ELSEVIER

Electroencephalography and clinical Neurophysiology 102 (1997) 228–239

The Utah Intracortical Electrode Array: a recording structure for potential brain-computer interfaces

Edwin M. Maynard, Craig T. Nordhausen, Richard A. Normann*

John Moran Laboratories in Applied Vision and Neural Sciences, University of Utah, Salt Lake City, UT 84112, USA

Received 18 October 1995; revised version received 8 September 1996; accepted for publication: 12 September 1996

Abstract

We investigated the potential of the Utah Intracortical Electrode Array (UIEA) to provide signals for a brain-computer interface (BCI). The UIEA records from small populations of neurons which have an average signal-to-noise ratio (SNR) of 6:1. We provide specific examples that show the activities of these populations of neurons contain sufficient information to perform control tasks. Results from a simple stimulus detection task using these signals as inputs confirm that the number of neurons present in a recording is significant in determining task performance. Increasing the number of units in a recording decreases the sensitivity of the response to the stimulus; decreasing the number of units in the recording, however, increases the variability of the response to the stimulus. We conclude that recordings from small populations of neurons, not single units, provide a reliable source of sufficiently stimulus selective signals which should be suitable for a BCI. In addition, the potential for simultaneous and proportional control of a large number of external devices may be realized through the ability of an array of microelectrodes such as the UIEA to record both spatial and temporal patterns of neuronal activation. © 1997 Elsevier Science Ireland Ltd.

Keywords: Microelectrodes; Multi-unit activity; Brain-computer interfaces

1. Introduction

There are many circumstances which can result in an individual's loss of control of skeletal musculature including traumatic injury and degenerative diseases of the central and peripheral nervous system. In extreme cases, the extent of these deficits is such that current assistive technologies which rely on peripheral muscle control cannot operate due to lack of input from the individual. For these people, a brain-computer interface (BCI) provides the potential for a new means of communicating with and controlling external devices. Various investigators have built BCIs which use EEG recordings to provide real-time information about the relative activity of different regions of the brain (Farwell and Donchin, 1988; Pfurtscheller et al., 1992; Wolpaw and McFarland, 1994). While these experiments have demonstrated the potential for EEGs in a BCI, early experiments using recordings taken from a single microelectrode implanted in motor cortex have also shown its potential as a source of control signals (Humphrey et al., 1970; Schmidt, 1980).

To control multiple devices or a single device with a number of degrees of freedom using individual microelectrode recordings, it would be desirable to implant a large number of these recording electrodes into motor cortex (Heetderks and Schwartz, 1995; Humphrey and Hochberg, 1995). The Utah Intracortical Electrode Array (UIEA), a silicon-micromachined structure, permits the simultaneous implantation of a large number of microelectrodes in a small region of cortex (100 electrodes per 16 mm^2 of cortex). However, unlike the microelectrodes in the experiments of Humphrey and Schmidt which generally recorded the activity of a single neuron, an electrode of the UIEA preferentially records from a small population of neurons in close proximity to its tip. Due to this fundamental difference in the nature of the recordings, we set out to determine if recordings from the UIEA might be suitable for potential application in a BCI.

* Corresponding author, Department of Bioengineering, 506 Biopolymers Research Building, Salt Lake City, UT 84112, USA. Tel.: +1 801 5817645; fax: +1 801 5818966.

PII S0921-884X(96)95176-0

EEG 95176

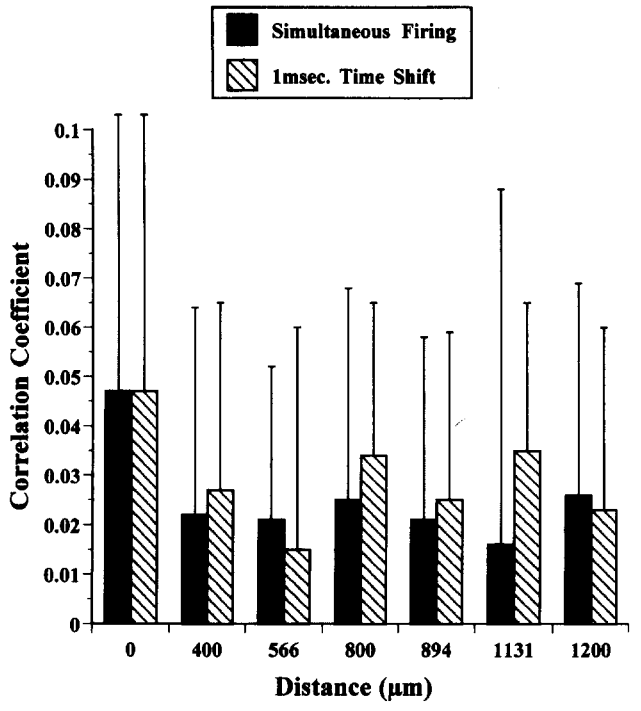

Fig. 9. Cross-correlation coefficients and distance. Cross-correlation coefficient as a function of distance between pairs of electrodes. Solid bars are correlation coefficients relating to simultaneous firing. Line filled bars are the correlation coefficients relating to the expected level of coincident firing using a 1 ms shift predictor. Values calculated for a distance of 0 μm are the autocorrelation coefficients evaluated with a 1 ms time shift.

ing the cross-correlation of two activity trains evaluated at a zero time shift (measuring the level of coincident activity) to the cross-correlation evaluated at a short interval shift (1 ms), it is possible to determine the relative contribution of the common population. If the coincident level is greater than would normally be expected due to the stimulus evoked nature of the responses, then one could conclude that a significant portion of the population of recorded neurons is common to two adjacent electrodes.

Claims by other researchers as to extent of extracellular potential spread from cortical neurons (Nelson and Frank, 1964) confined our interest to a maximum inter-electrode distance of three electrodes (1200 μm). Comparisons were made not only in the rostral-caudal and medial-temporal direction but also along the diagonals. Fig. 9 summarizes the cross-correlation coefficients as a function of the distance between two electrodes; the squares are for cross-correlations evaluated at a 0 ms (within a 1 ms window) time shift, the circles are for cross-correlation coefficients evaluated at a 1 ms time shift. A statistical analysis of these two distributions shows that they are the same ($P < 0.01$). Because the level of coincident activity is not statistically greater than would be expected from two electrodes which had stimulus-evoked activity present on them, we can conclude that there does not exist a significant population of neurons common to adjacent pairs of electrodes.

4. Discussion

4.1. Multiple electrode recording structures

Recent advances in the silicon micromanufacturing processes used to fabricate multi-electrode structures now allow researchers to place a large number of recording electrodes into cortex in a variety of configurations (Kruger, 1983; Jones et al., 1992; Hoogerwerf and Wise, 1994). For any of these structures to be considered as a recording substrate for a BCI, it must be shown to be biocompatible. While this particular topic has not been investigated here, the biocompatibility of silicon-based structures and the UIEA in particular have been addressed elsewhere (Stensaas and Stensaas, 1978; Schmidt et al., 1993). These studies have shown that the extent of the host reaction with silicon structures, in general, is modest and that the response to an implanted UIEA was an isolated region of gliosis (20–40 μm) in the region of the electrode and the formation of a small (<9 μm) fibrous capsule. In addition, there were no indicators of a chronic inflammatory response or edema. Another way to evaluate the biocompatibility of the UIEA is to examine the stability of recordings obtained over a long period of time. Recording studies of chronically implanted electrode arrays in auditory cortex have shown that it is possible to record auditory evoked activity after extended periods of time (Rousche, in preparation). These anatomical and physiological results suggest that the UIEA itself is a biocompatible structure. Problems with the UIEA in the chronic preparations have been attributed to lead failures, tethering forces from the lead wires, or electrode insulation failures rather than damage to the cortical tissue due to the insertion process or presence of the UIEA (Nordhausen and Rousche, personal communication).

4.2. Multi-unit recordings and microelectrode arrays

One of the fundamental features of recordings obtained with microelectrodes is that they generally record from a population of neurons. With single microelectrodes or arrays of microelectrodes with individually positioned electrodes, single unit recordings can be obtained by carefully positioning the electrode close to a neuron. The UIEA and other structures with fixed electrodes do not have this capacity and, thus, require the use of sophisticated algorithms to extract single units from the multi-unit recording. As the SNR of a recording decreases, these algorithms can and do become complex and computationally intensive. While the neuron represents the smallest unit in the nervous system, recent research has shown that it is the ensemble activity of groups of neurons that results in complex behaviors such as movement. This has been shown in a number of experiments conducted in motor cortex where population vectors can accurately predict a number of movement parameters (Humphrey et al.,

unit recordings with the low response variability of EEGs.

To test this, a simple task was designed where neural responses were used to determine the position of a cat's gaze. As described in Methods, a cat's gaze was directed to various positions on a tangent screen. A different computer then correlated the stimulus position with the evoked neural response and attempted to determine the position of the cat's gaze. Fig. 7A shows a representative trace of the activity used to control this system; it is consistent with the type of multi-unit activity generally recorded by the electrodes of the UIEA. Fig. 7B,C shows plots of the cat's gaze (dark trace) and the computers estimate of its gaze (light trace) Experiments were performed for runs where the evoked responses were thresholded at increasing threshold voltages (7C was highest threshold). The performance of the system at the different thresholds was compared using the mean squared error (MSE) between the known position of the eye and the computed position of the eye. As can be seen from these two charts, increasing the threshold voltage had the effect of enabling the system to better track the position of the eye. Increasing the threshold voltage past a certain point, however, does not result in continued improvements in system performance. This is shown in Fig. 8 which plots normalized MSE for runs of the system performed at a number of threshold voltages. As is evident from the graph, at very low and very high threshold values, the system performs less well when compared to intermediate threshold values.

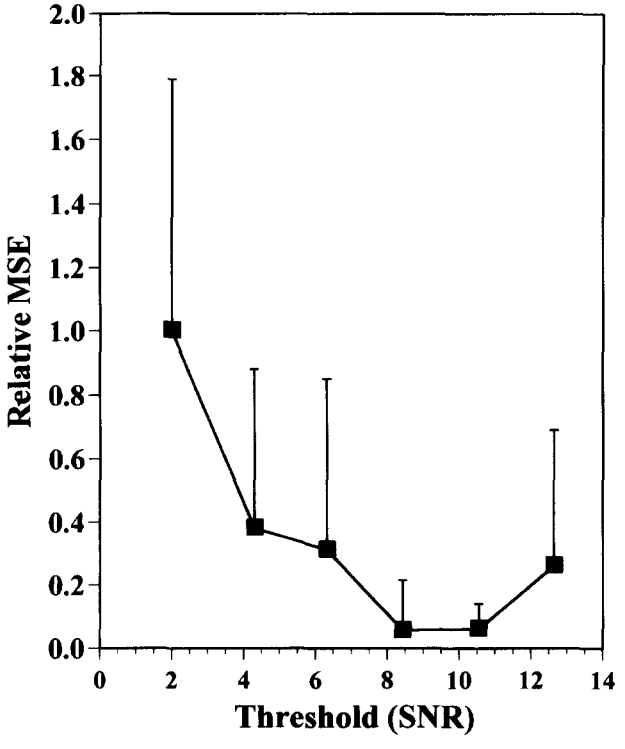

Fig. 8. Summary of tracker performance. Shows the performance (measured in MSE) of the control task for various values of the threshold voltage. Normalization of the MSE was performed to permit comparison between different populations of neurons.

These results demonstrate that the responses recorded by the UIEA, when appropriately thresholded, have the properties of high stimulus specificity and low response variability. In particular, depending on the desired features of the response, the ability to select threshold which results in the desired response characteristics is a significant advantage over other recordings which contain exclusively single unit or EEG-type responses.

3.4. Recording zone independence

Technical challenges associated with the fabrication of high density electrode arrays and the invasiveness of implanting a large number of electrodes into cortex drive the desire to implant an optimal number of electrodes into the target region of cortex. Optimality can be defined as enough electrodes to extract enough information from cortex to accomplish the control goals. Maximizing the amount of information retrieved from the UIEA is not only achieved by changing the size of the constituent neuronal population (through thresholding) but also by eliminating redundant information. Redundancy in this case is manifested as evoked responses from the same neuronal population being recorded by two adjacent electrodes. This is undesirable since it means that the electrodes are not being used in the most efficient manner possible.

There are two means of evaluating whether the electrodes of the UIEA record from independent populations. The first is to use models of extracellular action potential spread in cortical tissues developed by other researchers to predict the range at which action potentials could be detected. These models, with typical values for parameters such as extracellular currents, extracellular resistivity, and RMS noise of the UIEA (4.5 μV) imply that an amplitude detection algorithm such as we have used cannot detect action potentials much beyond 30 μm and 70 μm (Rall, 1962; Humphrey et al., 1978). This would suggest that two electrodes minimally spaced at 400 μm centers, such as is the case with the UIEA, cannot record from populations of neurons which overlap.

The second method of verifying recording independence is to analyze the amount of coincident activity present on adjacent pairs of electrodes as a means of determining electrode independence. The stimulus evoked neural responses on each electrode can be modeled as the sum of two random stochastic processes: one which results from a unique population of neurons around the electrode and another which represents a population of neurons common to two electrodes. To isolate the coincident activity resulting from a common population of neurons is impossible without knowledge of the underlying generating processes, but a method exists for the determination of the existence of a significant common population. By compar-

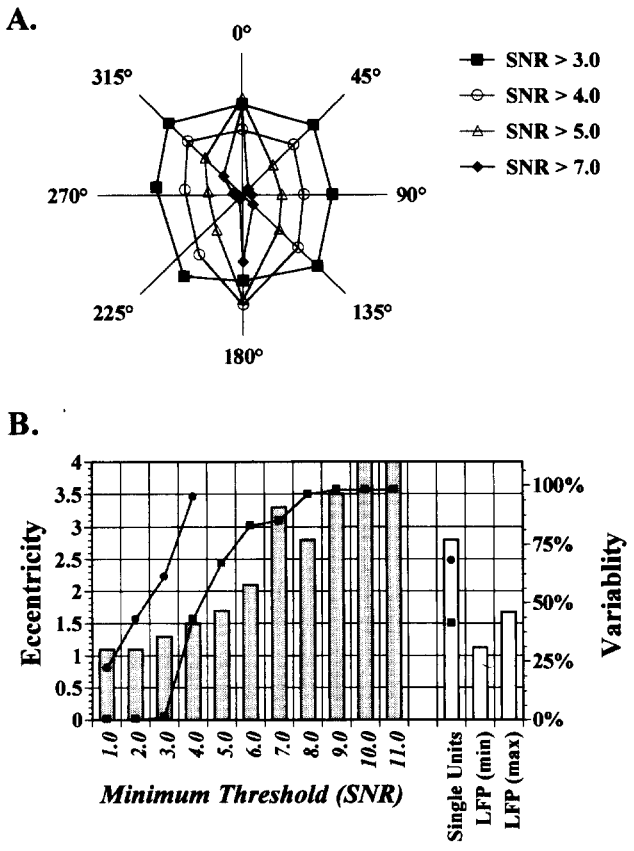

Fig. 6. Thresholding and orientation specificity. (A) Example of the orientation specificity of a single population of neurons as a function of threshold level. (B) Summary of orientation specificity. Shows the effect of threshold (minimum SNR) on eccentricity (filled bars), coefficient of variance (solid circles), and missed passes (solid squares). Eccentricities determined for extracted single units and local field potentials (open bars) are given for comparison.

olded with a very high threshold (SNR > 7.0). In contrast, local field potentials show the same lack of orientation selectivity as responses thresholded at low values (Victor et al., 1994). These results further suggest that single-unit recordings or multi-unit recordings with a high threshold provide the most specific information about the orientation of the light bar stimulus.

Fig. 6B also shows the two measures of response variability: the coefficient of variance (filled circles) and the percentage of passes which did not evoke a detectable response (filled squares). Similar to the results with receptive field size, low thresholded multi-unit activity showed the least amount of variability (C_v = 22%; missed passes = 0%). Results from multi-unit activity thresholded at higher levels (SNR > 4.0; C_v = 95%; missed passes ~43%) were similar to those obtained using single unit recordings (C_v = 68%; missed passes = 41%). These results suggest, as did the results for receptive field size, that multi-unit responses can provide information about the stimulus in a reliable fashion.

3.3. Using thresholded responses as inputs to external devices

These results suggest that low thresholded multi-unit responses (SNR > 2.0) which have similar characteristics to EEG and VEP recordings do not contain very specific information about various features of the response (location and orientation) but their responses manifest low variability, making them easier to reliably detect. High thresholded responses (single units, SNR > 5.0) contain very specific information about the stimulus but are quite variable for sequential presentations of the same stimulus. This would suggest that the reliable signals which a BCI would need might be achieved using multi-unit activity thresholded at an intermediate value. These signals would combine the high stimulus specificity of single

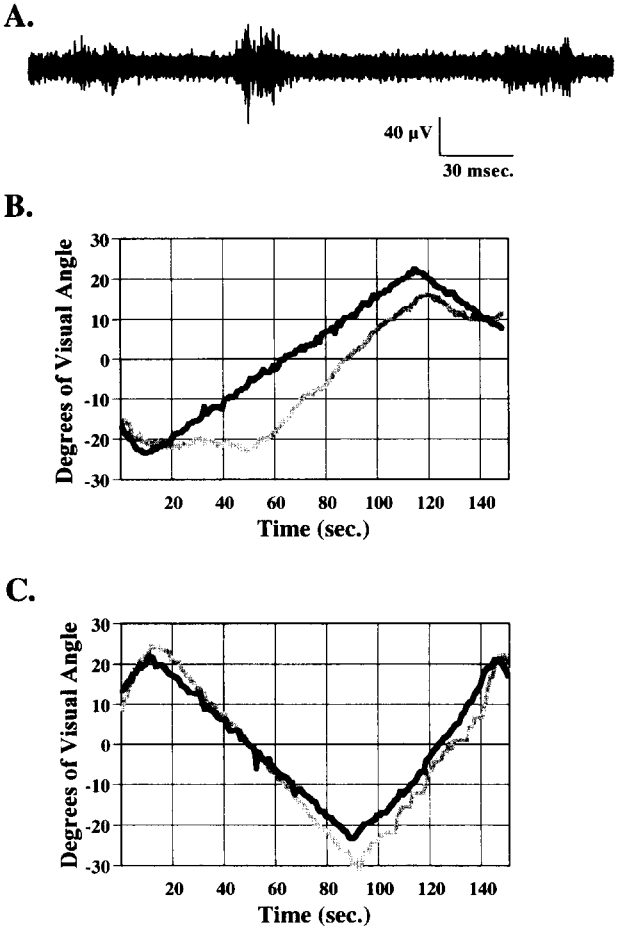

Fig. 7. Tracker performance. (A) Example of neural activity used in the eye tracking task. (B) Performance of task for threshold level set to accept action potentials with SNR>4.0 (MSE = 9.2°). The dark line is the known position of the eye and the light line is the computer's determination of eye position from the evoked responses. (C) Performance of task for threshold level set to accept action potentials with SNR>8.1 (MSE = 4.2°). The dark line is the known position of the eye and the light line is the computer's determination of eye position from the evoked responses.

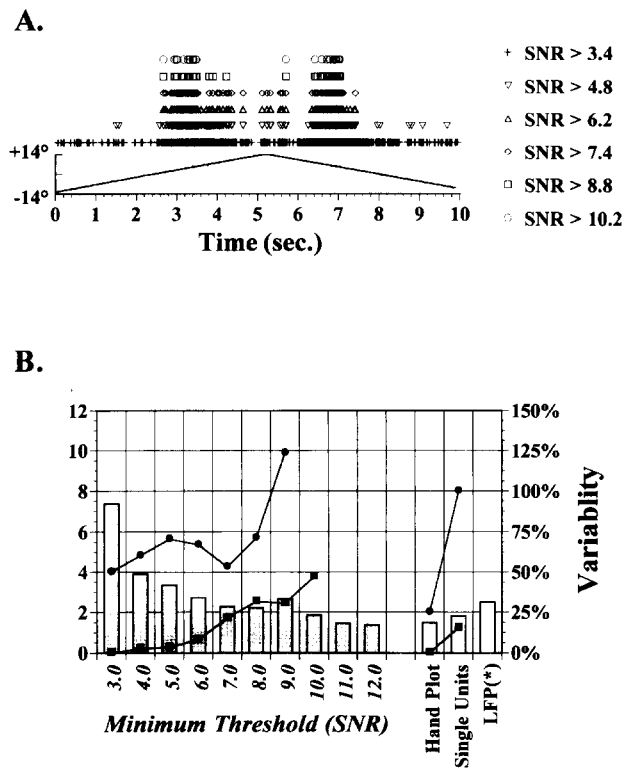

Fig. 5. Thresholding and receptive field size. (A) Traces of a multi-unit evoked response thresholded at a number of increasing threshold voltages. A symbol indicates the detection of a superthreshold event. Stimulus position is shown in lower graph. (B) Summary of receptive field size. The effect of threshold (minimum SNR) on receptive field size (filled bars), coefficient of variance (solid circles), and missed passes (solid squares). Open bars show similar measures for hand plotted receptive fields, single unit recordings, and local field potentials (LFP).

paper). The most reliable response possible would have a coefficient of variance of zero and no missed presentations. Fig. 5A illustrates the effect of threshold level on the specificity of the evoked response to a particular stimulus. The lowest trace shows the position of a light bar stimulus as it is moved across the receptive field of the electrode being recorded from. At low SNR, there is neural activity evoked for all positions of the light making it difficult to determine the position of the light based on the presence of increased neural activity. However, at high thresholds, the presence of a neural response becomes localized to specific positions of light. Using the methods described above, these responses would be processed to compute the mean receptive field size, coefficient of variance (C_v), and percentage of missed presentations.

Fig. 5B is a summary of the results for receptive field sizes computed from four separate recording sessions using 16 electrodes of the UIEA. At the lowest threshold (SNR > 3.0) the mean apparent receptive field size (filled bars) is large (~7.5° of visual angle). Increasing to the highest threshold (SNR > 12.0) results in the receptive field size decreasing to 1.2° of visual angle. Single units extracted from these recordings had receptive field sizes averaging 1.5°. These results can be compared to receptive fields sizes measured for local field potentials (LFP) (Victor et al., 1994). Local field potentials reflect summated activity from a significantly larger populations of neurons than those recorded by the UIEA.

Fig. 5B also shows the coefficient of variance (C_v) (filled circles) and the percentage of passes which did not evoke a detectable response (filled squares). At low thresholds (SNR > 3.0), the coefficient of variance is low (C_v = 50%) and increases with increasing threshold. The coefficient of variance for extracted single units was nearly twice that of the multi-unit responses (C_v = 100%). The percentage of presentations which failed to evoke a measurable response also increased with the threshold. Recordings thresholded at low levels tended to have measurable responses on all stimulus presentations. Recordings thresholded at higher levels manifested a greater number of presentations that failed to evoke a detectable response. The percentage of missed presentations for extracted single units from these records lay in the middle of the range of the multi-unit responses.

These results clearly demonstrate that the more selective recordings (taken with high thresholds) have the smallest receptive field size. However, in most cases, these recordings manifest the highest degree of response variability. In contrast, the non-selective recordings (taken with low thresholds) have the largest receptive field sizes and the lowest coefficient of variance. For the purposes of predicting stimulus location based on evoked activity, this would suggest that very selective recordings, although they best indicate the position of the light, are extremely variable and thus have low reliability.

Fig. 6A shows the effect of thresholding on the orientation tuning curve of a multi-unit recording taken from feline visual cortex. At low thresholds, no orientation preference is manifest in the orientation tuning plot. As the threshold increases, the response becomes significantly more tuned to bars of light at a particular orientation. The eccentricity of this plot (ratio of the highest average firing rate to lowest average firing rate) reflects the selectivity of the response to the given orientation. Tuning curves with an eccentricity close to 1.0 respond similarly to all orientations while higher eccentricities have significantly stronger responses in the preferred orientation.

Fig. 6B summarizes the results for orientation selectivity. Response eccentricity is computed from four recording sessions using 16 electrodes of the UIEA. At the lowest threshold (SNR > 1.0) the mean eccentricity (filled bars) is close to 1 (1.1). Increasing to the highest threshold (SNR > 11.0) results in the eccentricity increasing to 4.0. This result can be compared with the orientation selectivity of single units from these records (open bar). Orientation selectivity measured using these responses corresponds closely in size to multi-unit activity thresh-

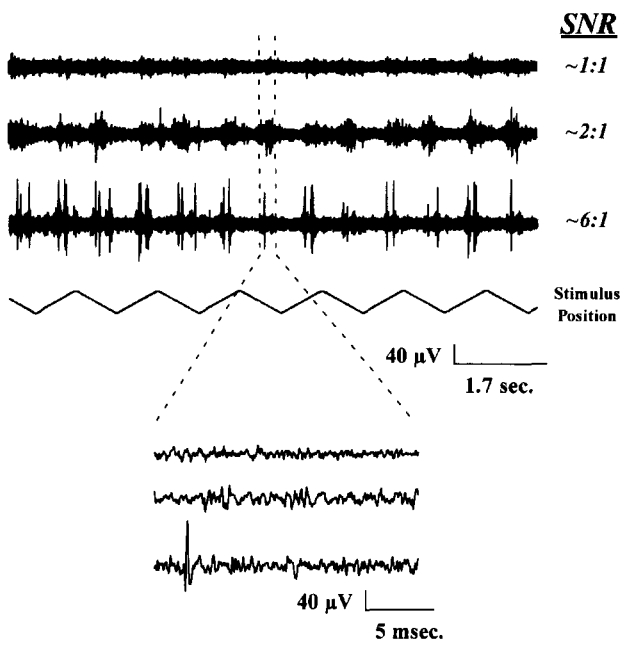

Fig. 3. Typical signals recorded by the UIEA. Traces of recordings taken from a UIEA implanted in cat visual cortex showing signals with three different signal-to-noise ratios (SNR). Stimulus position is shown for a bar of white light moving through a 27.4° displacement.

the energy of the action potential waveform divided by the energy of the noise. A SNR calculated in this manner is generally larger than would result from using peak-to-peak values for the action potential and noise. Beneath these traces, a small region of each trace has been expanded to show the temporal qualities of the recordings. This figure shows that the UIEA generally records from populations of neurons rather than individual neurons and that the SNR and temporal resolution of some recordings permit the extraction and identification of single units. Fig. 4A is a histogram which shows the distribution of action potential SNRs for 238 separable units present during two experiments where a UIEA was implanted in feline visual cortex. Action potentials with a SNR less than 2 are not represented since the amplitude-based threshold detection scheme becomes saturated quickly as the threshold approaches the background. The mean SNR of the action potentials recorded by the UIEA is 6.1:1 with the highest probability of occurrence in the range of SNR greater than 2:1 and less than 6:1. Recordings with SNR above 10, while not as common, can be obtained. The average number of separable neurons for each electrode of the UIEA was 3.4 (Nordhausen et al., 1996). Each point in Fig. 4B represents the number of separable neurons present in recordings from the UIEA that have been processed at different thresholds. The expected number of neurons in a recording at a given threshold level is given by the dashed curve in Fig. 4. Although all of the subsequent results are reported in terms of SNR, this can be used to approximate the number of units present in the recording.

3.2. Receptive field characterization and thresholding

For a neural signal to be considered as an input to a BCI, the BCI must be able to identify a particular pattern of neural activity with a particular stimulus. The ability to do this arises from the specificity of the neural response to a stimulus and its reproducibility. In the context of a BCI, information content can then be interpreted as the ability of the stimulus to be determined from the evoked neural activity. Thus, a response which is evoked only for a particular stimuli would have a high information content; if that particular stimuli always evoked the same response, it would be said to have a high reliability. In a BCI, the most desirable signals are those which are both stimulus specific and highly reproducible.

To quantify the specificity of the response to a particular stimulus, we report the measured receptive field characteristic as a function of the number of units in the response. The reliability of the response is reported as the coefficient of variance of the measured receptive field characteristic and the percentage of stimulus presentations which did not evoke a response ('missed trials' for the remainder of this

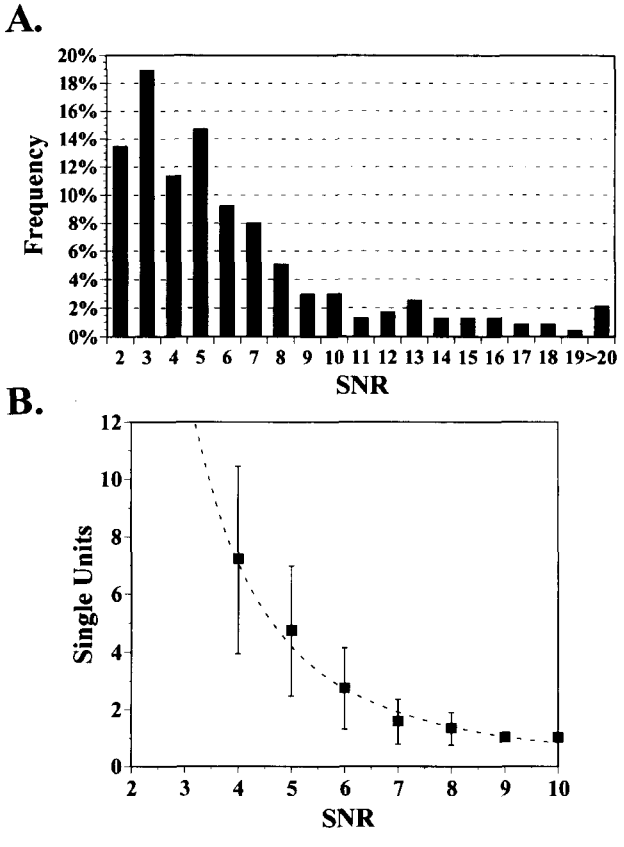

Fig. 4. (A) Single unit SNR distribution. Distribution of the amplitudes of single action potentials. SNR calculated as the peak amplitude divided by the noise level. (B) Plot of the SNR as it relates to the number of separable units in the recording. Each point represents the number of separable units on one electrode of the UIEA. These results predict that at a SNR of 1.0, there would be 183 neurons recorded from.

no more action potentials could be detected. At each threshold level, the mean receptive field size and eccentricity of the tuning curve were calculated. The coefficient of variance (standard deviation divided by the mean) of each of these measurements was also calculated.

2.5. Single unit classification

Because of our interest in comparing the response properties of multi-unit to single unit responses, it was necessary to separate and classify single units in our recordings. To extract single-unit activity from the multi-unit recording, individual action potentials for later classification were first extracted from the records of neural activity using the thresholding procedure described above. These action potential waveforms were peak-aligned and then passed to a template generation program. Templates were formed by a user utilizing either time-amplitude information or the results of an iterative 'principle components analysis' (PCA) process (Schmidt, 1984). An action potential template consisted of a waveform description and a measure of the noise on the electrode for which the template was made. After classifying as many action potentials as possible and creating templates for them, these templates were then used to automatically classify the remaining action potential records. An action potential classification was accepted if its Euclidean distance from the template which it best fit was less than the noise present on that electrode. While certainly not the optimal detection and classification scheme for neural signals, researchers have shown that the difference from optimality is only a couple of dB (Bankman et al., 1993). Once a number of single units had been classified, their receptive fields were characterized in the same manner as multi-unit responses. No attempt was made to extract all possible neurons and only readily separable units were used in the analysis.

2.6. Eye position detection task

To test the ability of multi-unit responses to be used as input in an external process a task was designed whereby a computer would attempt to predict the location of an animal subject's gaze using only stimulus-evoked neural responses. Fig. 2 shows a block diagram of this system. An anesthetized and paralyzed cat has a UIEA (A) implanted in visual cortex and is looking at a tangent screen on which is projected a bright bar of white light (D). Since the eyes are paralyzed and do not move, the cat's gaze can be directed to a known position on the screen using a set of mirrors placed between the cat and the screen (F). A computer (E) controls the angle of one mirror and, thus, can direct the cat's gaze within a 40° arc of visual angle on the tangent screen. The position of the light stimulus is controlled by another computer (C) which is responsible for rapidly and continuously sweeping the light bar across the screen and through the cat's gaze. As

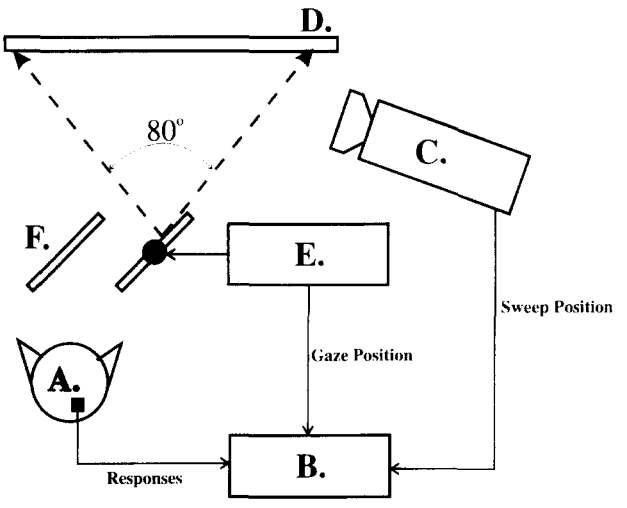

Fig. 2. Schematic of eye tracking task. Schematic of a system which uses visually-evoked responses recorded by the UIEA to determine the position of a cat's gaze. Cat (A) looks at computer-controlled stimulus (C) presented on screen (D). Mirrors (F) control position of cat's gaze. Computer (B) correlates stimulus information and evoked responses to determine position of the gaze.

the light bar moves through the receptive field of the neurons from which the recordings are made, the evoked responses are recorded by the UIEA. After amplification, the computer (B) extracts the action potentials based on the detection threshold and correlates the activity profile with the position of the stimulus from (C). From this correlation and using only the neural responses, the computer (B) attempts to determine where computer (E) has positioned the eye. The mean square error between the known position, given by (E), and the predicted position, given by (B), is computed. The mean square error between these two values is used as the figure of merit.

In the four animals in which receptive field characterizations could be made, the eye position determination task was performed. This task consisted of the computer (E) slowly sweeping the position of the cat's gaze on the tangent screen while the computer (B) attempted to determine the position of the eye based on the evoked responses and stimulus position. For each threshold level, the computers would be allowed to run in excess of 30 min which generally resulted in more than 250 determinations of the gaze position. Because the responses seemed to lose vigorousness after many sequential presentations, the stimulus was removed for a 30 min period between runs.

3. Results

3.1. Recording capabilities of the UIEA

For the UIEA to be a successful candidate recording structure for a BCI, it must possess suitable recording characteristics. Typical stimulus-evoked responses from visual cortex with three different SNR are shown in Fig. 3. The SNR in this figure is calculated as the ratio of

sia assessed using the previously described indicators. Breathing rate and minute volume were adjusted to maintain the end tidal CO_2 levels between 25–35 mmHg. The positions of the optic disks for both eyes were located using back projection techniques and the cat was objectively refracted by bringing the retinal blood vessels into focus on a tangent screen using external lenses.

2.3. Amplification and acquisition

The equipment to acquire signals from the UIEA permits the parallel recording of neural activities of a subset of 16 electrodes from the available 100 electrodes. These 16 signals were passed through a 16-channel custom built high-gain differential amplifier (AC gain 50 000, 4 μV peak-to-peak noise, corner frequencies 300 and 9 kHz, 100 dB/dec roll-off). The conditioned signal from each electrode was then sampled in parallel with the other 15 electrodes to 12-bit resolution at 20 000 samples/s using a commercially available data acquisition board (WIN-30D, UEI, Cambridge MA) and stored to hard-disk for off-line processing. Different sets of 16 electrodes were selected to provide recordings from all functioning electrodes.

2.4. Receptive field characterizations

In four feline preparations, two features of the receptive fields of neurons in cat striate cortex, size and orientation specificity, were characterized as a function of the expected number of cells in the recording. Consistent with established methods, the visual stimulus was a thin bar of bright white light (0.33° of visual angle) moved through the receptive fields of the cell population being recorded (Hubel and Wiesel, 1962). The brightness, orientation, length, and velocity of the bar could all be adjusted by computer control. The brightness of the bar was set at a level which produced a vigorous response for each presentation of the stimuli.

Before characterizing the receptive field size, its location was first mapped using a hand-held pantoscope. After the general region of the receptive field was located and the visual stimulus centered on that position, the bar of light would be swept back and forth across the receptive field at an optimal velocity for a minimum of 10 passes. To control for possible adaptive phenomena in the responses, the stimulus was repeated 5 times with a rest period of 5 min between each presentation. This stimulation was repeated for each electrode in the array which had stimulus driven responses. The sweep length of the visual stimuli was set at 20° of visual angle to ensure that it swept the entire receptive field.

The first step in computing the receptive field size of a population of neurons was to extract the action potentials by thresholding the raw stimulus-evoked activity. The action potential waveforms which were saved consisted of the 5 data points before and 15 data points after the threshold crossing. For a sampling rate of 20 000 samples/s, this corresponds to a 1 ms window. In addition, the position of the visual stimuli when the action potential occurred was also recorded and saved. An activity profile was produced by correlating the probability of the occurrence of an action potential with the position of the visual stimuli an activity profile was produced. This activity profile was further processed using a moving average filter to smooth the response.

$$f(x) = Ae^{\frac{(x-x')^2}{B}} \tag{1}$$

$$w = 2\sqrt{B \cdot \ln\frac{1}{2}} \tag{2}$$

Eq. 1 was then fitted to the smoothed activity profile using the Simplex search algorithm for fitting equations with multiple variables. Although the assumption of a unimodal Gaussian distribution for the response profile was generally adequate, a closeness-of-fit criteria was used to exclude runs where the fitting algorithm failed to converge. Eq. 1 has three parameters: A describes the peak activity, B describes the extent of the activity in terms of the position of the stimulus, and x' locates the center of the receptive field with respect to stimulus position. From this fit, the size of the receptive field was calculated as shown in Eq. 2. This equation uses the half-amplitude point of the response as delineating the edges of the receptive field and most closely resembles the process used for hand-plotting visual receptive fields. This process was repeated for each presentation of the stimulus in both directions and produced at least 20 measurements of the receptive field size.

Once the receptive field was located and its size determined, the orientation of the stimulus was changed in 45° increments and at each orientation 10 passes of the stimulus were made in both the left-to-right and right-to-left direction (20 presentations at each orientation). Five series of presentations were made at each orientation separated by 5 min intervals. In computing the orientation tuning of a multi-unit response, the magnitude of the response was correlated to the orientation of the stimulus. The magnitude of the response for a given stimulus orientation was taken to be the number of action potentials which occurred in the receptive field normalized by the receptive field size. This results in an average measure of the activity while the stimulus was in the receptive field. This response magnitude was then plotted in polar fashion and the eccentricity of the orientation tuning plot calculated as the ratio of the long to short axis of the polar response plot.

To examine the effect of the size of the cell population on the characterizations performed above, the original data sets collected were thresholded at a number of successively increasing threshold levels starting with a threshold level of twice the noise level and stopping at a level where

The suitability of neural recordings for use in a BCI is a function of the ability to extract the desired neural phenomena and interpret its meaning. EEG-based BCIs typically use a component of EEG recordings called the mu wave. These systems control external devices (i.e. cursors on a video screen) by detecting changes in the power of the mu wave in the recordings; these changes are interpreted as the user's desire to perform some action. A microelectrode-based BCI would detect individual action potentials in its recording and determine whether changes in the neural activity were indicative of the user's desire to perform an action. Whereas changes in the mu wave result from the efforts of the user, proper attribution of changes in neural activity patterns to a stimulus (i.e. user's desire) in microelectrode recordings is confounded by the inherent variability present in trains of action potentials (Dean, 1981).

In this paper, we compare the ability of small populations of neurons recorded by electrodes of the UIEA to reliably convey information from single neurons and large populations of neurons. Visual receptive field properties of neural populations in the striate cortex of cats are characterized as a function of the number of neurons recorded in the population. Further, a simple task is described which uses the responses of these neuronal populations to detect the position of a visual stimulus. Based on the results of these experiments and the ability of the UIEA to simultaneously sample neural activity from a large number of spatially distinct neuronal populations, we believe that further investigation in the use of the UIEA as a recording substrate for a BCI is warranted.

2. Methods

2.1. The Utah Intracortical Electrode Array

The micromanufacturing processes used to build the UIEAs in these experiments have been described in detail elsewhere (Jones et al., 1992; Nordhausen et al., 1994). However, a brief physical description of the array is given below. The UIEA is built from a single 0.2 mm thick silicon substrate that measures 4.2 mm on a side. One-hundred electrically isolated needles arranged in a 10 by 10 grid project out from this substrate (Fig. 1). Each needle is 1.2 mm long and is chemically etched to produce a sharp tip. Layering platinum, titanium-tungsten and platinum provides an electrical interface between the electrode and the surrounding tissues. The electrodes are coated with polyimide and the tips exposed to provide the desired electrode impedance (80–150 kΩ). For these experiments, the length of the exposed surface of each electrode was fixed at approximately 50 μm. Finally, gold contact pads on the back of the array provide electrical access to each electrode.

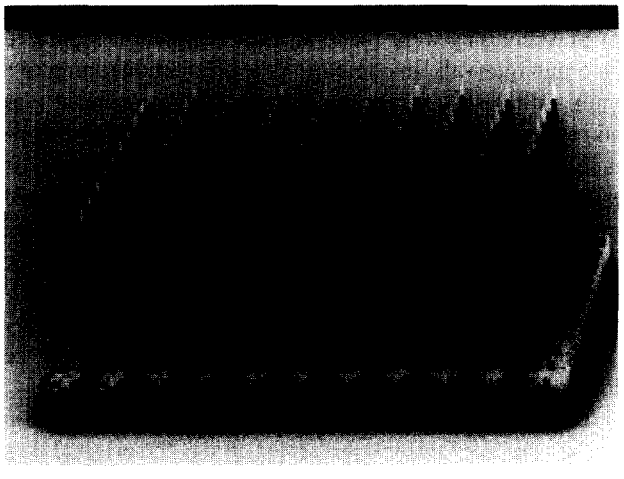

Fig. 1. SEM of the Utah Intracortical Electrode Array. Structure is 4.2 mm on a side and the electrodes are 1.2 mm in length.

2.2. Animal preparation and anesthesia

To minimize acute cerebral edema, the cat was subcutaneously injected with 1cc of dexamethasone (4 mg/ml, American Regent Laboratories) one day prior to the experiment and intravenously the morning of the experiment. At this time, a 1 cc injection of atropine (1 mg/ml, Elkins-Sinn) was administered to assist in the maintenance of a clear air passage. Anesthesia was induced with a 1 cc intramuscular injection of ketamine (10 mg/kg, Ketalar, Parke-Davis) and maintained with Halothane. Throughout the entire experiment, the status of the animal was monitored using EKG, end tidal CO_2, and eye and ear twitch response to mechanical stimulation.

To prepare the cortex for implantation of the array, a small craniotomy was performed using a rotary burring tool (Dremel). After removal of the bone, the dura was reflected and the cortex exposed. The area 17/18 boundary of visual cortex was identified using various anatomical landmarks rather than by probing with a wire electrode (this would result in unnecessary damage to cortex). After positioning, the UIEA was pneumatically inserted into cortex as described by Rousche (Rousche and Normann, 1992). A Pt/Ir reference electrode (2 mm exposed length) was placed in an adjacent gyrus and the cat grounded through a needle placed in the neck muscles.

Final preparation of the animal consisted of covering of exposed cortex with a thin plastic film to prevent cortical dehydration. Zero correction contact lenses were used to prevent dehydration of the corneas. After the stability and depth of anesthesia had been established, typically after 6–8 h, and immediately prior to the initiation of experiments, the cat was paralyzed by IV injection of pancuronium bromide (1 mg/ml, Gensia Laboratories) and the head mechanically immobilized. Paralysis was maintained by periodic injection (~0.2 mg/h) of the paralytic agent. To verify adequate anesthesia, the cat would periodically be allowed to recover from paralysis and the depth of anesthe-

1970; Humphrey, 1972; Georgopolous et al., 1992a; Georgopolous et al., 1992b; Schwartz, 1993). Based on this, one must begin to question the need for expending significant efforts to extract single unit activities from these multi-unit recordings.

The results in this paper show that the judicious selection of the thresholding voltage used to extract superthreshold events from a recording of multi-unit neural activity can result in response features (stimulus selectivity and variability) which are consistent with single unit derived results; without the need for complex extraction and classification algorithms (Hubel and Wiesel, 1962; Duysens et al., 1982; Jones and Palmer, 1987; Lohmann and Reitboeck, 1988). These results also highlight one significant problem with single unit activity: its inherent variability. The design constraints for a BCI require that the signals used contain specific information and convey it in a reliable fashion, this argues against the use of single unit activity records. We have shown, through an analysis of receptive field characteristics and a detection task, that suitably thresholded multi-unit activity can provide the requisite qualities of information content and reliability.

The independence of the population of neurons recorded by each electrode of the UIEA is also important in maximizing the amount of information extracted from the patterns of activity. This independence is a function not only of the propagation of action potential waveforms through cortex but also the size and spacing of cortical efferent columns in motor cortex. Researchers have determined the size of these cortical efferent pathways (comprising the corticospinal tract) to have diameters between 300 μm and 1.0 mm (Asanuma, 1975; Murray and Coulter, 1981) and a separation spacing between 500 μm and 1.0 mm (Jones and Wise, 1977). Minimum feature sizes, the size of a region where all penetrations evoke an identical motor activity for a given level of injected current, were determined to be 500 μm to 2.0 mm (Kwan et al., 1978). Using these anatomically and functionally derived values, proper electrode spacing would place one electrode in each cortical efferent column: this would suggest an electrode spacing on the order of 500 μm. Placing the electrodes closer would not necessarily increase the number of independent control channels since it has been shown that the neurons present in one of these columns are all related to the performance of a particular task.

4.3. Validity of sensory corticies as models for motor cortex

The results presented in this paper were derived from a UIEA implanted in feline visual cortex. To apply these results to the development of a BCI, it is necessary to validate that these results would be consistent with a UIEA implanted in motor cortex. This assumption can be justified from the standpoint of the anatomy of motor cortex and the recordings themselves.

Functionally, the visual and motor corticies perform widely different roles. Visual cortex is a sensory structure and thus is primarily concerned with inputs to the brain. Motor cortex, on the other hand is primarily an output structure. This fundamental difference has its effects on two anatomical features of these corticies: the types and densities of cells and the lamellar structure of the neocortex. Visual cortex is much more densely packed than motor cortex, and by some estimates contains nearly 5 times the cell density as motor cortex in primates (Cragg, 1967; Beaulieu and Colonnier, 1989). The distribution of cell types in feline visual and motor cortex are identical (Winfield et al., 1980). For these two reasons, we would expect recordings from motor cortex to have fewer separable cells but these units would have larger amplitudes. Because visual cortex receives most of its input from the thalamus, it has a larger layer IV than motor cortex. The efferent nature of motor cortex results in a larger layer V structure where a majority of the pyramidal cells reside.

Neither of these two features has the effect of changing the results of this study. A larger layer V can be compensated in the UIEA by increasing the length of the individual electrodes without changing the recording characteristics. Likewise, these electrodes will still record multi-unit activity which will have all of the same characteristics as responses from visual cortex. The use of a sensory cortex for this study was chosen to greatly facilitate the experimental methods. Primarily, the advantage of a sensory cortex is that it is easy to correlate the neural activity with a particular stimulus and it avoids the necessity of complex behavioral training paradigms.

Another reason we can expect multi-unit recordings from visual cortex to be similar to those from motor cortex is that multi-unit activity represents a type of temporal and spatial integration. In cortex, neurons near to each other have significant relationships in their response to a given stimulus. Since each neuron present in a multi-unit response processes the stimulus in a subtly different manner; the combined response of a large number of neurons is functionally independent of slight differences in stimulus characteristics. This was seen as a reduction in the response specificity to bars of a particular orientation with decreasing threshold and a decrease in the response variability. As more neurons were recruited into the response by lowering the threshold, more integration of the stimulus was performed. Multi-unit activity in motor cortex can be expected to have the same characteristics, except that the stimulus is no longer visual or auditory but the volitional intent to perform a motor movement. This integration could prove to be greatly beneficial in that the user will not be required to develop extremely stereotyped strategies for generating specific patterns of activity relating to a command. Further, this integration may mitigate the ef-fects of plastic changes in the functional organization of motor cortex which may result from use of a BCI.

The development of a BCI based on recordings from implanted microelectrode arrays will involve significant research in a number of areas. While this paper has demonstrated that the UIEA might provide a suitable platform for implementation in a BCI, it has not addressed the process of associating a pattern of activity with a desired command. We have shown that the signals recorded by the UIEA can be used to perform a simple task, to extrapolate this result into a functional assistive technology will require a better understanding of the temporal and spatial organization of information processing in motor cortex.

Although there is currently substantial interest in the development of EEG-based BCI systems for individuals with severe motor impairments, we believe that the future of this technology lies with arrays of electrodes which penetrate into motor cortex, such as the UIEA. We look forward to additional studies that will explore the two-dimensional space and time representations of primary motor cortex. These studies will provide further design specifications for future generations of multielectrode arrays which will be even better suited for use as a BCI. In spite of the invasiveness of the UIEA compared to an EEG lead, it is a biocompatible structure and it is capable of acquiring a great deal more information about volitional motor intent than an EEG. This increased amount of information should result in the ability of individuals to control multiple external devices with great precision.

Acknowledgements

The authors would like to thank Vinh Ngo and Pat Rousche for their invaluable assistance. Financial support for this work was provided by the NSF (Grant BCS 9110767) and donations from John A. Moran and the W.M. Keck Foundation.

References

Asanuma, H. Recent developments in the study of the columnar arrangement of neurons within the motor cortex. Physiol. Rev., 1975, 55 (2): 143–155.

Bankman, I.N., Johnson, K.O. and Schneider, W. Optimal detection, classification, and superposition resolution in neural waveform recordings. IEEE Trans. Biomed. Eng., 1993, 40 (8): 836–841.

Beaulieu, C. and Colonnier, M. Number of neurons in individual laminae of areas 3(beta), 4(gamma), and 6a(alpha) of the cat cerebral cortex a comparison with major visual areas. J. Comp. Neurol., 1989, 279: 228–234.

Cragg, B.G. The density of synapses and neurons in the motor and visual areas of the cerebral cortex. J. Anat. 1967, 101 (4): 639–654.

Dean, A.F. The variability of discharge of simple cells in the cat striate cortex. Exp. Brain Res., 1981, 44: 437–440.

Duysens, J., Orban, G.A. and Verbeke, O. Velocity sensitivity mechanisms in cat visual cortex. Exp. Brain Res., 1982, 5: 285–294.

Farwell, L.A. and Donchin, E. Talking off the top of your head: toward a mental prosthesis utilizing event-related brain potentials. Electroenceph. clin. Neurophysiol., 1988, 70: 510–523.

Georgopolous, A.P., Ashe, J., Smyrnis, N. et al. The motor cortex and the coding of force. Science, 1992a, 256: 1692–1695.

Georgopolous, A.P., Schwartz, A.B. and Kettner, R.E. Neuronal population coding of movement direction. Science, 1992b, 233: 1416–1419.

Heetderks, W.J. and Schwartz, A.B. Command-control signals from the neural activity of motor cortical cells: joy-stick control. In: A. Langton (Ed.), Proc. RESNA '95 Annu. Conf., Vol. 15, 1995, pp. 664–666.

Hoogerwerf, A.C. and Wise, K.D. A three-dimensional microelectrode array for chronic neural recording. IEEE Trans. Biomed. Eng., 1994, 41 (12): 1136–1146.

Hubel, D.H. and Wiesel, T.N. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. J. Physiol., 1962, 160: 106–154.

Humphrey, D.R. Relating motor cortex spike trains to measures of motor perfromance. Brain Res., 1972, 40 (1): 7–18.

Humphrey, D.R. and Hochberg, L.R. Intracortical recording of brain activity for control of limb prostheses. In: A. Langton (Ed.), Proc. RESNA '95 Annu. Conf., Vol. 15, 1995, pp. 650–658.

Humphrey, D.R., Corrie, W.S. and Rietz, R. Properties of the pyramidal tract neuron system within the precentral wrist and hand areas of primate motor cortex. J. Physiol. (Paris), 1978, 74 (3): 215–226.

Humphrey, D.R., Schmidt, E.M. and Thompson, W.D. Predicting measures of motor performance from multiple cortical spike trains. Science, 1970, 170 (959): 758–762.

Jones, E.G. and Wise, S.P. Size, laminar and columnar distribution of efferent cells in the sensory-motor cortex of monkeys. J. Comp. Neurol., 1977, 175: 391–438.

Jones, J.P. and Palmer, L.A. The two-dimensional spatial structure of simple receptive fields in cat striate cortex. J. Neurophysiol., 1987, 58 (6): 1187–1211.

Jones, K.E., Campbell, P.K. and Normann, R.A. A glass/silicon composite intracortical electrode array. Ann. Biomed. Eng., 1992, 20: 423–437.

Kruger, J. Simultaneous individual recordings from many cerebral neurons: techniques and results. Rev. Physiol. Biochem. Pharmacol., 1983, 98: 177–233.

Kwan, H.C., MacKay, W.A., Murphy, J.T. et al. Spatial organization of precentral cortex in awake primates. II. Motor outputs. J. Neurophysiol., 1978, 41 (5): 1120–1131.

Lohmann, H. and Reitboeck, H.J. Visual receptive fields of local intracortical potentials. J. Neurosci. Methods, 1988, 25: 29–44.

Murray, E.A. and Coulter, J.D. Organization of corticospinal neurons in the monkey. J. Comp. Neurol., 1981, 195: 339–365.

Nelson, P.G. and Frank, K. Extracellular potential fields of single spinal motoneurons. J. Neurophysiol., 1964, 27: 913–927.

Nordhausen, C.T., Maynard, E.M. and Normann, R.A. Single unit recording capabilities of a 100 microelectrode array. Brain Res., 1996, 726: 129–140.

Nordhausen, C.T., Rousche, P.J. and Normann, R.A. Optimizing recording capabilities of the Utah Intracortical Electrode Array. Brain Res., 1994, 637 (1-2): 27–36.

Pfurtscheller, G., Flotzinger, D., Mohl, W. et al. Prediction of the side of hand movements from single-trial multi-channel EEG data using neural networks. Electroenceph. clin. Neurophysiol., 1992, 82(4): 313–315.

Rall, W. Electrophysiology of a dendritic neuron model. Biophys. J. 1962, 2: 145–167.

Rousche, P.J. and Normann, R.A. A method for pneumatically inserting an array of penetrating electrodes into cortical tissue. Ann. Biomed. Eng., 1992, 20: 413–422.

Schmidt, E.M. Single neuron recording from motor cortex as a possible source of signals for control of external devices. Ann. Biomed. Eng., 1980, 8 (4-6): 339–349.

Schmidt, E.M. Computer seperation of multi-unit neuroelectric data: a review. J. Neurosci. Methods, 1984, 12: 95–111.

Schmidt, S., Horch, K. and Normann, R.A. Biocompatibility of silicon-

based electrode arrays implanted in feline cortical tissue. J. Biomed. Mater. Res., 1993, 27: 1393–1399.

Schwartz, A.B. Motor cortical activity during drawing movements: population representation during sinusoid tracing. J. Neurophysiol., 1993, 70 (1): 28–36.

Stensaas, S.S. and Stensaas, L.J. Histopathological evaluation of materials implanted in the cerebral cortex. Acta Neuropathol. (Berl.). 1978, 41 (2): 145–155.

Victor, J.D., Purpura, K., Katz, E. et al. Population encoding of Spatial frequency, orientationm and color in Macaque V1. J. Neurophysiol., 1994, 75 (5): 2151–2166.

Winfield, D.A., Gatter, K.C. and Powell, T.P.S. An electron microscopic study of the types and proportions of neurons in the cortex of the motor and visual areas of the cat and rat. Brain, 1980, 103: 245–258.

Wolpaw, J.R. and McFarland, D.J. Multichannel EEG-based brain-computer communication. Electroenceph. clin. Neurophysiol., 1994, 90 (6): 444–449.