinated axons, glia, neurons and the blood vessels (see Alberts et al. (25) and Carpenter (20)). Although the pia is not very strong in relation to the brain, it possesses a formidable resistance. We have observed that it is often easier to deform the brain than it is to penetrate the pia. The term "dimpling" is commonly used to describe the situation where the electrodes are pushing the brain cortex in without penetrating it (see Nicolelis (2)). Obviously, the more electrodes there are in an array, the more pronounced this effect becomes. In addition to potentially injuring the brain tissue, dimpling is obviously a source of error in the determination of depth measurements. Ideally, if dimpling could be eliminated, the electrodes would move in relation to the pial surface, allowing effective and accurate electrode placement.

### 2.2.1. Cyanoacrylate Adhesives

To minimize the detrimental influences of brain movement at the electrode implantation site, we have employed the use of cyanoacrylate adhesives, applied directly to the pial surface in rats and primates. An air-powered dental drill is first used to create a craniotomy over the area of interest. The craniotomy is made large enough to allow a perimeter of at least 1 mm around the electrode array. The dura is carefully removed over the entire craniotomy. Next, petroleumbased ointment such as bacitracin ointment is applied to the exact site of electrode implantation, either on a small piece of Gelfoam (Upjohn, Kalamazoo, MI) or as a droplet encompassing the wires of the array. This step is intended to protect the electrode entry site from the cyanoacrylate adhesive since the cyanoacrylate will not stick to areas covered with petroleum ointment. It is vital to limit the application of the petroleum ointment to the site on the cortex that the electrodes will penetrate, so that it does not counteract the intended use of the cyanoacrylate. The cyanoacrylate adhesive is then applied judiciously to the zone of the pia surrounding the electrode array and onto the cranial surface. Care is taken prior to this step to remove as much cerebrospinal fluid as possible using rounded fine tip suction cannulas that can be touched to the surface of the pia without injury. Once a thin layer of the cyanoacrylate adhesive is applied, the polymerization can be visually observed to occur over a period of few minutes. After the petroleum ointment at the target site is removed, a small window of exposed pia with surrounding cyanoacrylate adhesive adhered to the pia can be seen (Figs. 6A, 6B). Next, the electrodes are advanced with the micropositioner at approximately 0.05- to 0.1-mm increments (Fig. 6C). By following the crescendo of neuronal firing, and using depth measurements, the electrodes are advanced until maximal neuronal firing is obtained or a depth of 1.8 to 2.0 mm is reached.

In the frontal cortex of Long–Evans rats, this sequence achieved reliable penetration of the cortex with electrode arrays. In a series of five rats, we observed that the onset of neural activity began reliably at approximately 0.6-mm depth from the pia surface (range: 0.2-0.9 mm). In all cases, beyond the point where neuronal activity was first detected, changes in depth resulted in changes in neuronal firing patterns, suggesting that the electrodes were moving in relation to the brain.

It appears that the cyanoacrylate glue creates a bond between the rigid skull and the pia. This is suggested by the fact that the previously pulsating brain, a phenomenon that is enhanced at the frontal pole in the



**FIG. 6.** Cyanoacrylate glue technique. (A) After application of the cyanoacrylate glue to the pia surrounding the target site, the polymerized glue can be seen surrounding an open area of exposed cortex with a small blood vessel running through it. (B) Under higher power, the vessel and cortex appear uninjured. Cerebrospinal fluid can be seen bulging, per normal, from beneath the glued area. (C) A 16-microwire array (NBLabs, Dallas, TX) is seen penetrating the pia and cortex through the window in the glue. The cyanoacrylate adhesive glues the pia of the brain to the skull of the craniotomy, creating a hardened crust that provides a rigid structure that prevents dimpling.

rat, no longer moves in relation to the skull after application of the glue. Further, attempting to remove the glue once it is applied to the pia results in tearing of the cortical surface with resultant bleeding of cortical vessels. By adhering to the pia, the glue can be thought of as enhancing the natural surface structural elements of the brain by making the pia a more rigid structure and fixing it in relation to the skull. As a result, the electrodes penetrate the pia instead of dimpling it, and thus move into the brain reliably.

Although pathological confirmation is required, we feel that there is minimal injury to the cortex or cortical blood vessels with this technique. After application, the cortical vessels and yellow brain tissue can be seen unchanged through the translucent cyanoacrylate adhesive. Further, long-term recordings in rats are comparable to techniques not using cyanoacrylate adhesive.

#### 2.2.2. Collagenase

To implant microelectrode arrays and bundles in the rats, we have found that the glue procedure alone appears sufficient for penetrating the pia. In other words, structurally stabilizing the pia allows the electrodes to penetrate the pia of the rat with very little dimpling. For array or bundle implantation in owl and squirrel monkeys, however, using glue alone is only intermittently successful in penetrating the thicker pial layer of the primate. In our experience thus far, approximately 25% of the time glue alone is sufficient for penetrating the monkey pia with little dimpling. To increase the success rate, we complement the glue technique by removing the main structural element at the surface of the brain posing the most resistance to multielectrode penetration, namely, the pia (remembering that the dura has already been removed). In the past we have physically removed the pia using a sharp hook before lowering the microwires. We have recently attempted to remove the pia enzymatically, to limit injury to deeper layers. Collagen appears to be the major structural material of the pia layer, providing this resistance (20-24); see Alberts et al. (25)); and thus, the application of collagenase, the enzyme that breaks down collagen, would remove the main source of resistance, allowing for easier implantation. We have found that the application of collagenase does indeed break down the pial layer, allowing for easier implantation of electrode arrays in the primate cortex.

Our procedure for applying collagenase to the pia layer, in both monkeys and rats, is the following. After the craniotomy is made, and the dura and arachnoid layers are surgically removed (see above), a collagenase solution is applied. First, we make a solution consisting of 20 mg/ml crude collagenase (Sigma, St. Louis, MO), 0.36 mM calcium chloride, in 50 mM Hepes buffer, pH 7.4, at 37°C (26-28). This buffered solution is then mixed with an equal amount of KY jelly. We have successfully applied the collagenase solution to the cortical surface in two different ways. First, collagenase is applied using a rectangular block of Gelfoam (approximately 2 mm wide  $\times$  3 mm long  $\times$  4 mm tall) that was soaked in the collagenase mixture. The collagenasesoaked Gelfoam block is then placed directly on the surface of the pia, covering a little more than the surface area equal to the area of the electrode array (approximately  $2 \times 3$  mm for a  $2 \times 8$  array from NBLabs that has an overall surface area covered by the electrodes of approximately  $1 \times 2$  mm). The collagenase mixture is kept on the surface of the pia for approximately 15 min. During this time, the glue is applied to the remaining exposed pia surface and to the surrounding skull surface, both inside the craniotomy and just onto the surface of the skull (approximately 2 mm beyond the diameter of the craniotomy). Afterward, the Gelfoam is lifted off, and the area is washed with a constant stream of saline that is removed using suction for approximately 2 min.

The second procedure that we use for applying collagenase to the pia consists of directly soaking the electrode tips in the collagenase mixture. The collagenasesoaked electrode tips are lowered onto the surface of the pia, held there for approximately 15 min, and then lowered into the cortex. The glue is applied to the remaining exposed pia surface and to the surrounding skull surface.

Our overall approach for placing microelectrode wires into cortical targets involves maximizing those factors that promote entry and minimizing those factors that hinder it. We feel that by enhancing the structural integrity of the pia in areas surrounding the exact site of implantation we are able to minimize dimpling. By enzymatically breaking down the pia at exactly the site of electrode entry, we are able to minimize the resistance that the electrodes meet at the surface of the brain. In combination, these two techniques have significantly improved our ability to place electrodes into the cortex without physically disturbing the cortex with sharp instruments.

Since each species and each individual animal is different with respect to the resistance of the brain to electrode penetration, we have developed an algorithm of successively more aggressive steps in penetrating the cortex with multielectrode arrays. We have found that for rats, using the cyanoacrylate adhesive technique is sufficient for consistant reliable penetration. However, with squirrel monkeys, where the brain coverings are clearly thicker, adjuncts are necessary. In our squirrel monkey surgeries we typically employ both the cyanoacrylate adhesive technique and collagenase from the outset. If we fail to achieve recordings, signaling failure to penetrate the cortex, we remove the electrode, and apply successively higher concentrations of collagenase. In our experience this technique has completely eliminated the need for sharp dissection of the pia, which we feel is more injurious to the cortex.

# 3. ACQUISITION AND RECORDING OF NEURONAL ENSEMBLES, SENSORY STIMULI, AND BEHAVIOR

After the animals have recovered from the multielectrode implantation surgery, they are ready to be included in experiments in which simultaneous recordings are made of populations of neurons and of several stimulus and behavioral parameters.

# 3.1. Data Acquisition and Recording of Populations of Single Units

Neural signals are recorded using a 128-channel many-neuron acquisition processor, the MNAP (Plexon Inc., Dallas, TX), which can theoretically record from 512 single units simultaneously. This system provides direct computer control over amplification, filtering, signal selection, spike-waveform discrimination, and storage for each channel. Single-unit isolation is carried out both on-line using the software program SortClient and again off-line using the software program Offline-Sorter (both from Plexon Inc.). Individual units are discriminated using multiple time-voltage windows, principal components analysis of spike waveforms, and interspike interval histograms. Figure 7A shows a raster plot of simultaneous recordings from populations of units in several cortical areas during a segment of a session from a motor control experiment in the monkey (8) (also see Fig. 2 for examples of waveforms).

# 3.2. Acquisition and Recording of Stimuli and Behavior

In this section we describe devices that we have developed or integrated into our current recording systems to make increasingly more detailed studies of (1) tactile discrimination in rats and (2) cortical motor control of hand and arm movements in rats and monkeys.

#### 3.2.1. Tactile Discrimination Tasks in Rats

To study the mechanisms and processes involved in the acquisition and recognition of different tactile stimuli by the trigeminal somatosensory system in the awake, behaving rat, we developed a novel behavioral paradigm where rats are trained to discriminate between different tactile stimuli using only their facial whiskers. With this task, we are able to present stimuli to the rats' whisker pad in a highly controlled, yet flexible manner. Using this task, we are able to record neuronal ensemble activity from multiple locations within the trigeminal system while rats actively sample these behaviorally relevant stimuli with their whiskers. These recordings can also be combined with reversible inactivation of the barrel region of SI cortex with the  $\gamma$ -aminobutyric acid (GABA) agonist muscimol. By quantitatively comparing ensemble responses in thalamus or brainstem before and after SI inactivation, we will be able to assess the role of corticofugal projections in active somatosensory discrimination.

The behavioral whisker discrimination task is designed so that rats use their facial whiskers to discriminate between two different possible stimuli. Rats (Long-Evans, adult, male, approximately 300 g), are trained to enter a darkened rectangular chamber and poke their nose into a small hole located at the end of the chamber (Fig. 8). Directly in front of this hole is a variable-width aperture. The width is set by computer-controlled stepping motors. The aperture is positioned so that the mystacial whiskers are in contact with it when the rats' nose is poked into the hole. Once the rats have poked their nose into the center hole and sampled the aperture with their whiskers,



**FIG. 7.** (A) Example raster plots of neural ensemble spike trains simultaneously recorded from posterior parietal (PP), primary motor (MI), dorsal premotor (PMd), and ipsilateral primary motor, and dorsal premotor cortices in an owl monkey, obtained from a segment of a session in a motor control experiment. Each horizontal line depicts the activity of a single unit (i.e., presumed single neuron), and each dot represents a single firing of the unit. (B) Hand position in *X*, *Y*, and *Z* dimensions, as measured by the Shape Tape, described in the text.

they are then trained to signal whether the aperture is wide or narrow by poking their nose into one of two reward holes located in a different section of the training box. Optical infrared sensors are used to determine accurately when the whiskers are in contact with the aperture. All mechanical and electronic devices are controlled with a PC running MedPC behavioral control software (Med Associates, St. Albans, VT).

# 3.2.2. Electromyogram Recording

To study the subtle and complex relationships of neuronal ensembles in different brain structures with individual muscles, it will be necessary to record from multiple muscles, neurons and brain areas simultaneously



FIG. 8. Schematic diagram of the behavioral training apparatus for sensory discrimination in the rat. The width of the variable-width aperture is set by rotating the computer-controlled stepper motors. Moving each side of the aperture inward, toward the center, results in a narrower aperture; moving them outward results in a wider aperture. At the start of each training session, a rat is placed in the outer reward chamber with the sliding door closed. When the door is opened, the rat enters the center discrimination chamber and pokes its nose into the center nose poke and samples the variable-width aperture with its facial whiskers. After poking its nose into the center nose poke (detected by interrupting an infrared photobeam), the rat then backs out into the outer reward chamber and pokes its nose into either the left or right reward nose poke to receive a water reward: left nose poke if the aperture is narrow, right nose poke if the aperture is wide. Immediately after the rat pokes its nose into either the left or right nose poke, the sliding door between the outer reward chamber and the center discrimination chamber is closed. A new trial begins when the sliding door is again opened.

(29). To this end, our laboratory has been able to record electromyographic (EMG) activity of several muscles simultaneously with the recording of neural ensembles (7; see Loeb and Gans (30)). For instance, we have recorded from EMG electrodes chronically implanted in the rat triceps brachii (to study elbow extension), and trapezius (to study rotation of the forelimb at the shoulder), while concurrently recording from an array implanted in the rat motor cortex (7). Bipolar patch electrodes (MicroProbe, Potomac, MD) were implanted subcutaneously; and the EMG signals were amplified (SA Instrumentation, Encinitas, CA), filtered at 20-500 Hz, sampled at 1000 Hz, and input directly into the same data file as the neuronal signal recordings. An EMG recording system is currently being integrated into our monkey neurophysiological setup.

#### 3.2.3. Motor Control Experiments in Monkeys

To study such topics as how motor cortical areas (dorsal premotor, primary motor, and posterior parietal cortices, for instance) encode hand and arm position, or to study the changes in response properties of neurons during motor learning, we use the following devices. Accurate 3D hand and arm position and orientation measurements are obtained using a new type of technology, the Shape Tape (Measurand, Inc., Fredricton, New Brunswick, Canada), shown in Figs. 9A and 9C. This recording device consists of a flexible tape attached to the arm that contains sensors based on optical fibers that signal the bending and twisting of the tape along its length. The resulting analog signals are sampled at 200 Hz, saved in the same data file as the neural recordings, and later converted to 3D positions (see Figs. 9B and 9D). The device does not interfere with the neuronal recordings and it can provide higher temporal resolution (200 Hz) than comparable optical or magnetic systems. Example recordings of hand position in 3D space are shown in Fig. 7B.

Other behavioral experimental apparatus, such as devices that provide visual cues and juice reinforcement to the monkeys or record the position of manipulanda moved by the monkey, are controlled in real time by a personal computer running the Tempo (Reflective Computing, St. Louis, MO) behavioral data collection system. In addition, all experimental sessions are videotaped, with the video frames being time-stamped with a 1-ms resolution by a video timer system that receives a timing signal from the MNAP recording system.

# 3.2.4. New World Monkey Handling and Chair Training Procedures

Before experiments can be conducted on monkeys, they must learn to sit comfortably in a restraining chair. We have developed handling and chair training procedures for working with New World monkeys in awake, behaving experiments, which ensure the successful training of all monkeys, regardless of disposition, and which allow several different people to work with the monkeys, without introducing bias. To remove monkeys from their home cages, we lure them into transport boxes (Primate Products, Miami, FL) with small pieces of fruit. Once in the test room, we have designed a restraining chair that fastens onto the transport box. The restraining chair is also designed to be initially in the form of a box, which then transforms in successive stages into a restraining chair. The monkeys learn to step from the transport box directly into the restraining "box." We then place the restraining box in the sound attenuating test chamber. The back of the restraining box is designed to be moved forward, toward the front, progressively restraining the monkey. We first move this back wall forward halfway, leaving the monkey a little space to move. The monkeys are then trained to lift their heads up through the neck brace opening at the top front of the box. When the monkeys lift their heads through this opening, the neck brace is closed, holding the monkey in place. The adjustable back wall is then moved the final distance, to hold the monkey firmly in place in the restraining chair. If there is any difficulty training an individual animal to lift its head through the neck brace opening, we have also designed small poles that hook the front loop on the monkey's harness, helping to coax the monkey to lift its head through the neck brace opening. Finally, one or both of the front doors on the chair are removed, to allow free movement of either arm. The head stages are then plugged into the monkey's implanted electrode arrays, and the experiment is begun. By designing a chair that transforms from a box into a restraining chair in successive stages, we can train all of the New World monkeys gradually to acclimate them to restraining, which ensures that every monkey will be readily chair-trained, and at the same time it standardizes our training procedures so that several people can train the same animals using exactly the same training sequence.

# ANALYSIS OF NEURONAL ENSEMBLES, SENSORY STIMULI, AND BEHAVIOR

The recordings from neuronal ensembles in multiple brain sites, together with the recordings of numerous measures of stimuli and behavior, generate unprecedented amounts of neurophysiological data, providing a significant challenge in the attempt to uncover the underlying structure inherent in these enormous data sets. Thus, we continue to investigate ways in which mathematical models and statistical tools can be used to conceptualize and analyze these data sets. In this section we provide a brief description of the main set of analysis tools that we currently use to understand the functional response properties of the neurons, as they relate to other neurons, stimuli, and behavior (Fig. 10).

#### 4.1. Single Neuron Analysis

Analysis of the neural ensemble recordings begins with an investigation of the response properties of each of the recorded neuronal units (commercially available software packages designed for such analyses include Neural Explorer a.k.a. Nex, from Plexon, Inc.). In particular, raster plots and peristimulus time histograms (PSTHs, or more generally, PETHs, where the "E" stands for stimulus or behavioral "events"). (Figs. 11A and 11C) graph neural firing patterns of individual cells in relation to individual stimulus or behavioral events of interest in the experiment, allowing one to identify the receptive fields of each of the neurons (e.g., Krupa et al. (31) and Nicolelis et al. (32); see Nicolelis et al. (9)). Figure 11, for example, shows the activity of a single unit around the time of the mechanical stimulation of a digit on the hand of an owl monkey. Clearly the plots reveal a strong potential relationship between the neural activity of this unit and the stimulation of the finger, and further statistical tests can then be conducted to verify the relation. Two statistical tests are commonly used to assess the significance of the firing pattern: (1) the Kolmogorov-Smirnov test, which tests whether a sample distribution diverges from a theoretical distribution by comparing the cumulative frequency functions of the two distributions, can be used to determine whether the distribution of counts in the histogram is nonrandom, or (2) Student's *t* test, which can be used to compare the average firing frequency between control and experimental periods (9, 49, 55).

To obtain precise estimates of the latency of the neuronal response to a stimulus or behavioral event, cumulative frequency histograms (CFHs) can be used (Fig. 11B) (9, 32, 44). Latencies are identified by isolating the exact time in which the firing frequency distribution diverges from a random distribution with P < 0.01 (using the Kolmogorov–Smirnov test).

Then, to begin to characterize the relationships between neurons, it is often useful to make pairwise comparisons between pairs of units, using such standard analysis techniques as cross-correlation analysis (see Gerstein (45)).

#### 4.2. Visualizing the Response of the Neural Ensemble

Following the analysis of single neurons and pairs of neurons, it is useful to visualize the spatiotemporal patterns of neural ensemble firing, and here we describe two graphing schemes used extensively in the laboratory (31, 32). Population peristimulus time histograms (PPSTHs, or PPETHs, where the "E" again represents stimulus or behavioral "events") allow one to visualize how an entire population of neurons respond as an ensemble to an experimental event of interest, such as to the mechanical stimulation of a primate's hand and digit (Fig. 12A). In these 3D graphs, the individual neurons line the x axis, ordered in some meaningful way such as rostrally to caudally, the peristimulus time is depicted on the y axis, and again some measure of the firing frequency of the neurons is depicted along the z-axis, such as the instantaneous firing rate (in spikes per second) or the spike counts per time bin. Usually the data averaged over 100 to 300 trials are plotted, being smoothed with a spline algorithm (see de Boor (46)). Again, such plots provide a useful glimpse at how the population ensemble responds around the time of the experimental event (32).

If there is a known spatial arrangement of the neurons, such as when they are recorded with evenly spaced electrode arrays, spatiotemporal population maps (SPMs) can be used to visualize how the spatial pattern of a neural ensemble varies over time in response to a behavioral or stimulus event (31, 32). As shown in Fig. 12B, an SPM is a series of graphs depicting the spatial arrangement of the neurons (along the x and y axes), along with a measure of the instantaneous firing of the neurons either color-coded or depicted along a z axis in a 3D plot. This representation helps the experimenter to detect potential spatial and temporal patterns in the neuronal ensemble activity in relation to a stimulus or behavioral event. These observed patterns can then be analyzed numerically in greater detail, as described in the sections below.

## 4.3. Linear Time Series Analysis

After investigating the response of the neural ensembles graphically, we begin a series of quantitative analyses of the data. For continuous stimulus or behavioral variables, such as hand position during a monkey's reach, it is often useful to begin with a linear analysis. The model we use is an extension of the basic linear regression equation, y = ax + b, in which the neuronal inputs, *x*, and the behavioral outputs, *y*, are considered as a time series (33). The relationship between neural signals and behavior should not be limited to being

synchronous in time, since, for instance, the firing of a neuron at time *t* can potentially influence behavior at time  $t + \Delta t$ ,  $t + 2\Delta t$ , and so on. As seen in Fig. 13, **X**(*t*) is a matrix of neural signals, with each column corresponding to a single neuron and each row corresponding to an epoch of time in the trial (for instance, a 100-ms time segment). In the case of three-dimensional behavioral data, such as are obtained using the Shape Tape, in which the *x*, *y*, and *z* dimensions of hand position are recorded, **Y**(*t*) is a matrix of three columns, one for each dimension. As can be seen in Fig. 13, for each neuron, a separate weight *a* is obtained for each time lag in a series of time lags *u* between neuronal firing and hand position, representing the different associations

that may exist between the firing of a neuron and behavior at different time lags in the trial. These weight functions a(u) are called impulse response functions. The boundaries of the time lag u, m and n in the equation are set by evaluating the impulse response functions statistically over data sets, and are chosen such that statistically significant coupling is captured by the equation. We calculate impulse response functions and the **Y**-intercepts b using a frequency-domain method that is described elsewhere in more detail (8, 33).

To evaluate the linear relationship between the neural activity and behavior, such as hand position, coherence spectra are calculated (see Wessberg *et al.* (8)). Spectra and cross-spectra for signal pairs are calculated



**FIG. 9.** (A, C) The Shape Tape in two different configurations. The white velcro strips at the end of the blue "tape" attach the tape to the New World monkey's wrist (farthest left strip in (A)) and upper arm. The pen in (A) provides scale. (B, D) Corresponding 3D reconstructions of the Shape Tape configurations.

by Fast Fourier transforming segments of neuronal firing frequency and continuous hand position data, such as over experimental trials, and averaging over all trials. The coherence spectrum is

$$|R_{xy}(\lambda)|^2 = rac{|f_{xy}(\lambda)|^2}{f_{xx}(\lambda)f_{yy}(\lambda)}\,,$$

where  $f_{xy}(\lambda)$  is the cross-spectrum between the neural and behavioral signals *x* and *y*, and  $f_{xx}(\lambda)$  and  $f_{yy}(\lambda)$  are the power spectra for the neural and behavioral signals, respectively. Thus, the coherence spectrum is the squared absolute cross-spectrum between the neural and behavioral signals, normalized by their autospectra. On a scale from zero to one, the coherence spectrum represents the degree of linear coupling between the neural and behavioral signals as a function of frequency; and the significance of the obtained linear relation can be evaluated using standard statistical methods (33, 34). In the case of 3D behavioral data, the coherence spectrum is calculated separately for each behavioral dimension. The coherence spectra are also used to calculate the impulse response functions, which in turn are examined to characterize the relationships between the neurons and stimuli or behavior in even greater detail (8, 33).

4.4. Categorical and Nonlinear Analysis Using Artificial Neural Networks

Artificial neural networks (ANNs) are one of the more powerful tools for investigating how well spatiotemporal patterns of neural ensemble activity relate to sensory stimuli or to behavior (32, 47; see Nicolelis *et al.* (9)). ANNs are well suited to use in large multivariate problems where the significance of each variable is unknown, because they are efficient at recognizing subtle or complex patterns in the data; and unlike many other analytical techniques, they require no *a priori* assumptions about the structure or distribution of the data. Indeed, ANNs have consistently outperformed other classification techniques we have used to analyze neural ensemble firing patterns.

We use ANNs in two main types of analyses that determine the extent neuronal activity relates to (i) categorical stimulus or behavioral variables such as the direction of a monkey's reach, with, for example, two directions defined as left and right, or (ii) continuous stimulus or behavioral variables such as the continuous hand position of a monkey in 3D space during reach. Here we provide a brief description of the specific ANN we typically use for each of the two main types of analysis (see Nicolelis *et al.* (9) for a more thorough introduction to the use of ANNs for neuronal ensemble analysis).



**FIG. 10.** The general order of procedures we follow to analyze the neuronal ensemble data. Each item is described in the text. Note that we list the preprocessing procedures as a distinct step in the analysis process; however, the specific preprocessing steps used for each type of analysis will be described in the section describing the particular analysis.

To assess the relationship of neuronal ensemble activity with categorical stimulus or behavioral variables, such as left versus right reach directions, we have used several types of ANNs (see Refs. 9, 35–37, 48), but here we describe the optimized learning vector quantization (OLVQ) ANN, a competitive network that is supervised during training, which we most commonly use (available in the Matlab Neural Network Toolbox) (7–9, 32, 47, 48). The OLVQ ANN is a multi-layer, feedforward network with full connectivity (see Figs. 14A and 14C) (35, 48). The input layer is defined by the data input vector. The hidden layer contains two artificial neural units (ANU) for each stimulus or response class to be classified. For instance, if a left or right arm movement is to be predicted, there are four hidden units. Each



FIG. 11. (A) A raster plot of neuronal activity for one neuron in the ensemble of recordings from somatosensory cortex SII in the owl monkey after mechanical stimulation of a finger. Starting at the top, each row displays each discharge of the neuron during one trial. (B) A cumulative frequency histogram (CFH) showing onset of firing. Each point is the cumulative frequency (the sum of the previous frequencies) minus the product of two terms, the overall average firing frequency multiplied by the number of previous bins. Latency is shown by the arrow and is identified by isolating the exact time in which the firing frequency distribution diverges from a random distribution with P < 0.01 (using the Kolmogorov–Smirnov test). (C) The peri-stimulus time histogram (PSTH) for the neuron. The *x*-axis represents time around the stimulation of the hand, and the y-axis depicts firing frequency (alternatively, the y-axis may depict the number of neuronal discharges, i.e., "spike" counts, per time bin). (From Nicolelis, M. A. L., Stambaugh, C. R.,. Brisben, A., and Laubach, M., 1999, Methods for Simultaneous multsite neural recordings in behaving primates. In Nicolelis, M. A. L. (Ed.), Methods for Neural Ensemble Recordings. Boca Raton: CRC Press. pp 121-156. With permission.)

ANU has a corresponding weight vector, and each ANU's output value from the hidden layer is determined competitively. The ANU with the weight vector nearest in Euclidean distance to the input vector is the "winning" ANU, sending a value of 1 to the output layer, with the "losing" ANUs sending 0. The output layer is a linear layer with one ANU per response class: one ANU for left movements and one ANU for right movements. The hidden layer ANUs must learn to assign neural inputs to the left or right movement direction, so two hidden units are designed to "choose" the left direction, and two to "choose" the right direction. This is achieved by setting the weight between the hidden unit and its assigned output unit to be 1, with the weight to the other output 0. The hidden ANUs must then learn to recognize the input vectors that correspond to the hidden unit's designated output ANU: if the neural activity associated with movement to the left is input to the network, the "left" designated ANUs must learn to recognize this neural activity pattern. Because this learning is supervised, the ANUs are rewarded for recognizing correct patterns, and penalized for recognizing incorrect ones, using Kohonen's learning rule (35; also see Demuth and Beale (48)). The weight vectors are moved closer to the input training vectors that are classified correctly and away from input training vectors that are classified incorrectly. In this way the ANUs learn to recognize the appropriate input vectors.

There are three main steps in using the OLVQ ANN: initialization, training, and testing. During initialization, the weight vectors for all ANUs are initialized to the same value, but each ANU is assigned a particular class that must be recognized, for instance, whether the monkey should reach left or right. During training, a set of neural signals are repeatedly input to the network and the ANU must learn to recognize the patterns in the neural signals that corresponds to the ANU's preassigned class, whether it be left or right movements. Finally, the OLVQ network is tested by presenting the network with trials that were not used during training to determine how well the network is able to predict the correct corresponding class, whether the monkey's movement was to the left or to the right.

We regularly use the leave-one-out cross-validation method for training and testing of the network (7, 49). In this method the network is trained on all trials making up the data set, except for one trial that is set aside, or left out. After training, the network is evaluated on how well it predicts, for instance, reach direction from the neuronal activity on the "left-out" trial. This procedure is repeated so that each trial in the data set takes a turn being left out. This method provides a thorough evaluation of the performance of the network, and thus of the potential relationship between the neuronal activity and the stimulus or behavioral variable of interest.

Just as in the linear time series model, the ANN can take into account the possible coupling of neural signals to a stimulus or behavior at times earlier or later than the actual time of stimulus onset or behavioral responses. Thus neural activity around the time of the stimulus or behavior is presented to the neural network. We typically organize the neural signals into single-trial peri-event histograms of a particular bin size (e.g., from 2- to 500-ms bins) and then normalize these





**FIG. 12.** (A) Population peri-stimulus time histograms (PPSTHs) of the responses of an ensemble of somatosensory cortex area SII neurons to the mechanical stimulation of a finger tip (a) and the back of the hand (b) of an owl monkey. Individual neurons line the *x* axis, the peristimulus time is depicted on the *y* axis, and the magnitude of neuronal firing in number of standard deviations away from the spontaneous firing rate is depicted along the *z* axis. (B) Spatiotemporal population map (SPM) depicting the spatial pattern of a neural ensemble in SII cortex of a monkey as it varies over time in response to the mechanical stimulation of the tip of a finger. The diagram in the lower half of B illustrates how each graph in the map is constructed. The *z* axis depicts the magnitude of neuronal firing in terms of the number of standard deviations away from the spontaneous firing rate (A & B from Nicolelis, M. A. L., Stambaugh, C. R., Birsben, A., & Laubach, M., 1999, Methods for simultaneous multsite neural recordings in behaving primates. In: Nicolelis, M. A. L. (Ed.), Methods for Neural Ensemble Recordings. Boca Raton: CRC Press. pp 121–156. With permission.)

spike counts, so neurons with a relatively high baseline rate of firing do not dominate the analysis. For evey trial in the experiment, then, each neuron has a sequence of bins that represents the sequence of discharges of the neuron across the trial (see Fig. 14B). And for each trial, the input vector to the neural network is composed of the spike counts for every bin in the sequence, for every neuron in the ensemble. That is, on every trial, the input vector to the neural network is

$$I = \{C_{N_1t_1}, C_{N_1t_2}, \dots C_{N_1t_k}, C_{N_2t_1}, \dots C_{N_2t_k}, \dots C_{N_rt_k}\},\$$

where  $N_i$  represents each neuron,  $t_i$  represents each time bin in the trial, and  $C_{N_i t_i}$  represents the spike count for neuron  $N_i$  in time bin  $t_i$  (with r equal to the total number of neurons and k the total number of time bins in each trial).

For continuous stimulus or behavioral variables, such as hand position during a monkey's reach, a nonlinear analysis using ANNs is typically conducted after a linear relationship between the neurons and behavior has been characterized. Since ANNs can detect linear and nonlinear relationships between inputs and outputs, the performance of the ANN can be compared with that of the linear model, and the extent to which the ANN outperforms the linear model suggests that a nonlinear relationship between the inputs (neurons)



**FIG. 13.** Data matrix configuration and model for the linear time series analysis. **Y**(*t*), the predicted behavioral response, and  $\varepsilon(t)$ , the residual errors, are matrices, with three columns corresponding to the three dimensions; and **b**, the **Y** intercept, consists of three values, one for each dimension. Finally, there is one impulse response function, **a**, for every neuron and dimension combination (e.g., with 50 neurons, there are 150 impulse response functions).

and outputs (behavior) has been detected by the ANN (see Ref. 8).

We generally use a two-layer feedforward backpropagation network (Figs. 14A and 14C) (see Refs. 8, 9, 48). Again, the input layer is defined by the data input vector. We have had good success with 15-20 ANUs in the hidden layer, and we use 1 or 3 ANUs in the output layer for 1D or 3D hand position prediction, respectively. The hidden-layer ANUs have a sigmoidal tangent transfer function to the output layer, while the output-layer ANUs have a linear transfer function to enable the output from the network to take on any value (Figs. 14A and 14C) (48). We have used several different learning rules to train the weights in the network, but we have found particularly good success with the Powell-Beale conjugate gradient algorithm (48, 50). Additionally, an early stopping rule can be used to minimize overfitting of the data, and thus to provide better generalization to new neuronal inputs.

## 4.4.1. Quantifying the Results of the ANN with Information Theory

To evaluate the performance of the ANN, and thus to evaluate the extent of the relation of neural activity with a continuous stimulus or behavioral variable, such as hand position, we generally assess the size of the correlation coefficient between the predicted and actual hand positions (e.g., Ref. 8). For categorical stimulus or behavioral variables, this success is typically quantified as the percentage of correct predictions and the extent to which this percentage is statistically above chance. But there is another way to quantify the results obtained from the ANN, providing a more sensitive measure than percentage correct. The ANN can be thought of as measuring the amount of information about a sensory stimulus or behavior, such as reach direction, that is contained in or conveyed by the neural activity per trial. The amount of information can be calculated from the confusion matrix generated from the ANN (see Laubach et al. (7) and Pierce (54)). For example, if there are 100 trials in the data set, with 50 trials of the monkey reaching left and 50 trials of the monkey reaching right, with the ANN making 75% correct predictions for reaching left (i.e., 45/(45+15)) and 88% for reaching right (i.e., 35/(35+5)), the confusion matrix is

		Actual reach direction	
		Left	Right
Predicted reach direction by the ANN	Left Right	45 5	15 35

where the actual directions are those obtained during

the experiment for the 100 trials, and the predicted directions are those output by the ANN when presented with the neural activity recorded during the same 100 trials of the experiment.

Information is calculated by determining how much uncertainty about which direction the monkey will reach before the trial will be reduced when the monkey reaches. More specifically,

$$I = H_{\rm a} + H_{\rm b} - H_{\rm ab},$$

where *I* is the information about reach direction in the neuronal activity that is extracted by the ANN, and the  $H_{is}$  are "entropy" values that measure the specific uncertainties that are involved.  $H_{a}$  is a measure of the

A. The Artificial Neural Unit (ANU)



B. The Data Matrix for the ANN





**FIG. 14.** (A) The general model for an artificial neural unit (ANU). In competitive networks, such as OLVQ, the transfer function obtains the distance between the input and weight vectors and adds the bias value. The outputs from the transfer functions of all ANUs are then compared, and the ANU with the smallest distance between the input and weight vectors "wins." The final output of the winning ANU is set to 1, while all others are set to 0. For backpropagation networks, the transfer function receives as input the sum of the multiplied input and weight vector elements  $w_i t_i$  (that is  $w_1 t_1 + \cdots + w_n t_n$ ) added to the bias value, and the transfer function outputs the value *a* according to the defined function (such as sigmoidal tangent). (B) Data matrix configuration for the analyses using ANNs.  $N_i$  are the neurons in the ensemble,  $t_i$  are the time bins of spike counts for each neuron within each trial, and  $T_i$  are the trials in the experiment. Each row represents the input vector for each trial presented to the ANN. (C) The general model for the ANNs. See text for details.

uncertainty of whether the monkey will reach left or right on any given trial and is calculated from the confusion matrix above as

$$H_{\rm a} = \sum_{\rm a} - p_{\rm a} \log_2 p_{\rm a} = -1*50/100*\log_2(50/100) + -1*50/100*\log_2(50/100) = 1 \text{ bit.}$$

where  $p_a$  is the probability of the reach direction, and the sum is taken over the two actual reach directions. This value defines the amount of total information that is in the problem. Since each actual reach by the monkey totally resolves the uncertainty for that trial, the entire problem (of not knowing what the reach will be prior to any given trial) is solved with 1 bit of information. It is a 1-bit problem. Note that the total information of the problem would be reduced if the monkey had a bias toward one of the reach directions or if the numbers of trials for the two directions were not equal, since these biases slightly reduce the uncertainty of which direction the monkey may reach.

 $H_{\rm b}$  is the uncertainty of whether the ANN will predict that the monkey's reach will be left or right and is calculated from the confusion matrix as

$$H_{\rm b} = \sum_{\rm b} - p_{\rm b} \log_2 p_{\rm b} = -1*60/100*\log_2(60/100) + -1*40/100*\log_2(40/100) = 0.97095,$$

where  $p_b$  is the probability of the predicted reach direction, and the sum is taken over the two predicted reach directions.  $H_b$  can be thought of as measuring the bias of the ANN, falling less than 1 if the network tends to predict one of the directions more than the other.  $H_{ab}$ is the joint uncertainty of (a) which direction the monkey will reach and (b) which reach will be predicted by the ANN.  $H_{ab}$  measures whether the ANN predicts correctly, increasing with the errors of the ANN.  $H_{ab}$ , and finally, I are calculated as

$$\begin{split} H_{\rm ab} &= \sum_{\rm a} \sum_{\rm b} - p_{\rm ab} \log_2 p_{\rm ab} = -1*45/100*\log_2(45/100) \\ &+ -1*15/100*\log_2(15/100) + -1*5/100*\log_2(5/100) \\ &+ -1*35/100*\log_2(35/100) = 1.6751, \end{split}$$

where  $p_{\rm ab}$  is the probability of the combination of the actual and the predicted reach directions, the sums are taken over the actual and the predicted reach directions, and

$$I = H_{\rm a} + H_{\rm b} - H_{\rm ab} = 1 + 0.97095 - 1.6751$$
  
= 0.2959 bit.

Thus, 0.2959 bit is the amount of information found by the ANN in the neuronal activity regarding the monkey's reach direction. It is 29.6% (i.e., 0.2959/1\*100) of the total information in the problem.

Thus, information theory provides another way to quantify and to conceptualize the structure in the neuronal ensemble firing patterns that relates to the measured behavior or stimulus. It is especially useful in comparisons of different potential neuronal coding schemes. For instance, the amount of information found in the correlated firing patterns of neurons can be compared to the amount of information found in the average firing rates of individual neurons to determine how much more information may be found in the interactions between neurons (see Laubach *et al.* (7)). Additionally, using information theory as a means to quantify the success of the ANN at finding a relationship between neuronal activity and stimuli or behavior is superior to percentage correct when the number of trials of each response or stimulus class is highly unbalanced or when the animal has a significant response bias (see Laubach *et al.* (7) for a detailed discussion).

#### 4.4.2. Other Analyses Using ANNs

Combined with particular manipulations of the data, ANNs can also be used to investigate the potential encoding strategies of neural populations. These manipulations include (a) trial shuffling, (b) spike shifting, (c) systematic changes in the bin sizes, or bin clumping, and (d) systematic removals or additions of individual neurons in the ensemble, known as neuron dropping or adding (see Nicolelis et al. (9) for a more detailed discussion; Ghazanfar et al. (47); Nicolelis et al. (32)). In trial shuffling, the network is first trained on the original data set, and then the trials for each neuron from the same stimulus or behavioral response class, such as all trials for left reaches, are randomly shuffled so that the network receives an input vector formed by single neuron firing patterns obtained in different trials. This manipulation tests whether correlated firing of neurons is due to an actual interaction between neurons or simply due to independent but similar firing patterns in response to the stimulus or behavioral event. When the trials are shuffled, interactions between the neurons are disrupted, while the eventrelated activity remains unchanged. The difference between the success of the ANN before and after this manipulation measures the extent that correlated firing between neurons reflects a true interaction between the neurons.

In spike shifting, the ANN is first trained with the original spike trains obtained in the experiment. Then the ANN predictions are examined when the ANN is fed with spike trains on a given trial that are individually shifted earlier or later in time (see Fig. 15A) (6). This spike train shifting changes the specific timing of the spikes relative to the stimulus or behavioral event and relative to other neurons while keeping the overall firing pattern and overall firing rate across the trials unchanged. This manipulation thus tests whether such specific timing relative to the experimental event is significant; at the same time, spike shifting tests the specificity of coupling between neurons, measuring the potential reduction in the success of the ANN after the lag between neurons is changed as a result of the shifting procedure. To determine whether neurons in the ensemble are actually interacting, the effect on the

ANN predictions due to neuron decoupling must be separated from the effects due to changing the relative timing of the firing to the experimental event. One way to do this is to compare the results from the random spike shifting procedure with the results from a procedure in which the spike trains for all neurons are shifted the same amount—for instance, by the median shift of the random shift procedure. In this "median shift" manipulation, the coupling between the neurons is held constant (since all spike trains are shifted identically), whereas the specific timing of firing relative to the experimental event has been disrupted by the shift (6).

Bin clumping helps to identify temporal attributes of single neuron responses to the experimental events. For instance, fine temporal patterns in the neural discharges of cells can be lost if the neural activity is organized into bin sizes that are too large. Indeed, we continue to find that finer temporal patterns of neural discharge routinely harbor information about the experimental events, such as primate arm movements,



# A. The Spike Shifting Procedure

**B. Example Neuron Dropping Curves** 



#### Number of Neurons

**FIG. 15.** (A) An example of how spikes are shifted in the spike shifting procedure. The spike train for each neuron in a trial is randomly shifted left or right a random amount. (B) Hypothetical neuron dropping curves obtained by randomly removing neurons, one at a time, from an ensemble of neurons in dorsal premotor (A) and primary motor (B) cortex, and presenting the remaining neurons to the ANN to obtain the amount of information about hand position in the remaining neurons.

and this information is lost when these finer temporal patterns are not measured, such as when only overall firing rates are calculated. To be sure, however, there is a limit to how fine the temporal patterns can be before information again begins to be lost (7).

Finally, "neuron dropping" or "adding" analysis is a straightforward and powerful way of studying the relative contributions of specific neurons to the representation of the behavior or sensory stimulus (8, 32, 47; see Nicolelis et al. (9)). In addition, it allows us to assess the information in neural ensembles as a function of the number and types of neurons in the ensemble. Thus, the procedure also provides an elegant way of comparing different brain areas by conducting the analysis on neurons from specific areas separately and then comparing the results generated for each brain area. For instance, to analyze the relative information about primate hand position in dorsal premotor and primary motor cortex, Wessberg et al. (8) began with the entire ensemble of neurons from dorsal premotor cortex and then randomly removed the neurons, one at a time, from the ensemble, presenting the remaining neurons back into the ANN, which produced a neuron dropping curve similar to curve A in Fig. 15B. The same procedure was conducted with primary motor neurons, yielding a curve similar to curve B in Fig. 15. With such results revealed by the neuron dropping analysis, at least two important observations can be made: (i) that both brain areas contain information about hand position, suggesting that this information is widely distributed; and (ii) that the areas nevertheless differ in how much information they contain about hand position, suggesting that some specialization does exist in these brain areas. Finally, the neuron dropping curves can be fit by equations that can then be used to estimate the number of neurons one would need to record from to obtain and arbitrary amount of information about hand position, for instance, to obtain 0.9 bit of information (see Ref. 8). These predicted numbers of neurons for each brain area provide another means of comparing different brain areas; they also provide intriguing estimates of the number of neurons necessary for useful applications, such as for prosthetic devices (11).

# 4.5. Wavelet Packet Analysis and Temporal Patterns of Firing

Wavelet analysis is a relatively new and promising signal processing technique that decomposes a neural signal into components. The analysis can remove irrelevant neural activity, significantly reducing the size of large neuronal data sets. But perhaps more importantly, the analysis helps one to determine what specific aspects of the neural signal are related to a stimulus or behavior, such as to the deflection of a set of whiskers of a rat or the direction of reach of a primate. For instance, if a particular frequency component of the signal at a particular time during the trial is related to a behavioral response, such as reach direction, wavelet analysis can help to find this relationship by capturing this part of the neuronal signal as a wavelet feature.

To conduct a wavelet packet analysis on a neuronal data set, single-trial peri-event histograms of 1-ms bins are first smoothed over intervals of 5, 10, or 20 ms (i.e., low-pass filtered and resampled using Matlab routines, fir1, decimate, and filtfilt) and then normalized, again to curb the influence of neurons with a relatively high baseline rate of firing. The resulting processed spike trains are then decomposed using a wavelet method (available as public domain code for Matlab) (7).

The wavelet algorithm we use was developed by Buckheit and Donoho (39) and works in the following way. First, for each neuron, the normalized firing frequencies are averaged across all trials for the left direction, and again for the right direction (Fig. 16A). The average activity from right-direction trials is then subtracted from the averaged activity from left-direction trials, producing a "difference vector" (Fig. 16B). Wavelet features are then extracted from this difference vector, to obtain features from the neuronal signal that are potentially relevant to the discrimination between reach directions. To do this, a wavelet packet table is generated from the difference vector. The table is obtained by passing the difference vector through a lowpass filter and a high-pass filter, and downsampling (i.e., obtaining half the original number of signal values), generating two resulting signals such as in Row 2 of Fig. 16D (for descriptions of the filtering process, see Refs. 39, 40, 51). The filters are based on the Daubechies 4-point wavelet, which is a waveform of a particular shape and limited duration, depicted in Fig. 16C. In the filtering process, the limited-duration filters are passed successively across the "difference vector" signal, obtaining an amplitude value, or coefficient, at each time point in the table (in Row 2 of Fig. 16D) that represents the degree to which the shape of the filter matches the difference vector at that point in time in the trial.

These resulting signals (i.e., in Row 2 of Fig. 16D) are themselves passed through the filters, generating the signals in Row 3 of Fig. 16D. This procedure is repeated a number of times based on the number of bins per trial. Because the filters are essentially "stretched" at each iteration in the wavelet analysis procedure, each row in the resulting table represents the original signal at a progressively lower resolution, or coarser grain, representing progressively lower frequency attributes of the original neural signal.

The entire set of discrete points making up all of the signals in Fig. 16D constitute the complete wavelet packet table, providing a rich and detailed set of potential features of which to represent the original neural signal. To begin selecting the wavelet features, these points from the table are arranged into a column vector as shown in Fig. 16D. Next, the location in the vector with the largest amplitude value is taken as the location of the first wavelet feature: for instance, position 128 in the column vector in Fig. 16D. The largest amplitude value signifies that the difference vector was best matched by the shape of the filter waveform at that point in the trial, which in turn means that the feature best accounts for the difference between the left and right reach trials. Note that using this amplitude value, or coefficient, the wavelet feature that it represents can



**FIG. 16.** (A) The firing frequencies are averaged across all trials for the left direction, and again for the right direction for an example neuron. (B) A schematic of the "difference" vector that is produced from subtracting the right-direction trials from the left-direction trials. (C) The shape of the Daubechies 4-point wavelet. (D) A schematic of how the wavelet packet table is generated by passing the difference vector through a low-pass filter and a high-pass filter, downsampling (Row 1), and then iteratively passing the products of the filtering process back through the filters (Rows 2, 3, etc.). All of the generated values can be rearranged into a column vector as shown in the middle of D. The position in the column vector with the largest amplitude value, in this case, position 128, with amplitude 0.436, is chosen as the location of the first wavelet feature; and the wavelet feature itself can be reconstructed from the amplitude value as depicted at the right of D.

be reconstructed from the coefficient and a filter based on the "mother" wavelet (the Daubechies 4-point wavelet); an example reconstructed wavelet feature is shown in Fig. 16D.

Since all of the chosen wavelet features must be orthogonal, once a feature is chosen, it must be completely removed from the wavelet packet table before the next feature is selected. Once this procedure is completed, the second wavelet feature is chosen in the same way as the first: the location in the vector with the largest amplitude value, again meaning that the filter shape at that point in the trial is the next best at matching the difference vector, is taken as the location of the next wavelet feature. This procedure of selecting wavelet features is repeated a number of times based on the number of bins per trial. Only the most significant of these selected features are used in subsequent analyses.

This entire wavelet analysis is conducted for each neuron in the ensemble, generating a set of wavelet features that now represent the neuronal signal of each of the neurons. Once the wavelet features are obtained for each neuron, we need to determine how much reach direction information is contained in these features. First, for each trial of the experiment, that is, for each arm reach, the value for each wavelet feature is obtained. To do this, each neural signal over the trial is processed to obtain a wavelet packet table, as described above. Then the amplitude at each selected feature location (i.e., at each location that was previously chosen as the location of a wavelet feature for that neuron) is used as the value for the feature on that trial. For example, once the wavelet packet table is obtained for Neuron 1 on Trial 1, and the table is rearranged into a column vector, the value for the first feature would be the amplitude at position 128 in the column vector of Fig. 16D. Thus, one value is obtained per feature per trial; if there are, for instance, 100 trials and 80 neurons that are represented by a total of 175 wavelet features, the resulting data matrix will be a 100 trials by 175 features matrix of feature values on all of the trials (compared to the original data matrix of 100 trials by 3200 columns, assuming 80 neurons with 40 bins per trial).

This 100 trials by 175 features data matrix is then fed to the OLVQ ANN, trial by trial, to determine the amount of reach direction information that is contained in the features. Thus, on each ANN learning trial, 175 wavelet feature values from a given trial in the experiment (i.e., one row of the data matrix) are fed to the network, and the ANN must learn to predict the reach direction based on the wavelet feature values. We have consistently found that the wavelet features tend to harbor a great deal of information regarding stimuli or behavior, providing evidence that stimuli and behavior may be encoded in specific firing patterns of the neuronal activity at specific times during the trial (7).

#### 4.6. Multivariate Statistical Techniques

Multivariate techniques have proven exceptionally useful in analyzing large neuronal ensemble data sets. Not only have these techniques helped in managing the large data sets, by, for instance, reducing the high dimensionality, they have helped to reveal the nature of the relationships of the neurons to each other and to sensory stimuli and behavior.

#### 4.6.1. Principal Component Analysis (PCA)

PCA reduces the large numbers of original neural signals to a smaller number of derived "components" that account for most of the variance observed in the original data set (e.g., Ref. 41). These components represent dimensions of information embedded in the firing pattern of the neural population, and they may reflect functional associations between the neurons in the ensemble. In our experience, neurons with similar functional characteristics, such as those related to the deflection of a specific whisker of a rat, or neurons located in the same area of the brain do tend to have high coefficients on the same principal components, whereas neurons with dissimilar functional associations or from different areas of the brain tend to be clearly separated onto different principal components (e.g., Refs. 9, 32).

Each recorded neuron is treated as a separate variable in the PCA. Time series of the firing rates of each neuron (e.g., rates obtained in 10- to 25-ms bins over 10- to 30-min periods) are correlated with those of all other neurons in the population, generating a correlation matrix of all neurons. From this correlation matrix, a series of principal components is extracted. Each of these components is formed by the weighted linear sum of the firing patterns of individual neurons, with neurons contributing differentially to the different components, as reflected in the component weights. Note that neither the response properties of the neurons nor their anatomical location are made available to the PCA algorithm; nonetheless, as described above, these features are often pulled out by PCA.

# 4.6.2. Independent Component Analysis (ICA) and Correlated Neuronal Firing

One of the most important analyses to be conducted on neural ensemble data is a study of the potential correlations or interactions between neurons. Recently, ICA has been successfully applied to neural ensemble data, and it is currently an important statistical tool that we use to evaluate the degree and type of correlated firing patterns between single units (6, 7). ICA is a statistical technique that decomposes a large data set of mixed signals into separate, independent signals (42, 43) (See Fig. 17A), and because the activity of individual neurons or groups of neurons could in some cases be linear mixtures of independent signals, the neuronal data may be appropriate for ICA analysis. For example, one subset of neurons in the primary motor cortex of primates may be influenced by activity from one source of neurons in the thalamus, while another subset of neurons may be influenced by activity from neurons in other cortical areas, such as from the dorsal premotor cortex, while third subset of neurons may receive inputs



# B. ICA Procedure



FIG. 17. (A) A schematic illustrating the basic concepts underlying our use of ICA. In this conceptualization, one subset of neurons in the primary motor cortex of primates may be influenced by activity from one source of neurons in the dorsal premotor cortex, another subset of neurons may be influenced by activity from neurons in the thalamus, and a subset of these primary motor neurons may receive inputs from neurons in both the dorsal premotor cortex and thalamus. The highly schematized graphs to the right illustrate that the influence of the two sources on the firing pattern of a neuron may be separable. (B) For the ICA algorithm, the original data matrix of spike counts for each neuron in each time bin on every trial is multiplied by the matrix of principal component (PC) coefficients of each neuron on each PC, resulting in a matrix of PC scores for each bin on each trial (see Ref. 41). Blocks of rows of this data matrix are then fed to a single-layer feedforward ANN that will result in output ANUs that correspond to the independent components (see text for details). (C) To obtain the coefficients for each neuron on each independent component (IC), a matrix of the IC coefficients (the weights from the ANN) for each PC is multiplied by the matrix of PC coefficients for each neuron. This result is displayed in a Hinton diagram that depicts the size of the coefficient for each neuron on each IC. The size of each rectangle corresponds to the size of the IC weight (i.e., coefficient). The black color represents a positive value; white is negative. Neurons with large coefficients on the same ICs have strong correlated firing patterns. (D) The data matrix that was used to train the ANN is multiplied by the matrix of IC weights obtained from the ANN to produce a new data matrix in which each row represents each bin on every trial, each column represents an IC, and each value in the matrix represents the number of spikes due to each independent source that occurred during each time bin on each trial. A wavelet packet analysis is then conducted on the data matrix, and the resulting wavelet features are then fed to an OLVQ ANN to assess the amount of information that the ANN can extract from the wavelet features, derived from the independent components.

from neurons in both the thalamus and dorsal premotor cortex (Fig. 17A).

Like PCA, ICA decomposes the original neural signals into a set of components that may better characterize the overall neuronal activity. The principal components from PCA decorrelate the neural signals, but they do not necessarily separate independent sources. If the probability distributions of the signal sources (i.e., of the source firing rates) are gaussian, the principal components will be independent (see Lee (52) and Johnson and Wichern (41)). However, if the probability distributions of the sources are not gaussian, the principal components of PCA will not be independent. This is because PCA decorrelates the second-order statistical property (i.e., variance), but not higher-order statistical properties (such as kurtosis, which describes the size of the tails and peakedness of the distribution). Thus, signals separated by PCA may still be dependent; PCA, therefore, may not be the most accurate method for assessing the degree of correlated firing between neurons. ICA, on the other hand, decorrelates the higher-order statistical

#### C. Procedure to obtain IC weights for each neuron

properties as well, providing a truly independent separation of signals. Thus, with ICA, neurons with correlated firing patterns will be identified because they will have significant coefficients on the same independent component.

The main steps we use to conduct ICA on our neuronal data sets are the following. First, single-trial perievent histograms of 1-ms bins are smoothed over intervals of 5, 10, or 20 ms and then normalized. Next, PCA is conducted on the data set, reducing the dimensionality of the data and separating the signals with respect to the second-order statistical property (i.e., variance). To take the most significant principal components, only those with variance greater than one are retained (7, 41). One of the methods for ICA is then applied to the data (ICA algorithms are available as public domain code for Matlab; see Laubach et al. (6)). Here, we describe the extended ICA algorithm that we commonly use (6, 42, 43, 52). For the ICA algorithm, the data are arranged in a matrix in which each row represents each bin on every trial (with the total number of rows equal



# D. Procedure to determine Information content in the IC scores



FIG. 17—Continued

to the number of trials times the number of bins per trial), and each column represents a principal component (PC), with each entry in the matrix being the particular PC value or score for the bin on a given trial (see Fig. 17B). As described in the previous section, the PC score is the weighted sum of the firing patterns of individual neurons during each time bin on each trial, with neurons contributing differentially to the different components, as reflected in the component coefficients. Blocks of rows of this data matrix are then fed to a single-layer feedforward ANN with a hyperbolic tangent transfer function, which transforms the outputs according to the specified function (Fig. 17B) (48, 52). The number of inputs to the ANN, as well as the number of ANUs in the single output layer, which will become the independent components, is equal to the number of principal components. The weights of the output layer are adjusted according to the extended ICA information maximization (infomax) learning rule, resulting in output ANUs corresponding to each independent component. Bell and Sejnowski (42) showed that maximizing the joint entropy (i.e., uncertainty) of the outputs of the ANN results in a minimization of the mutual information between the outputs, producing independent output ANUs. The infomax learning rule used to change the ANU weights, therefore, was derived by maximizing the joint entropy of the outputs with respect to the weights.

After the training of the ANN, with each output ANU now corresponding to an independent component, three subsequent analyses are conducted. First, the weights from the ANN are the coefficients that characterize the extent to which each neuron reflects the independent component (IC)—that is, characterizing the extent the neurons are firing due to a common external source (Fig. 17C). These coefficients are examined using Hinton diagrams to characterize the patterns of correlated firing in the neuronal ensembles: neurons with nonzero coefficients on the same independent component have correlated firing patterns (Fig. 17C). Second, the weights from the ANN are used to determine the amount of information content regarding stimuli or behavior, such as reach direction, that can be found in the correlated firing patterns of the neurons. To do this, the data matrix that was used to train the ANN is multiplied by the matrix of IC weights obtained from the ANN to produce a new data matrix in which each row represents each bin on every trial, each column represents an IC, and each value in the matrix represents the number of neuronal discharges (i.e., spikes) due to each independent source that occurred during each time bin on each trial (Fig. 17D). To extract as much information about reach direction as possible from the correlated firing patterns and also to reduce

the dimensionality of the data, a wavelet packet analysis is then conducted on the data matrix (see Section 4.5). The resulting wavelet feature values are then fed to an OLVQ ANN to assess the amount of information that the ANN can extract from the wavelet features, derived from the independent components (Fig. 17D). The result quantifies the amount of information regarding a stimulus or behavior that is extracted by the ANN from the correlated firing of neurons in the ensemble.

Finally, a spike shifting analysis is conducted to examine whether the correlated firing patterns obtained from ICA reflect an actual interaction between neurons as opposed to simply reflecting similar response properties of the neurons to the stimulus or behavioral event (see Section 4.4.2).

#### 4.7. Real-Time Analysis of Neuronal Ensembles

In this final section on neural ensemble analysis, we describe algorithms and procedures that are being developed in our laboratory to conduct such analyses in real time to, for instance, predict the hand movements of monkeys from their neural activity immediately before the monkeys make the movements (8). Successful real-time predictions of behavior from cortical activity are critical for the development of such devices as prosthetic limbs, which may aid people with debilitating motor impairments (11).

Here we briefly describe how a linear analysis of neural signals and hand position can be conducted in real time (see Wessberg *et al.* (8)). First, recorded neuronal spike trains are sorted into 100-ms bins (10 Hz). To account for the influence of neural signals that occur earlier than the actual hand position signals, ten 100ms bins are obtained for each neuron, corresponding to lags up to 1 s. The linear model for real-time neural signals and hand position then is

$$\mathbf{Y}(t) = \mathbf{b} + \sum_{u=0}^{9} \mathbf{a}(u)\mathbf{X}(t-u) + \varepsilon(t),$$

where *t* and *u* are in units of 10 per second. Thus 10 values are obtained for each neuron, corresponding to the 10 time lag steps. In this case, the weights **a** and **Y** intercepts **b** are calculated using standard linear regression techniques (e.g., Ref. 41). We have found that the real-time predictions of hand position using this simplified model tend to be only slightly inferior to offline predictions using the linear model described above (see Section 4.3).

Because our real-time system was designed to be adaptive, the weights of the linear model can be continually updated throughout the experimental session. A separate computer is used to fit the model continually, using the previous 10 min of recordings, and the resulting updated weights are sent back to the data acquisition computer to predict the hand position "instantaneously" using the currently recorded neural firing patterns. Finally, the predictions from both this linear model and ANN models (such as the backpropagation ANN described in Section 4.4) can be broadcast in real time over the internet (using TCP/IP protocol) to a computer that can use the signal to control devices, such as a robot arm, in real time (8).

# 5. CONCLUSION

The growing excitement generated from the ability to record from populations of neurons across many brain areas in awake, behaving animals has increased the demand for improved methods for sampling, monitoring, and analyzing the electrophysiological data from these neural ensembles. In this article, we attempted to describe in sufficient detail many of the latest issues and the specific solutions we have been achieving. But this is only the beginning. It is clear that in the near future, movable multielectrode bundles and arrays will be used in primates, and the technology will need to be extended so that these movable systems can remain in place for long periods, perhaps up to years, before being moved again to record from a new set of neurons. Surgical techniques must continue to be developed to minimize damage during multielectrode implantation, as well as to reach targeted brain structures more accurately.

Technologies must also be developed to continue to miniaturize all components of the multichannel recording system; and work toward wireless recording systems should continue (53). And finally, there are many avenues to be explored in the challenging task of analyzing larger and larger amounts of data. One promising area, for instance, is the extension of time series analysis to such techniques as PCA. Time series analysis allows one to take time dependencies into account, such as when cortical activity at time *t* affects muscle (EMG) activity at time  $t + x\Delta t$ , due to the inherent conductance delays between cortical neuronal responses and those of the muscles. And then beyond all this, there is the hope that in the future, chronic multielectrode recording techniques could be used in human applications, particularly when severe impairments could be alleviated such as in the case where the brain could control a prosthetic arm (8, 11).

To reach these goals, we need to make advances like those presented in this article. As the basic methods in multielectrode physiology continue to improve, it promises to continue to vitalize the neurosciences, providing the ability to study the underlying nature of such higher brain processes as learning, problem solving, and emotion.

# ACKNOWLEDGMENTS

We thank Mark Laubach, Johan Wessberg, Pamela D. Beck, Chris Stambaugh, Larry Hawkey, and Michael Wiest from the Department of Neurobiology at Duke University, as well as Sherry W. Mohr and Richard L. Auten from the Department of Pediatrics at Duke University Medical Center. This work was made possible by grants from NIH (J.D.K., D.J.K., D.B.K., M.A.L.N.), the Human Frontier Foundation (D.C., M.A.L.N.), NSF, DARPA and ONR to M. A. L. Nicolelis.

# REFERENCES

- Lilly, J. C. (1958) *in* Biological and Biochemical Bases of Behavior (Harlow and Woolsey, Eds.), pp. 83–100, Univ. of Wisconsin Press, Madison.
- 2. Nicolelis, M. A. L. (Ed.) (1998) Methods for Neural Ensemble Recordings, CRC Press, Boca Raton, FL.
- Krupa, D. J., Brisben, A. J., and Nicolelis, M. A. L. (2001) J. Neurosci. Methods 104, 199–208.
- 4. Chicurel, M. (2001). Nature 412, 266-268.
- Hatsopoulos, N. G., Ojakangas, C. L., Paninski, L., and Donoghue, J. P. (1998) Proc. Natl. Acad. Sci. USA 95, 15706–15711.
- Laubach, M., Shuler, M., and Nicolelis, M. A. L., (1999) J. Neurosci. Methods 94, 141–154.
- Laubach, M., Wessberg, J., and Nicolelis, M. A. L. (2000) Nature 405, 567–571.
- Wessberg, J., Stambaugh, C. R., Kralik, J. D., Beck, P. D., Laubach, M., Chapin, J. K., Kim, J., Biggs, S. J., Mandayam, A. S., and Nicolelis, M. A. L. (2000) *Nature* 408, 361–365.
- Nicolelis, M. A. L., Stambaugh, C. R., Brisben, A., and Laubach, M. (1999) *in* Methods for Neural Ensemble Recordings (Nicolelis, M. A. L., Ed.), pp. 121–156, CRC Press, Boca Raton, FL.
- 10. Eichenbaum, H. B., and Davis, J. L. (Eds) (1998) Neuronal Ensembles: Strategies for Recording and Decoding, Wiley, New York.
- 11. Nicolelis, M. A. L. (2001) Nature 409, 403-407.
- 12. Kubie, J. L. (1984) Physiol. Behav. 32, 115-118.
- Kentros, C., Hargreaves, E., Hawkins, R. D., Kandel, E. R., Shapiro, M., and Muller, R. U. (1998) *Science* 280, 2121–2126.
- Freeman, J. A. (1969) Electroencephalogr. Clin. Neurophysiol. 26, 623–626.
- 15. Hubel, D. H. (1957) Science 125, 549-550.
- Snodderly, D. M. (1973) *in* Bioelectric Recording Techniques, Part A: Cellular Processes and Brain Potentials, (Thompson R. F., and Patterson, M. M., Eds.), pp. 137–163. Academic Press, New York.
- 17. Loeb, G. E., Bak, M. J., Salcman, M., and Schmidt, E. M. (1977) *IEEE Trans. Biomed. Eng.* **24**, 121–128.
- Miller, R. D., Miller, E. D., Jr., and Reves, J. G. (2000) *in* Anesthesia (Miller, R. D., and Miller, E. D., Jr., Eds.), Churchill Livingstone, Kent.
- Warner, D. S., Zhou, J., Ramani, R., et al. (1991) J. Cereb. Blood Flow Metab. 11, 794–802.

- 20. Carpenter, M. B. (1991) Core Text of Neuroanatomy, Williams and Wilkins, Baltimore.
- 21. Al-Mefty, O. (Ed.) (1991) Meningiomas, Raven Press, New York.
- Greenberg, R. W., Lane, E. L., Cinnamon, J., Farmer, P., and Hyman, R. A. (1994) Semin. Ultrasound CT MRI 15, 454–465.
- Alcolado, R., Weller, E. P., Parrish, E. P., and Garrod, D. (1988) Neuropathol. Appl. Neurobiol. 14, 1–17.
- 24. Ramsey, H. J. (1965) J. Cell Biol. 26, 323-333.
- Alberts, B., Bray, D., Lewis, J., Raff, M., Roberts, K., and Watson, J. D. (1994) Molecular Biology of the Cell, Garland, New York.
- 26. Seglen, P. O. (1976) Methods Cell Biol. 13, 29.
- 27. Schomberg, D., and Salzmann, M. (Eds.) (1991) Enzyme Handbook, Springer-Verlag, Berlin/New York.
- 28. Haralson, M., and Hassell, J. (Eds.) (1995) Extracellular Matrix: A Practical Approach, Oxford Univ. Press, New York.
- Porter, R., and Lemon, R. (1995) Corticospinal Function and Voluntary Movement, Oxford Univ. Press, New York.
- 30. Loeb, G. E., and Gans, C. (1986) Electromyography for Experimentalists, Univ. of Chicago Press, Chicago.
- Krupa, D. J., Ghazanfar, A. A., and Nicolelis, M. A. L. (1999) Proc. Natl. Acad. Sci. USA. 96, 8200–8205.
- Nicolelis, M. A. L., Ghazanfar, Faggin, B. A., Votaw, S., and Oliviera, M. O. (1998) *Neuron* 18, 529–537.
- Brillinger, D. R. (1981) Time Series: Data Analysis and Theory. Holden–Day, San Francisco.
- Halliday, D. M., Rosenberg, J. R., Amjad, A. M., Breeze, P., Conway, B. A., and Farmer, S. F. (1995) *Prog. Biophys. Mol. Biol.* 64, 237–278.
- 35. Kohonen, T. (1997) Self-Organizing Maps. Springer, New York.
- Bishop, C. M. (1995) Neural Networks for Pattern Recognition, Clarendon Press, Oxford.
- 37. Krogh, A., Palmer, R., and Hertz, J. (1992) Introduction to the Theory of Neural Computation, Perseus, New York.

- Krippendorff, K. (1986) Information Theory: Structural Models for Qualitative Data. Sage, Thousand Oaks, CA.
- 39. Buckheit, J., and Donoho, D. (1995) Proc. SPIE 2569, 540-551.
- Wickerhauser, M. V. (1994) Adapted Wavelet Analysis from Theory to Software, AK Peters, Wellesley, MA.
- 41. Johnson, R. A., and Wichern, D. W. (1998) Applied Multivariate Statistical Analysis, Prentice-Hall, Englewood Cliffs, NJ.
- Bell, A. J., and Sejnowski, T. J. (1995) Neural Comput. 7, 1129– 1159.
- 43. Lee, T. W., Girolami, M., and Sejnowski, T. J. (1999) Neural Comput. 11, 417-441.
- Nicolelis, M. A. L., Lin, R. C. S., and Chapin, J. K. (1997) J. Neurophysiol. 78, 1691–1706.
- Gerstein, G. L. (1999) *in* Methods for Neural Ensemble Recording (Nicolelis, M. A. L., Ed.), pp. 157–177, CRC Press, Boca Raton, FL.
- 46. de Boor, C. (1999) Spline Toolbox. Mathworks, Natick, MA.
- Ghazanafar, A. A., Stambaugh, C. R., and Nicolelis, M. A. L. (2000) J. Neurosci. 20, 3761–3775.
- Demuth, H., and Beale, M. (1998) Neural Network Toolbox, Mathworks, Natick, MA.
- Mosteller, F., and Tukey, J. W. (1977) Data Analysis and Regression: A Second Course in Statistics, Addison-Wesley, Reading, MA.
- 50. Powell, M. J. D. (1977) Restart procedures for the conjugate gradient method. *Math. Programming*, **12**, 241–254.
- 51. Misiti, M., Misiti, Y., Oppenheim, G., and Poggi, J. (1997) Wavelet toolbox for use with Matlab. Mathworks: Natick, MA.
- 52. Lee, T. W. (1998) Independent Component Analysis: Theory and Applications, Kluwer, Boston, MA.
- Nicolelis, M. A. L., Obeid, I., Morizio, J., Wolf, P. (2000) Soc. Neurosci. Abs. 26(1), 1228.
- 54. Pierce, J. R. (1980) An Introduction to Information Theory: Symbols, Signals and Noise, Dover, New York.
- 55. Siegel, S., and Castellan, N. J. (1988) Nonparametric Statistics for the Behavioral Sciences, McGraw Hill, New York.