

Neural prosthetic control signals from plan activity

Krishna V. Shenoy,^{1,2} Daniella Meeker,³ Shiyao Cao,⁴ Sohaib A. Kureshi,^{1,5,6} Bijan Pesaran,^{1,7}
Christopher A. Buneo,¹ Aaron P. Batista,^{3,8} Partha P. Mitra,⁹ Joel W. Burdick⁴ and
Richard A. Andersen^{1,3,CA}

¹Division of Biology; ³Computation and Neural Systems Program; ⁴Division of Engineering and Applied Science; ⁷Division of Physics, Mathematics and Astronomy, California Institute of Technology, Pasadena, CA 91125; ⁵Division of Neurosurgery, Duke University, Durham, NC 27707;

⁹Bell Labs, Lucent Technologies, Murray Hill, NJ 07974, USA

²Present address: Department of Electrical Engineering and Neurosciences Program, Stanford University, Stanford, CA 94305, USA

⁶Present address: Neurosurgical Medical Clinic, Inc., San Diego, CA 92103, USA

⁸Present address: Howard Hughes Medical Institute and Department of Neurobiology, Stanford University, Stanford, CA 94305, USA

^{CA}Corresponding Author: andersen@vis.caltech.edu

DOI: 10.1097/01.wnr.0000063250.41814.39

The prospect of assisting disabled patients by translating neural activity from the brain into control signals for prosthetic devices, has flourished in recent years. Current systems rely on neural activity present during natural arm movements. We propose here that neural activity present before or even without natural arm movements can provide an important, and potentially advantageous, source of control signals. To demonstrate how control signals can

be derived from such plan activity we performed a computational study with neural activity previously recorded from the posterior parietal cortex of rhesus monkeys planning arm movements. We employed maximum likelihood decoders to estimate movement direction and to drive finite state machines governing when to move. Performance exceeded 90% with as few as 40 neurons. *NeuroReport* 12:591–596 © 2003 Lippincott Williams & Wilkins.

Key words: Bayesian decoders; Brain–computer interfaces; Finite state machines; Maximum likelihood; Neural prosthetic systems; Posterior parietal cortex

INTRODUCTION

The parietal reach region (PRR) of the posterior parietal cortex (PPC) is located at an early stage in the sensory-motor pathway. It is closely related to sensory areas, particularly visual areas, and projects to limb movement areas within the frontal lobe [1,2]. Many properties of PRR make it an attractive source of plan activity to derive control signals for prosthetic systems [3,4]. First, PRR plan activity is selective for arm movements, as opposed to eye movements, and persists until a reach is initiated [5]. The persistence of activity during planning does not require an actual movement; in essence this area codes the ‘thoughts’ to move. Second, PRR plan activity is abstract, being represented in visual (eye-centered) coordinates and the activity within the spatial representation shifts with each eye movement to remain spatially invariant [6]. Moreover, cells in this area also carry eye position information in the form of a modulation of the eye-centered response fields and thus the goals of movements can be read out in other coordinate frames as well [7]. Finally, during sequential reaching to two memorized locations, PRR plan activity codes just the next intended reach [8]. This simplifies the interpretation of activity in this region for prosthetic control since plan activity reflects the upcoming movement, not any or all planned movements. These properties suggest that intended

movement activity from PRR may be well suited for generating high-level, cognitive control signals for prosthetic applications.

We report here the results of a computational investigation, using a database of PRR action-potential responses, to explore how high-level, cognitive control signals can be estimated from plan activity using maximum likelihood decoders and finite state machines. The resulting control signals are well suited for use with external devices such as robotic limbs or computer interfaces [9–13].

MATERIALS AND METHODS

Single neuron recordings: PRR spike data were obtained from a previous study from our laboratory [6], and surgical and recording techniques for acquiring single-neuron action potentials have been described previously [5,6]. All protocols were approved by the Caltech Institutional Animal Care and Use Committee.

Data analysis: We used maximum likelihood estimation, which is equivalent to Bayesian estimation with a uniform prior probability distribution, to estimate reach parameters. Our assumptions were Poisson spike statistics and statistical independence between cells, but explicit models of tuning to

the various parameters were not assumed [14]. To reconstruct the planned reach direction, we defined the scalar $x = (1, 2, \dots, 8)$ to be the reach direction and the vector $\mathbf{n} = (n_1, n_2, \dots, n_N)$ to be the spike count from each neuron (n_i) during a time interval (τ). Combining the expression for the conditional probability for the number of spikes \mathbf{n} to occur, given a plan to reach direction x , with Bayes' rule yields the following expression for the conditional probability of x given \mathbf{n} : $P(x|\mathbf{n}) = C(\tau, \mathbf{n})P(x)(\prod_{i=1}^N f_i(x)^{n_i} \exp(-\tau \sum_{i=1}^N f_i(x)))$. The normalization factor $C(\tau, \mathbf{n})$ ensures that the sum of the probabilities equals one. $P(x)$ is the prior probability for reaches in each direction, and is uniform by experimental design, and the mean firing rate of the i^{th} neuron while planning a reach to direction x is $f_i(x)$. The estimated reach direction, \hat{x} , was taken to be the one with the highest probability: $\hat{x} = \operatorname{argmax}_{x \in \{1, 2, \dots, 8\}} (P(x|\mathbf{n}))$.

Action potentials from 23 PRR neurons from monkey CKY, and 41 PRR neurons from monkey DNT, were analyzed. All analyses yielded similar results for both animals. We used cross-validation techniques to assess the performance of this estimation process. For each repetition of the simulation, and in each of the eight possible reach directions, a random subset of the total number of cells was selected to avoid a cell sampling bias. One trial was selected randomly, from each of the selected cells, and set aside for use as test data. With the remaining trials from the selected cells, we calculated the average firing rates for each cell while planning to reach to each target. This mean was used as the rate parameter λ in Poisson distributions. The probability that a particular selection of test data belonged to each of the multidimensional distributions from each direction was assessed, and thus the most probable (i.e., decoded or predicted) reach direction was selected for each repetition in the given direction. This process was repeated 1000 times in each of the eight reach directions and then normalized.

A similar procedure was used to estimate the response distributions for the time-course analyses, but with the following variations. After selection of the random subset of cells and the exclusion of a single random trial from each cell, the remaining trials were divided into 3 epochs: baseline, plan period, and pre-movement period (-600 to 0 , 300 to 1000 , and 1100 to 1350 ms, respectively, where 0 ms is the onset of the reach target and reaches began directly after the pre-movement period ends). The trials from each direction, for each cell, and in each epoch were concatenated, and the data were sampled with 250 ms long moving windows with 50 ms time steps. The baseline epoch was concatenated across all directions. Additionally the plan epoch was also sampled using 500 ms windows rather than 250 ms windows. The mean of each epoch was used as the parameter for the single multidimensional Poisson distribution for the baseline period, and for each of the eight multidimensional distributions for each direction in the three other epochs (the 250 ms sampled memory epoch, the 500 ms sampled memory epoch and the pre-execution period).

Test-data firing rates were measured in 250 ms windows, advanced 50 ms at each time step, through the duration of the test trial. The most probable condition (baseline, one of eight plan directions, or one of eight execution directions) was estimated independently in each time step as above.

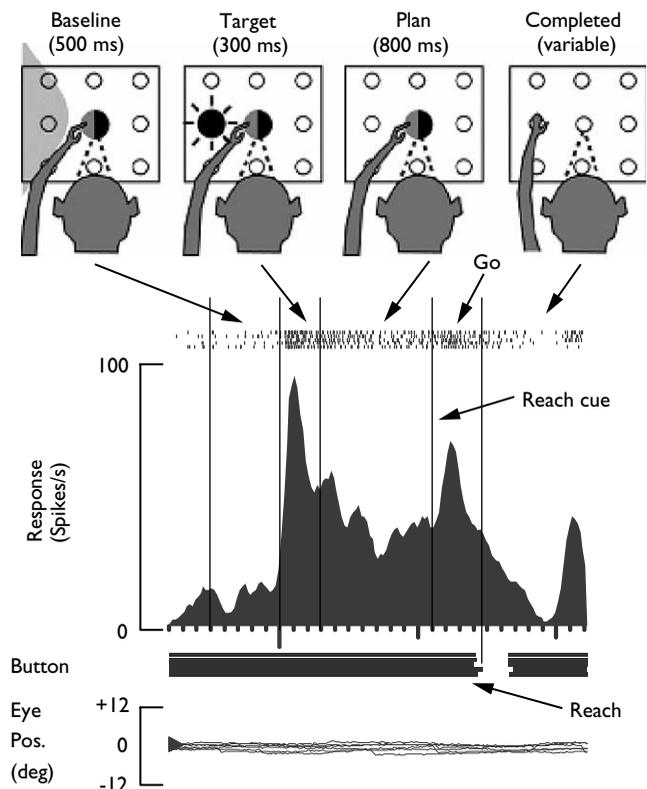


Fig. 1. Parietal reach region (PRR) neural activity during the delayed, center-out reaching task. The behavioral task consists of four stimulus/behavior periods (baseline, target presentation, plan and go). Icons depict eye (black semicircle and dashed lines) and hand positions (gray semicircles and arm), potential target locations (open circles), the neuron's region of maximum response or response field (large shaded region), and the location of the flashed target specifying the reach goal (black circle with emanating lines). The vertical line labeled reach cue indicates when the central eye and hand LEDs are extinguished. Spike times are indicated as vertical lines in the trial-by-trial rasters (five rows at top of data) and the peri-stimulus time histogram (PSTH) represents the average response. Horizontal bars indicate when push buttons were depressed, with reach onset in a particular trial corresponding to when the bar vanishes. Eye position traces are shown at the bottom. Data are aligned to target onset (vertical line separating baseline and target presentation periods), which is when the reach goal first becomes known. This neuron from monkey DNT preferred downward reaches; the response field is illustrated to the left in the icon for illustrative convenience.

RESULTS

Data were analyzed from a previous study in which action potentials, eye movements and push-button state were recorded while two rhesus monkeys performed a delayed center-out (eight push-button targets) reaching task [6]. Figure 1 plots the response of a PRR neuron during repeated reaches to the memorized location of a flashed visual target. Three periods of neural activity are of particular interest: a baseline period preceding target presentation, a plan period following target presentation but preceding the reach cue, and a pre-movement or go period following the reach cue but preceding the onset of the arm movement. Plan and go period activity levels vary with the location of the flashed visual target, which specifies the goal of the arm movement.

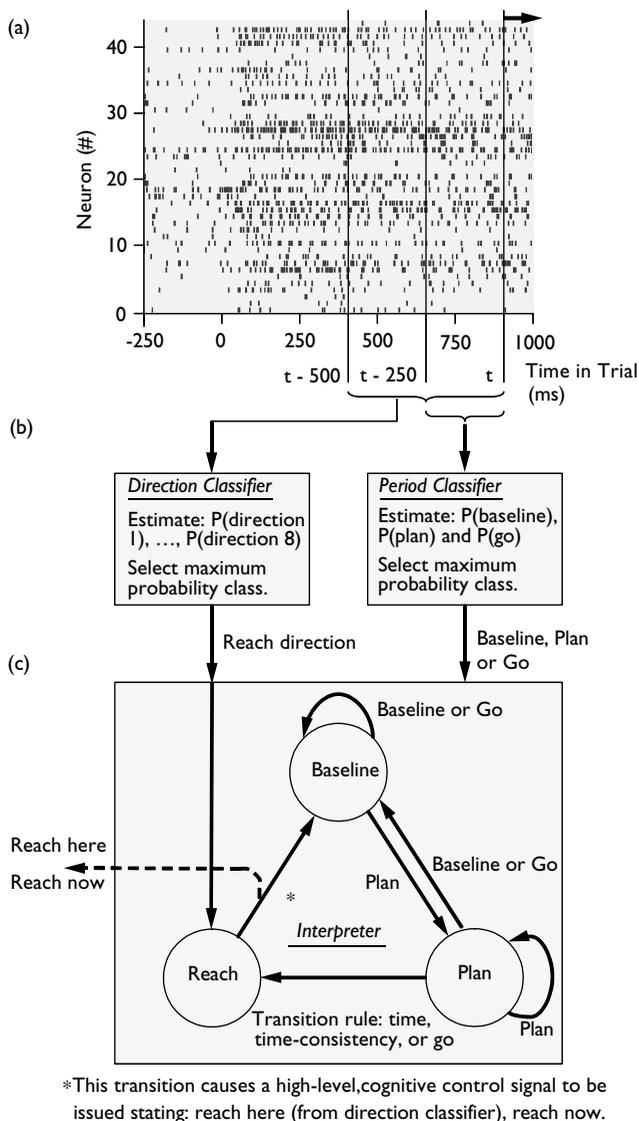


Fig. 2. Computational architecture for generating control signals from PRR plan activity. (a) Spike rasters from one trial for each of 41 PRR neurons. The goal visual target occurs at 0 ms. The onset of arm movement occurs after 1100 ms (not shown). (b) Classifiers use neural activity from fixed-width sliding analysis windows to estimate the direction of arm movement (direction classifier) and the current neural/behavioral period (period classifier). (c) The interpreter receives the stream of period classifications (i.e., baseline, plan or go) from the period classifier and the stream of movement direction classifications (e.g. downward reach) from the direction classifier. The interpreter consists of a finite state machine that transitions among three states (baseline, plan and reach) according to the period classification at each time step.

Most neurons are tuned for a particular goal direction, with other directions eliciting weaker plan period activity. The estimation and control algorithms for determining from neural measurements (1) when PRR is planning a reach, (2) when the animal intends to execute the planned movement and (3) in which direction the reach is being planned are illustrated in Fig. 2. Figure 2a illustrates the neural response from each neuron in the population throughout a representative (but simulated) delayed cen-

ter-out reaching task. The vertical line labeled *t* represents the current time, which would also indicate the time of the most recent data if the prosthetic system was running in real time. Figure 2b is termed the classifier and has two parts. The direction classifier uses neural data from the past 500 ms to estimate the probability that a reach is being planned to each of the eight directions, and the most probable reach direction is then selected. The period classifier uses neural data from the past 250 ms to estimate the probability that PRR is currently in a baseline, plan or go period (see Fig. 1), and the most probable class is then selected. Figure 2c is termed the interpreter. The interpreter must take in the series of baseline, plan and go classifications, generated by the period classifier as time evolves, and determine when a reach should be executed. It must also take in where the reach should be directed from the direction classifier and finally issue the high-level control signal stating: reach here, reach now.

The interpreter starts in the baseline state and, as shown in Fig. 2c, can transition to the plan state or return to the baseline state each time the period classifier issues another period classification. A baseline or go period classification keeps the interpreter in the baseline state, while a plan period classification advances the interpreter to the plan state. Once in the plan state, a baseline- or go-period classification will return the interpreter to the baseline state. The reason for this operating logic will become clear when we discuss below the possible rules for transitioning the interpreter from the plan state to the reach state. Once the reach state is achieved the interpreter automatically transitions back to the baseline state, and simultaneously issues a high-level, cognitive control signal commanding an immediate reach to the location given by the goal classifier (Fig. 2c, *).

The question of when to transition the interpreter from the plan state to the reach state, and subsequently triggering an arm movement, can be answered by considering the behavioral task instructions and go period classifications. We now summarize the logic and performance of three different transition rules to illustrate how increasingly sophisticated rules improve performance.

Time transition rule: If the behavioral task instruction to the subject is simply to plan a reach to a particular location for half a second, then a prosthetic system can safely execute an arm movement after detecting 500 ms of plan activity. In other words, the interpreter can transition from the plan state to the reach state when the period classifier issues 500 ms of contiguous plan classifications. The time transition rule is essential because the visual-cue onset response is similar to the movement onset response and could, without a rule enforcing a minimum plan duration, result in a premature and erroneous go period classification. Figure 3a shows the percentage of trials achieving the reach state, and thus executing a reach, for a range of population sizes. Figure 3b indicates the percentage of these trials that executed reaches in the correct direction for a range of population sizes. Ideally all trials would execute reaches, as all of our experimental data are from successful reach trials, and all trials would reach in the correct direction. Although this transition rule successfully executes reaches for most

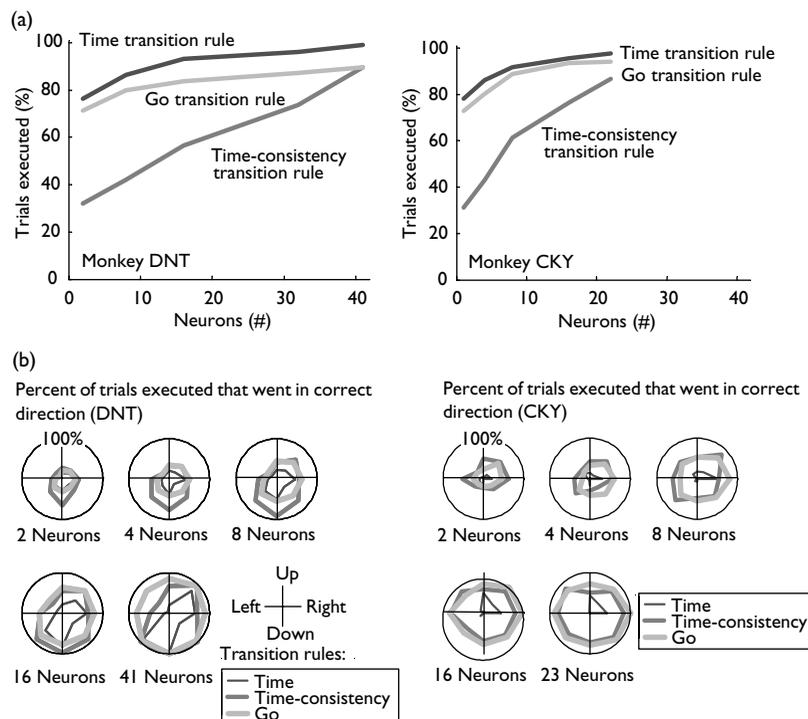


Fig. 3. Interpreter performance characteristics. The interpreter was characterized using the time, time-consistency, and go transition rules. (a) Percentage of trials achieving the interpreter's reach state, thereby triggering a reach, as a function of the number of neurons in the population. Perfect performance (100%) means that all trials executed a reach to some goal location, but not necessarily to the correct goal location. (b) Percentage of trials that executed a reach to some goal location that did reach to the correct goal location. Perfect performance (100%) is plotted as a circle in all sub-panels. Each sub-panel shows performance for a different number of neurons in the analysis population. In both (a) and (b) neurons from animals DNT and CKY were used to generate the performance curves appearing to the left and right, respectively.

trials (Fig. 3a), many of the reaches go in the wrong direction (Fig. 3b). These errors are due to direction classifier misclassifications, and are probably caused by low signal to noise ratios. If errors were caused by drifts in plan or volition then the prediction accuracy would not be expected to increase dramatically by adding more neurons to the estimate, as is seen in Fig. 3b.

Time-consistency transition rule: A simple extension of the prior transition rule can address these concerns by adopting the conservative view that it is better not to execute a reach at all than to reach in the wrong direction. By adding the constraint that the period classifier's plan classifications must also specify a given goal direction throughout the required 500ms plan period (i.e. cannot switch between reach goals during this period) we effectively impose a plan-stability requirement. Importantly, the period classifier, which employs a 250ms sliding window, can also estimate goal location using response models and estimation methods analogous to those in the familiar direction classifier. Figure 3 also summarizes the performance of this transition rule. As expected, fewer trials now execute reaches (Fig. 3a) but those that do tend to reach in the correct direction more often (Fig. 3b).

Go transition rule: While the previous two transition rules perform well for certain applications, and importantly they

do not rely on neural signals associated with movement execution, we would also like to be able to produce a larger absolute number of correct reaches. We can achieve this by replacing the previous stability constraint with a requirement that the period classifier issue a go period classification, after plan period classifications have been issued continuously for 500 ms, in order to transition from the plan state to the reach state. Using a neural go signal could afford the subject an additional opportunity to abort a planned reach by withholding the go command, or the possibility of reducing the length of the plan period on some trials. Figure 3 illustrates the performance. The period of time used by the direction classifier to estimate the reach direction, which is the 500ms directly preceding the go period classification, tends to be slightly later than with the previous two transition rules. This is because the go period classification can occur up to several hundred milliseconds after the plan duration criterion has been met. This accounts for the increased percentage of reaches to the correct location (Fig. 3b). This algorithm executes an intermediate number of reaches, compared with the other two transition rules (Fig. 3a), with good performance arising from the readily detected and classified go activity.

Besides the increase in spike activity, another potential source of movement onset information is the local field potential (LFP) [15–17]. To investigate this possibility we recorded LFPs while a monkey looked and reached toward

eight peripheral visual targets from a central starting position. In this additional (third) rhesus monkey a silicon micro-machined Utah electrode array with 22 active electrodes was implanted permanently within PRR for chronic recording and the LFP signals were filtered (15–25 Hz bandpass). The average power in this band is moderate around the time the central fixation and touch targets are illuminated, builds just before the peripheral targets specifying the saccade and reach goals become visible, and declines rapidly around the time of movement onset. Further examination revealed that power in this band is modulated by both saccadic eye movements and reaching arm movements, but that reaching arm movements tend to modulate the power to a greater extent than do saccadic eye movements, with power being reduced to nearly zero directly after reach onsets. This suggests that the dramatic reduction in low-frequency PRR LFP power at the time of movement onset could also contribute to a neural go-signal transition rule.

DISCUSSION

Inspired by the considerable success of cochlear implants, tremor-control devices and other neural-prosthetic systems aimed at delivering signals to the nervous system, research aimed at reading out neural signals for prosthetics applications has intensified in recent years [3,4,9–13]. Despite these advances, the field of prosthetic systems that interface with the CNS is still in its infancy. It is important to introduce new ideas about decoding movement parameters for possible use as prosthetic control signals since the full range of prosthetic applications is not yet known, and there does not yet exist a neural-prosthetic architecture that is optimal for all plausible prosthetic applications.

To explore the feasibility of using pre-movement neural signals from PRR to generate high-level cognitive control signals, we developed and tested the computational architecture presented in Fig. 2. This part of an envisioned neural prosthetic system estimates, from PRR neural activity, when an arm movement is being planned (period classifier), the direction of the planned movement (direction classifier), and when the arm should move (interpreter). The resulting computations issue a cognitive control signal with two parts: reach here and reach now. Thus PRR contains sufficient signals to operate a neural prosthetic system.

To our knowledge, this is the first application of a state machine model to predict cognitive states using neural activity; neural decoding algorithms typically use estimators that do not explicitly model internal dynamical states or the transitions between these states [12,13]. Importantly, we put forth the architecture shown in Fig. 2 simply to demonstrate the essential features of translating plan activity into control signals. For instance, we have used a relatively long (500 ms) analysis window for the direction classifier since averaging over this period reduces estimation noise and does not cause additional delay. Meanwhile, the period classifier analysis window is relatively short (250 ms) since it is important to track the period (state) reasonably quickly so that the interpreter evolves in a timely fashion. Similarly, the interpreter transition rules investigated here merely represent three simple possibilities. We do not claim that this is the ideal estimation-control architecture or that any of the

parameters (e.g. analysis window widths) are optimal. In fact, we fully expect architectures that propagate the full probabilistic nature of the internal state transitions (e.g., with hidden Markov models), adapt along with neural adaptation or plasticity, and incorporate truly optimal parameter values will only perform better. Our goal here is simply to demonstrate the basic approach needed to produce useful control signals from plan activity.

Possible attributes of PRR for neural-prosthetic control:

One open and central question is whether neural representations that are present during natural arm movements, and are employed in current prosthetic-arm research systems, remain completely intact following injury and disease or, alternatively, suffer at least some degeneration that would complicate their use in prosthetic control [18]. Given that PRR is more closely linked to the visual system, and more distant from motor areas effected by paralysis, it is possible that it remains more intact following paralysis. Shoham and colleagues [19] have reported residual topography in motor cortex related to the will to move in partially paralysed patients. No doubt the most direct method of assessing the integrity of areas will come from cell recordings in paralyzed patients receiving prosthetic implants.

The parietal cortex is believed to participate naturally in ongoing visuomotor coordination and adaptation [20]. Such cortical plasticity could help improve prosthetic system performance by continually, and quickly, adjusting for visual-prosthetic misalignments and by countering neural sampling biases, whether created by less than optimal surgical placement of electrodes or by a representational bias in cortex. A rapid and high degree of plasticity would enable patients to control a variety of devices including robotic devices that are very different from the human body, computers for communication, and autonomous vehicles. Plasticity will also be useful in allowing patients to effortlessly adjust to changes in the recordings that result from the usual small drifts of the electrodes in the brain.

In many forms of paralysis patients also lose somatosensation. Somatosensory and proprioceptive feedback are important for error correction for motor behavior, as is vision. Since vision generally remains after injuries or diseases resulting in paralysis, and PRR is strongly and directly linked to visual cortex [1,21] it is likely that PRR will still receive appropriate error signals for motor learning. We have recently shown that PPC maintains appropriate spatial register between proprioceptive and visual modalities, both for visual-to-limb coordinate transformation and for taking in retinal error signals [22].

The cognitive quality of PRR activity also has possible advantages. The persistence of planning activity, which does not require the execution of a movement, may be easily tapped in paralyzed patients who may still be able to activate this planning area, even though they cannot execute movements. At least in motor cortex this planning-related activity, which precedes movement-related activity, appears not to be as robust as in parietal cortex.

Finally, the use of cognitive signals may reduce the number of neurons required for a given prosthetics application. This reduction is possible if relatively few, and high-level, parameters are estimated from the cognitive

activity and the signal to noise ratios are enhanced by averaging over the movement planning period. That performance reaches ~90% with as few as 40 neurons with this simple, non-optimized classifier-interpretor architecture is already encouraging.

Prosthetics system design using cognitive control signals:

In order to produce complex movements we envision delivering high-level control signals to a reasonably sophisticated prosthetic controller capable of generating arm movement trajectories and capable of using inverse models of the prosthetic or electrically stimulated arm to achieve the desired movement dynamics. While the idea of such an intelligent prosthetic controller might sound daunting at first, industrial robotics combine state of the art machine vision and learning to achieve impressive levels of path planning, grip force control and safety. The patient would have the ability to plan an arm movement to an object, have the controller guide the arm to that location and perhaps automatically grasp the object and, finally, the person could plan a subsequent arm movement to the desired location where the object could be released. Cognitive control signals could potentially cooperate with lower-level motor-cortical signals to further optimize control of arm movement prostheses and cognitive control signals could also potentially contribute to existing communication link systems [9,10] due to the expected versatility of cognitive control signals as discussed above. Though using cognitive control signals may require more sophisticated prosthetic system engineering than is currently employed, these control signals may offer important advantages and reduce overall system complexity.

CONCLUSION

Despite the fact that numerous cortical regions and classes of neural activity contribute during the planning and execution of natural arm movements, neural prosthetic systems research has focused almost exclusively on neural activity present during natural arm movements (perimovement activity). We report here the first computational

demonstration that activity present before, or even without, natural arm movements can be readily decoded to produce useful control signals. Neural activity from the parietal reach region, which represents the goal location of the upcoming arm movement, may prove to be an ideal source of prosthetic control signals by virtue of its relative isolation from the motor periphery and direct involvement in sensory-motor coordination and adaptation.

REFERENCES

1. Johnson PB, Ferraina S, Bianchi L and Caminiti R. *Cerebr Cortex* **6**, 102–119 (1996).
2. Andersen RA, Snyder LH, Bradley DC and Xing J. *Annu Rev Neurosci* **20**, 303–330 (1997).
3. Shenoy KV, Kureshi SA, Meeker D *et al.* *Soc. Neurosci Abstr* **25**, 383 (1999).
4. Meeker D, Shenoy KV, Cao S *et al.* *Soc Neurosci Abstr* **27** (2001).
5. Snyder LH, Batista AP and Andersen RA. *Nature* **386**, 167–170 (1997).
6. Batista AP, Buneo CA, Snyder LH and Andersen RA. *Science* **285**, 257–260 (1999).
7. Cohen YE and Andersen RA. *Neuron* **27**, 647–652 (2002).
8. Batista AP and Andersen RA. *J Neurophysiol* **85**, 539–544 (2001).
9. Wolpaw JR, Birbaumer N, Heetderks WJ *et al.* *IEEE Trans Rehab Eng* **8**, 164–173 (2000).
10. Kennedy PR, Bakay RAE, Moore MM *et al.* *IEEE Trans Rehab Eng* **8**, 198–202 (2000).
11. Wessberg J, Stambaugh CR, Kralik JD *et al.* *Nature* **408**, 361–365 (2000).
12. Seruya MD, Hatsopoulos NG, Paninski L *et al.* *Nature* **416**, 141–142 (2002).
13. Taylor DM, Helms-Tillery SI and Schwartz AB. *Science* **296**, 1829–1832 (2002).
14. Zhang K, Ginzburg I, McNaughton BL and Sejnowski TJ. *J Neurophysiol* **79**, 1017–1044 (1998).
15. Donoghue JP, Sanes JN, Hatsopoulos NG and Gaal G. *J Neurophysiol* **79**, 159–173 (1998).
16. Murthy VN and Fetz EE. *Proc Natl Acad Sci USA* **89**, 5670–5674 (1992).
17. Pesaran B, Pezaris JS, Sahani M *et al.* *Nature Neurosci* **5**, 805–811 (2002).
18. Turner JA, Lee JS, Martinez O *et al.* *IEEE Trans Neural Syst Rehab Eng* **9**, 154–160 (2001).
19. Shoham S, Halgren E, Maynard EM and Normann RA. *Nature* **413**, 793 (2001).
20. Clower DM, Hoffman JM, Votaw JR *et al.* *Nature* **383**, 618–621 (1996).
21. Blatt GJ, Andersen RA and Stoner GR. *J Comp Neurol* **299**, 421–445 (1990).
22. Buneo CA, Jarvis MR, Batista AP and Andersen RA. *Nature* **416**, 632–636 (2002).

Acknowledgements: We thank B.L. Gillikin, J. Wynne, J. Baer and J.S. Pezaris for expert advice and assistance in designing and performing the surgery to implant a chronic electrode array in PRR. We also thank E.M. Maynard, B.W. Hatt, R.A. Normann, N.G. Hatsopoulos and J.P. Donoghue for surgical and electrode-array advice related to this surgery. We thank M. Sahani for writing the real-time behavioral control and data collection software HYDRA2, D.J. Dubowitz for pre-surgical MRI and R.A. Normann for lending us LFP amplifiers. Finally, we thank B.L. Gillikin for veterinary assistance, Cierina Marks for administrative assistance and Viktor Shcherbatyuk for computer assistance. This work was supported in part by NIH, DARPA, ONR, Sloan Center for Theoretical Neurobiology at Caltech, McKnight Foundation, NSF Engineering Research Center at Caltech, James G. Boswell Neuroscience Professorship (R.A.A.), and Burroughs Wellcome Fund Career Award in the Biomedical Sciences (K.V.S.).