INFORMATION THEORETIC BOUNDS ON NEURAL PROSTHESIS EFFECTIVENESS: THE IMPORTANCE OF SPIKE SORTING

Ilan N. Goodman and Don H. Johnson

Department of Electrical and Computer Engineering Rice University, Houston, TX 77251-1892

ABSTRACT

We compute the capacity of neural prostheses using a vector Poisson process model for the neural population channel. For single-electrode stimulation prostheses, the capacity is proportional to the size of the population being stimulated, the same value that results when each neuron is stimulated individually. In contrast, when gross recordings are used in control prostheses, the capacity is much less than it is when each neuron's output is treated separately. Consequently, spike sorting, whereby gross recordings are sorted into their constituent spike trains, is crucial to the performance of neural control devices. By computing the capacity of the neural population channel with spike sorting, we find that false positives cause a far greater reduction in capacity than either missed spikes or mislabeled spikes. Thus, a good spike sorting algorithm for neural prostheses should be biased against committing false positives, even at the expense of altering the spike train statistics.

Index Terms— Neural prosthetics, information capacity, spike sorting, neural populations.

1. INTRODUCTION

Advanced techniques for electrode implantation, neural recording, and signal processing have enabled a suite of new applications for neural prostheses and brain-machine interfaces. These devices generally fall into two categories. *Stimulation devices* consist of an electrode (or electrodes) exciting a population of neurons. This category includes sensory prostheses such as cochlear implants and bionic eyes, and implantable stimulators used to treat a variety of conditions such as Parkinson's and clinical depression. *Control devices* rely on extracellular recordings from multiple neurons to control a mechanical device, and include motor prostheses such as bionic limbs, and brain-computer interfaces that have restored communication to some patients with neurodegenerative diseases [1].

Currently, no general consensus has emerged as to the "best" approach to designing neural prostheses, with researchers devising a variety of novel techniques with varying degrees of success. However, though many existing prostheses may perform qualitatively well, evaluating their performance objectively requires knowledge of the ultimate performance benchmark. In other words, we need to know how well the operational system works compared to the optimal ideal, given constraints such as current stimulation limits, number of electrodes, and spike-sorting reliability. From an information theoretic viewpoint, how well neural prostheses perform depends on, in the stimulation case, how effectively information from the stimulator can be conveyed through a gross excitation of a population, and in the control case, how accurately the interface can extract information from multi-unit recordings. In both cases, this capability is determined by the system's information capacity.

Previously, we derived the capacity for a population of neurons modeled as a vector Poisson process, and applied those results to study neural prostheses of both types [2, 3]. That analysis showed that when using only a single input to stimulate an entire neural population (for example, with a single implanted electrode), the best achievable performance is theoretically the same as that resulting from stimulating each neuron separately. In contrast, for neural control devices, having access to the neurons' individual outputs is crucial. If the control device has access only to the superposition of multiple neural responses, the capacity is greatly reduced. Consequently, separating gross recordings into their constituent spike trains, a process commonly known as spike sorting, is an important first step to achieving good performance. Sorting errors reduce capacity from the ideal, though different types of errors -mislabeling, missing spikes, and inserting spikes —may have different effects on the system. Here, we extend our previous results to show how sorting errors alter the capacity of the neural control channel.

2. THE CAPACITY OF NEURAL PROSTHESES

We model a neuron (or population of neurons) as a point process communication channel. The channel produces a sequence of events N_t that encodes an input signal X_t according to an intensity $\mu(t; \mathcal{H}_t)$, which represents how the instantaneous event rate depends on the input X_t and on the process's history \mathcal{H}_t . Assuming a stationary input, the capacity of this

Email: {igoodman,dhj}@rice.edu



Fig. 1. Several configurations of parallel channels for vector and common input cases are shown.

channel is the maximum time-averaged mutual information between X_t and N_t . The maximization is with respect to all input probability distributions that yield a class of intensities reflecting the channel characteristics. Here we consider only constraints on the maximal intensity, though other constraints could be incorporated as well [2].

Kabanov [4] derived the capacity of the single point process channel when the intensity is constrained as $0 \le \mu(t; \mathcal{H}_t) \le \lambda_{\max}$:

$$C^{(1)} = \frac{\lambda_{\max}}{e \ln 2}$$
 bits/s.

We use the notation $C^{(1)}$ to denote the capacity of a single neuron. Kabanov's derivation showed that the capacity of *any* point process satisfying the same intensity constraint cannot exceed the Poisson process's capacity.

We extended this result to find the capacity of a population of neurons. The simplest channel to analyze is the one shown in Figure 1a. In this model, each neuron in a population of size M has both independent inputs and independent outputs. Simple manipulations using the properties of mutual information show that the total capacity of this population equals the sum of the individual capacities. Thus, in a homogeneous population, $C^{(M)} = MC^{(1)}$. When the population has *connection-induced dependencies* that model lateral connections among neurons, the capacity actually increases by a quantity that is a function of the amount of dependence present in the population [2].

Figure 1b models single-electrode neural stimulation, wherein all M neurons respond to a common stimulation signal. Surprisingly, when the neurons are conditionally independent (no lateral connections), the capacity of this channel is identical to the independent input case. Thus, if the proper stimulating signal could be found, single-electrode stimulation works as well as individual stimulation.

From a neural control perspective, these two cases represent the ideal in which the control device has direct access to the spike trains produced by each neuron in the population. In practice, however, it is usually infeasible or impossible to record directly from individual neurons [5]. Figure 1c depicts a more realistic scenario, in which the aggregated activity of an entire population of neurons is recorded as a single output signal, $Y = \sum_{m} Y_{m}$. This situation models single electrode extracellular recordings, as well as EEGs and local field potentials arising from the concurrent activity of large numbers of neurons. Calculations show that when the neurons are conditionally independent, $C^{(M)} \leq 1.577C^{(1)}$, regardless of population size. Consequently, not separating an aggregate recording into its constituents drastically reduces the information that can be gleaned.

A common technique to improve the effectiveness of neural control devices is to record from multiple sites simultaneously, using a multielectrode array, for example. The corresponding channel is depicted in Figure 1d. Here, L recordings are obtained, each from a different subpopulation of the M total neurons. For simplicity, we consider aggregate outputs from L equal-sized, overlapping subpopulations. For large populations with no connection-induced dependencies, $C^{(M)} \leq (2L-1)1.577C^{(1)}$, indicating that multiple aggregated recordings can increase capacity. The upper bound is achieved when the subpopulations overlap substantially. This result breaks down when the factor (2L-1)1.577 approaches M.

3. SPIKE SORTING

Although using multi-site recording can mitigate the negative effect of response aggregation on the ultimate performance of control prostheses, even in the optimal case the resulting capacity is far less than that obtained when each neuron's response can be recorded separately. An alternative (or complementary) technique is to separate response signals from gross recordings: spike sorting. A variety of spike sorting algorithms have been devised using techniques such as template matching, principal components analysis and various clustering algorithms, and each method suffers different detection and classification error rates [6]. To determine the effect spike sorting errors have on the optimal performance of a neural control device, we investigate how different kinds of errors in spike sorting affect the system's capacity.

Unfortunately, finding the probability distribution for a



Fig. 2. Spike sorting channels with different types of sorting errors are shown in the top panel. Labels represent the neural response within a single time bin; for example, 0, 0 indicates neither neuron produced a spike. The bottom panel shows the corresponding ratio of the capacity with spike sorting to the optimal capacity for both the common input (dashed) and independent input (solid) cases. Capacity ratios are shown for $\rho = 0$ (blue), $\rho = 0.4$ (green), and $\rho = 0.8$ (red).

vector Poisson process with spike sorting errors is extremely difficult, which prevents us from computing the capacity directly. Instead, we divide time into spike sorting analysis bins and consider the analogous discrete-time process, the Bernoulli process, for both the population and the spike sorting channel. At each bin index n, the Bernoulli process Y(n)equals either zero or one, independent of any other bin. For the Bernoulli channel, the probability that the channel output equals one at bin n is governed by the stochastic input X(n):

$$P[Y(n) \mid X(n)] = \begin{cases} X(n), & Y(n) = 1, \\ 1 - X(n), & Y(n) = 0. \end{cases}$$

The input is constrained as $0 \le X(n) \le x_{\max}$, reflecting the intensity constraint for the point process channel. We further assume that the input is stationary and consists of statistically independent values from bin to bin.

To extend this model to a population channel, we write the vector Bernoulli distribution using the Sarmanov-Lancaster expansion [7]. For simplicity, we consider only the case M = 2 here. The conditional joint distribution is

$$\begin{split} P(Y_1, Y_2|X) = & P(Y_1|X) P(Y_2|X) \\ & \cdot \left[1 + \rho \frac{(Y_1 - \mathsf{E}[Y_1|X]) \left(Y_2 - \mathsf{E}[Y_2|X]\right)}{\sqrt{\sigma_{Y_1|X}^2 \sigma_{Y_2|X}^2}} \right] \end{split}$$

where ρ is the simple correlation coefficient between the pair of Bernoulli random variables Y_1 and Y_2 , which models the connection-induced dependence between neurons. This construction generalizes easily to M variables using higher order correlation coefficients.

To consider the effect of spike sorting on the overall capacity of the system, we pass the vector Bernoulli process modeling the population output through a *spike sorting channel*. Several options exist for studying the errors introduced by the spike sorting channel, a few of which are depicted in Figure 2 for the case M = 2. Since the output of the cascade is also a vector Bernoulli process, we can find the capacity of the combined channel using the same techniques we used for finding population capacities. Though the discussion here is limited to the M = 2 case, we did investigate the effects for larger populations, and the results are qualitatively similar.

In Figure 2a, the spike sorter commits only errors in labeling single spikes; that is, with probability α , a spike from neuron 1 is mistakenly said to have come from neuron 2, and vice versa. The lower panel in Figure 2a shows the ratio of the channel's capacity to the capacity of the population channel when each response signal is completely known. When the neural population has a single common input, labeling errors have no effect and the capacity ratio is always 1. In contrast, when the population is separately stimulated, the capacity ratio decreases to a worst-case value of around 0.6.

Connection-induced dependence has almost no effect on the capacity in this case.

In the second spike sorting channel, depicted in Figure 2b, simultaneously occurring spikes are incorrectly labeled as being a single event produced by either neuron with probability β , and both spikes are missed entirely with probability $1-2\beta$. The correct two-spike event is never detected in this model. In the common input case, the capacity reduction is linearly proportional to the error rate and the correlation ρ between neurons. In the separate input case, β has almost no effect on the capacity.

Finally, in the channel depicted in Figure 2c, all spikes are correctly detected and sorted, but false positives (extra spikes) are reported and assigned to each neuron with probability γ . Here, the behavior for the common input and separate input case are similar. In both cases, false positives substantially reduce the channel capacity, dropping it to just 10% of its original value in the worst case. Moreover, the capacity ratio drops sharply with γ , approaching its lowest value when $\gamma = 0.2$. Interestingly, connection-induced dependence between neurons actually mitigates the effect of false positives; for example, when $\rho = 0.5$, the worst case capacity ratio increases to around 0.4.

4. CONCLUSION

The capacity defines the ultimate fidelity limits of information transmission by any system [2]. Computing the capacity of the simple population structures revealed that neural stimulation devices are not fundamentally constrained by the stimulation technique; electrically stimulating an entire population of neurons with a single electrode can, theoretically, achieve the same performance as an array of electrodes that stimulates each individual neuron separately. In contrast, we found that neural control devices can suffer severe performance degradation when only gross recordings are used, and multi-electrode recording can only mitigate this effect to a limited extent. In other words, spike sorting is essential to the successful use of neural control devices.

Our capacity analysis of spike sorting channels revealed that the success of spike sorting for neural control depends heavily on the type and frequency of errors that are committed. Most crucial to this performance are false positives, which severely degrade the capacity, even when the error rate is small. Missing spikes also degrades the capacity, but the effect is only linearly proportional to the error rate, and the reduction is only severe when the connection-induced dependence is high. Finally, mislabeling spikes only decreases capacity in the independent input case, whereas in the common input case mislabeling causes no loss in capacity. Though we analyzed the effects of different sorting errors separately here, the channels in Figure 2 can be combined. For example, a recent paper [8] reported average spike sorting error rates of 23% false positive and 30% false negative. In that case, the capacity can be less than 20% of its optimal value —not significantly better than that obtained by a single unsorted gross recording.

It is important to note that maximizing capacity for the spike sorting channel does not necessarily preserve the statistics of the original spike trains. This non-intuitive result suggests that different strategies should be employed for spike sorting during experimentation versus spike sorting for prostheses. While accurate spike train reconstruction is important for experimental analysis, for neural control applications spike sorting should be intentionally biased in order to reduce the rate of false positives, thereby increasing the capacity and improving the device's best-case performance.

5. REFERENCES

- M. Nicolelis, "Actions from thoughts," *Nature*, vol. 409, pp. 403–407, 2001.
- [2] D. H. Johnson and I. N. Goodman, "Inferring the capacity of the vector Poisson channel with a Bernoulli model," *To appear in Network: Computation in Neural Systems*, 2007.
- [3] I. N. Goodman and D. H. Johnson, "Information theoretic bounds on the effectiveness of neural prosthetics," *BMC Neuroscience*, vol. 8 (Suppl 2), p. S8, 2007.
- [4] Y. Kabanov, "The capacity of a channel of the Poisson type," *Theory Prob. and Applications*, vol. 23, pp. 143– 147, 1978.
- [5] A. B. Schwartz, X. T. Cui, D. J. Weber, and D. W. Moran, "Brain-controlled interfaces: Movement restoration with neural prosthetics," *Neuron*, vol. 52, pp. 205–220, October 2006.
- [6] M. S. Lewicki, "A review of methods for spike sorting: The detection and classification of neural action potentials," *Network: Computation in Neural Systems*, vol. 9, pp. R53–R78, 1998.
- [7] I. N. Goodman and D. H. Johnson, "Orthogonal decompositions of multivariate statistical dependence measures," *Proc. 2004 International Conference of Acoustics, Speech, and Signal Processing*, May 2004.
- [8] F. Wood, M. J. Black, C. Vargas-Irwin, M. Fellows, and J. P. Donoghue, "On the variability of spike sorting," *IEEE Transactions on Biomedical Engineering*, vol. 51, no. 6, pp. 912–918, June 2004.