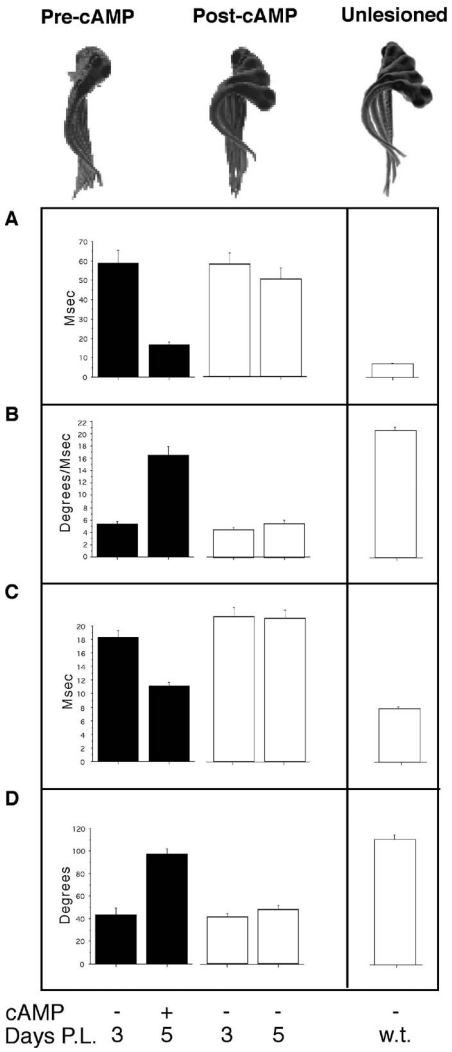

Fig. 5. Recovery of behavior. (Top) The escape bend of a fish before regeneration, after regeneration, and in an unlesioned fish. Images are shown every 2 msec after the start of the turn until the maximum of the bend. Bottom panels quantify the escape performance before and after cAMP-induced regeneration (mean + SEM). Performance measures included (**A**) response latency, (**B**) peak angular velocity, (**C**) duration, and (**D**) maximum angle of the bend. These performance measures are shown for a group of five fish (five trials each) studied before (black bar, 3 days post-lesion) and after (black bar, 5 days post-lesion) cAMP treatment and for a control, untreated group (white bars) over the same time course ($P < .0001$ in every case for treated versus control). White bars on the right show performance measures from wild-type (w.t.) fish at 9 days.

7. J. L. Goldberg, M. P. Klassen, Y. Hua, B. A. Barres, *Science* **296**, 1860 (2002).
8. W. D. Snider, F. Q. Zhou, J. Zhong, A. Markus, *Neuron* **35**, 13 (2002).
9. A. J. Udvadia, R. W. Koster, J. H. Skene, *Development* **128**, 1175 (2001).
10. H. M. Bomze, K. R. Bulsara, B. J. Iskandar, P. Caroni, J. H. Skene, *Nature Neurosci.* **4**, 38 (2001).
11. A. Buffo *et al*., *J. Neurosci.* **17**, 8778 (1997).
12. B. Zheng *et al*., *Neuron* **38**, 213 (2003).
13. T. Takami *et al*., *J. Neurosci.* **22**, 6670 (2002).
14. T. Becker *et al*., *J. Neurosci.* **18**, 5789 (1998).
15. T. Becker, M. F. Wullimann, C. G. Becker, R. R. Bernhardt, M. Schachner, *J. Comp. Neurol.* **377**, 577 (1997).
16. S. J. Zottoli *et al*., *Prog. Brain Res.* **103**, 219 (1994).
17. S. Neumann, F. Bradke, M. Tessier-Lavigne, A. I. Basbaum, *Neuron* **34**, 885 (2002).
18. H. T. Kao *et al*., *Nature Neurosci.* **5**, 431 (2002).
19. J. Qiu *et al*., *Neuron* **34**, 895 (2002).
20. J. R. Fetcho, *J. Neurophysiol.* **67**, 1574 (1992).
21. K. S. Liu, J. R. Fetcho, *Neuron* **23**, 325 (1999).
22. D. A. Ritter, D. H. Bhatt, J. R. Fetcho, *J. Neurosci.* **21**, 8956 (2001).
23. D. M. O'Malley, Y. H. Kao, J. R. Fetcho, *Neuron* **17**, 1145 (1996).
24. K. Haas, W. C. Sin, A. Javaherian, Z. Li, H. T. Cline, *Neuron* **29**, 583 (2001).
25. D. I. Lurie, D. S. Pijak, M. E. Selzer, *J. Comp. Neurol.* **344**, 559 (1994).
26. D. Cai *et al*., *Neuron* **35**, 711 (2002).
27. Y. Gao, E. Nikulina, W. Mellado, M. T. Filbin, *J. Neurosci.* **23**, 11770 (2003).
28. H. Song *et al*., *Science* **281**, 1515 (1998).
29. M. Nishiyama *et al*., *Nature* **424**, 990 (2003).
30. A. J. Shaywitz, M. E. Greenberg, *Annu. Rev. Biochem.* **68**, 821 (1999).
31. O. Steward, B. Zheng, M. Tessier-Lavigne, *J. Comp. Neurol.* **459**, 1 (2003).

32. We thank S. Palma and R. Grady for their care of the animals; S. Kishore, H. Patzelova, and D. McLean for assistance with experiments; D. McLean and P. Brehm for helpful comments on the manuscript; and S. Halegoua for experimental advice. This work was supported by grants from the New York State Spinal Cord Injury Research Trust Fund and the NIH NS 26539.

Supporting Online Material

www.sciencemag.org/cgi/content/full/305/5681/254/DC1
Materials and Methods
Fig. S1
Movie S1

29 March 2004; accepted 8 June 2004

Cognitive Control Signals for Neural Prosthetics

S. Musallam, B. D. Corneil,* B. Greger, H. Scherberger,† R. A. Andersen‡

Recent development of neural prosthetics for assisting paralyzed patients has focused on decoding intended hand trajectories from motor cortical neurons and using this signal to control external devices. In this study, higher level signals related to the goals of movements were decoded from three monkeys and used to position cursors on a computer screen without the animals emitting any behavior. Their performance in this task improved over a period of weeks. Expected value signals related to fluid preference, the expected magnitude, or probability of reward were decoded simultaneously with the intended goal. For neural prosthetic applications, the goal signals can be used to operate computers, robots, and vehicles, whereas the expected value signals can be used to continuously monitor a paralyzed patient's preferences and motivation.

Neural prosthetics are being designed to record brain activity related to intended movements from the sensorimotor pathway of paralyzed patients and to use these signals to control external devices. It would be valuable to determine what parameters can be decoded and used for prosthetic applications. Previous research has concentrated on extracting the online (real-time) intended trajectories of the hand by recording signals primarily, but not exclusively, from the motor cortex (*1*–*5*). This study explores whether a higher level signal of the goal of a movement can be decoded for prosthetic control. For example, a goal signal indicates the intention to reach for an apple, whereas a trajectory signal would indicate the intended direction of the hand movement during the reach. Another high-level signal of interest is expected value, which is used for making decisions. For instance, if an individual has two potential reach goals, an apple and an orange, and the subject prefers apples over oranges, there are signals in his or her brain that indicate this preference and will influence the decision to reach for the apple instead of the orange. We refer to this approach of extracting high-level signals as cognitive based; intended trajectories can also be considered among this group of signals, although at a lower level.

Recordings were made at points along a major pathway for visually guided movement which begins in the extrastriate visual cortex (*6*) and passes through the parietal reach region (PRR) and area 5 to the dorsal premotor cortex (PMd) and then to the primary motor cortex (*7*, *8*). Although PRR is specialized for reaching movements (*9*, *10*), it represents the goals of the reach in visual coordinates (*11*). This visual representation indicates that the planned movement is at an abstract level and codes the intended goal rather than how to move the hand. Further emphasizing its cognitive nature, this goal signal is present when

Division of Biology, California Institute of Technology, Mail Code 216-76, Pasadena, CA 91125, USA.

*Present address: Department of Physiology, Pharmacology, and Psychology, University of Western Ontario, London, ON N6A 5K8, Canada.

†Present address: Institute of Neuroinformatics, University/Eidgenössische Technische Hochschule (ETH) Zurich, 8057 Zurich, Switzerland.

‡To whom correspondence should be addressed. E-mail: andersen@vis.caltech.edu

animals plan but withhold movements in the dark; that is, the planning activity exists apart from any visual or movement-related signals. Area 5 has also been shown to encode movement intention in both visual and limb coordinates (*12*). Although less is known about the reference frame of PMd, it appears that at least a subset of cells have similar properties to PRR (*13*–*15*).

In three monkeys, 64 and 32 electrode arrays were implanted in the medial intraparietal area (MIP) (a component of PRR) and area 5, respectively (*16*). Only cells from MIP were used for decoding in two monkeys (monkeys S and C), whereas a small minority of area 5 neurons were included for monkey O. Monkey S also had 64 electrodes implanted in PMd in a separate surgery.

Experiments were initiated 2 weeks after array implantation and each daily experimental session consisted of 250 to 1100 trials (median number of trials was 819, 726, and 361 for monkeys S, C, and O, respectively). Each session was divided into a reach segment for collecting a database and a brain control segment to decode the position of a cursor on a computer screen. Each session started with the reach segment, during which monkeys performed 30 (or 20 for PMd) memory guided reaches per direction. This task required them to reach to a flashed cue after a delay of 1.2 to 1.8 s (memory period, Fig. 1A). The go signal was the offset of a central green target (*17*). Neural data recorded during successful reaches were added to a database to be used in the brain control segment. The brain control trials began similarly to the reach trails, with a cue flashed in the periphery. However, during the memory period the movement intention was decoded with 900 ms of neural data beginning 200 ms after the cue offset. When the correct goal was decoded, a cursor was placed at the intended reach location and the subjects were rewarded. Trials were aborted if a hand movement occurred during the memory period. If the wrong target was decoded, then monkeys were instructed by the offset of the central green target to reach to the cued location (*17*). Trials in both the reach and brain control segments were aborted whenever any eye movements occurred that placed the gaze outside of a 5° centrally located window.

In different sessions, the database was either not updated after the end of the reach segment (frozen database), or updated after successfully decoded brain control trials (adaptive database). The adaptive and the frozen database yielded similar decode results (fig. S2), indicating that the database update was unnecessary. Notably, the adaptive database eventually contained only brain control trials without leading to loss in performance. This indicates that training sets for cognitive prosthetics can be obtained in paralyzed patients who do not have the ability to reach.

The use of memory period activity ensured that only the monkeys' intentions were being decoded and not signals related to motor or visual events. An example of the absence of motor-related signals during brain control trials can be seen in Fig. 1B. The memory activity of this cell is present for both reach and brain control trials, but the motor burst, which occurs about 1.4 s after the onset of the memory period in reach trials (red), is absent during successfully decoded brain control trials (black).

Three representative performances of online cursor control with intentional signals from three separate sessions from monkey S are shown in Fig. 1, C and D. On the basis of only memory period activity from eight PRR neurons, four targets were correctly decoded with 64.4% accuracy with 250 brain control trials, and six targets were decoded with 63.6% accuracy with 275 brain control trials in separate sessions (Fig. 1C). On the basis of recordings of 16 cells from PMd, eight targets were decoded with 67.5% accuracy with 310 brain control trials (Fig. 1D). Figure 1E describes the overall performance of the three monkeys across all sessions for brain control trials.

The performance can likely be improved by increasing the number of cells. Although many neurons were tuned during the visual and/or motor period of the task, they were not used during brain control trials unless

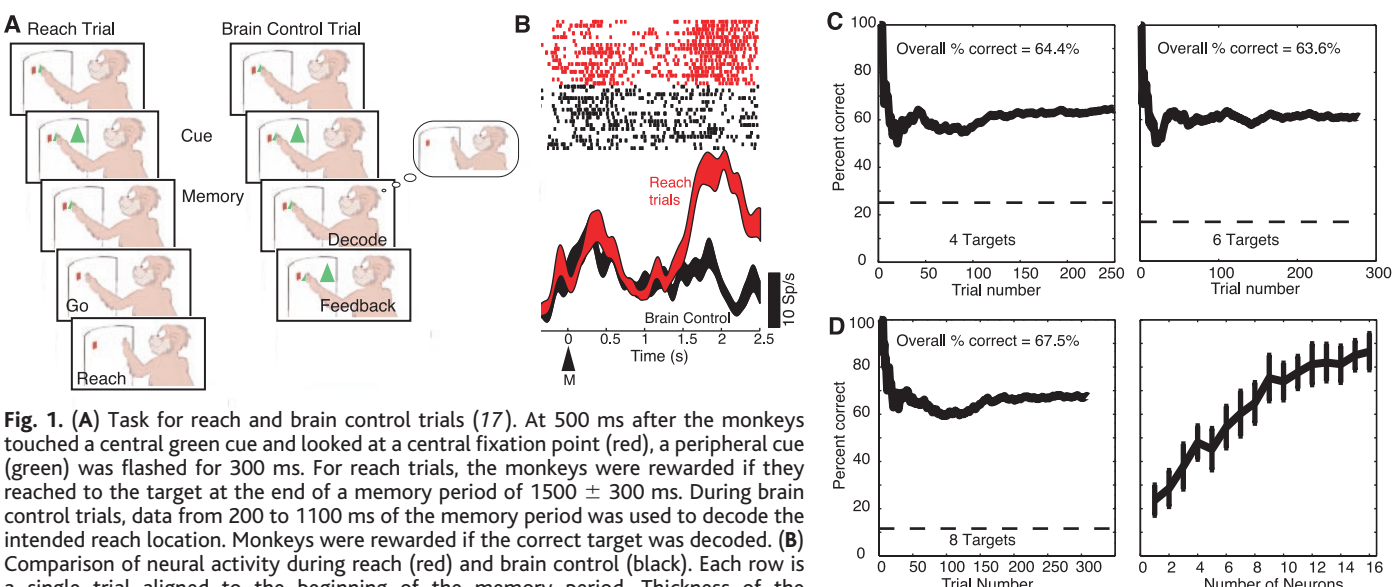

Fig. 1. (**A**) Task for reach and brain control trials (*17*). At 500 ms after the monkeys touched a central green cue and looked at a central fixation point (red), a peripheral cue (green) was flashed for 300 ms. For reach trials, the monkeys were rewarded if they reached to the target at the end of a memory period of 1500 ± 300 ms. During brain control trials, data from 200 to 1100 ms of the memory period was used to decode the intended reach location. Monkeys were rewarded if the correct target was decoded. (**B**) Comparison of neural activity during reach (red) and brain control (black). Each row is a single trial aligned to the beginning of the memory period. Thickness of the poststimulus-time histogram (PSTH) represents the standard error calculated with the bootstrap method. M, start of memory period; Sp, spikes. (**C**) Cumulative decode performance of monkey S during brain control trials for four targets and six targets on the basis of eight neurons from the parietal cortex. Dashed line indicates chance performance. (**D**) (Left) Cumulative performance of a brain control session with 16 neurons recorded from the dorsal premotor cortex of monkey S. (Right) Offline decode with the same data, showing the effect of the number of cells on decode performance. Notably, the number of neurons that can achieve a high success rate remained relatively low. (**E**) Mean success rate across all sessions for three monkeys. Values are the percentage of successfully decoded brain control trials. Number in parentheses is the standard deviation of the distribution of success rates. NS, number of sessions; *, recordings from dorsal premotor cortex; all other recordings are from parietal cortex.

E

	Monkey S	NS	Monkey C	NS	Monkey O	NS
4 Targets	45.0 (10.5)	62	34.2 (5.0)	81	43.2 (17.1)	13
5 Targets	48.1 (7.3)	10	30.6 (2.9)	7	59.3 (0.2)	2
6 Targets	37.1 (11.1)	10	25.6 (5.8)	2	31.2 (14.7)	6
*4 Targets	75.2	1				
*8 Targets	68.2 (1.3)	2				

they showed significant tuning during the memory period as assessed by an analysis of variance (*17*). The right panel of Fig. 1D shows the effect of the number of PMd neurons on the performance of the decode offline. Neurons were randomly chosen without replacement from the pool of 16 used in the brain control trials. The improved performance offline, compared with the left panel of Fig. 1D, is due to the larger training sets used (*17*). Increasing the number of neurons improves performance. However, the total number of neurons still remained relatively low. A similar result was found for PRR neurons with an offline decode of cells that were recorded one at a time (*18*).

Shorter decode intervals than 900 ms can also be used for the online decoding (*17*) (fig. S6). Offline analysis on the subset of sessions that yielded decode rates greater than 60% on 900-ms intervals showed that a 100-ms interval (200 to 300 ms of the memory period) decreased the performance by a mean (±SD) of 14.3 ± 6.1%. Thus, it is likely that increasing the number of cells will result in very fast and accurate online decodes.

Significant learning resulted in improved performance of the brain control task for PRR recordings over a period of weeks. The percentage of trials successfully decoded from the parietal cortex in monkeys S and C for all sessions with four targets (250 to 1100 brain control trials per session) is shown in Fig. 2A. Not enough sessions were available for monkey O and the PMd recordings from monkey S to permit a similar analysis. For monkeys S and C, our ability to decode their intentions was initially poor, hovering just above chance level. However, it continuously improved over the course of a number of weeks. Regression coefficients of the performance as a function of session number for monkey S and monkey C were 0.5 and 0.08 percentage points per session, respectively. Both of these positive regression coefficients are significant ($P < 0.01$ for monkey S and $P < 0.02$ for monkey C).

Over the course of all the sessions, the amount of information carried by neurons in the brain control task increased more than the amount of information during the reach segment of the task. We calculated the mean of the mutual information of the neurons in the memory period during the reach segment and during the brain control segment for each of the first 68 sessions for monkeys S and C (*17*). The mutual information measure quantifies the degree of spatial tuning and can be used as a metric that can describe any change in the degree of tuning. Data from monkey S are shown in Fig. 2B. This analysis yielded two points per session: the mean information during the reach segment while constructing the database (120 reaches, red points, Fig. 2B) and the information during the initial 120 brain control trials immediately following the reach segment (black). The information carried by cells recorded during the same session increased more during the brain control segment than during the reach segment over the course of 68 sessions ($P < 0.01$). A similar result was seen in monkey C (*17*).

Another series of experiments examined whether expected value could also be decoded from PRR activity. We ran a variant of the memory reach task in which cue size indicated the amount, probability, or type of reward that monkeys would receive upon completion of a successful trial. Only one aspect of the reward (amount, probability, or type) was varied in a single session. Cue size was randomly varied trial by trial and the interpretation of cue size was varied across sessions so that a large cue represented the more desirable and less desirable rewards on different days.

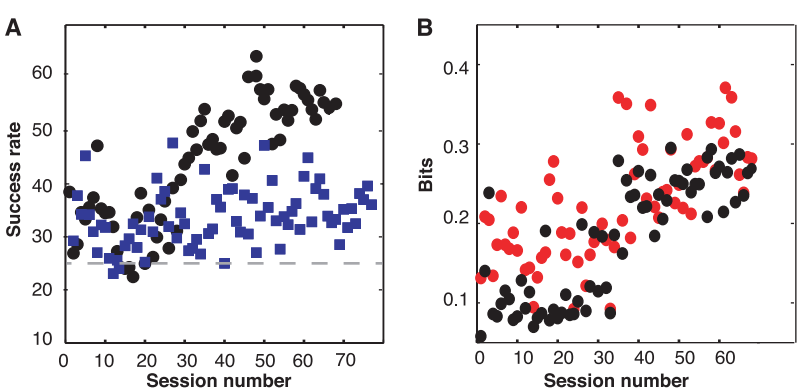

Fig. 2. (**A**) Overall success rates for decoding movement intention from four possible target positions. Black circles, monkey S; blue squares, monkey C. The number of brain control trials varied from 250 to 1100 trials. (**B**) The mean mutual information of the cells from monkey S, whose activity was used to build the database (red) and perform the brain control task (black), is depicted for all 68 sessions. For each session, a selection of cells was chosen on the basis of significant tuning. These cells were then used in the brain control trials. The mutual information of these cells was calculated for the 120 reach trials and the subsequent 120 brain control trials.

PRR neurons increased their spatial tuning during brain control and reach trials when the preferred type of reward (orange juice versus water) was indicated (Fig. 3, A and B) (*19*). The prior knowledge of a high probability of receiving a reward or the impending delivery of a large volume of reward also increased the tuning of these cells (Fig. 3, C and D). The latter two effects were observed for all three animals. The increased activity is unlikely to be due to attention, given that no increase to the expected delivery of the nonpreferred reward was recorded when it was aversive [0.076 M NaCl (*20*)]; the response to the saline was similar to the response of a neutral (water) reward. In addition, the increased activity during the memory period for preferred rewards was not related to an associated increase in muscle activity (*17*) (fig. S3).

For the brain control trials, those ending with the delivery of the preferred reward carried more information than trials ending in nonpreferred rewards (nonpreferred reward: median, 0.062; 95% confidence interval, 0.0571 to 0.0671; preferred reward: median, 0.091; 95% confidence interval, 0.077 to 0.097) (Fig. 4A). Accordingly, the increased information encoded during the preferred reward condition should improve the success with which movement intentions could be decoded. To test this assertion, two independent decodes were run online and in parallel during the brain control task: one for the preferred reward and one for the nonpreferred reward. Within a given experimental session, a single aspect of the reward (size, probability, or type) was varied. The preferred and nonpreferred rewards were randomly interleaved on a trial-by-trial basis (*17*). An example of the performance of these dual decodes, which used six neurons and varied reward size, is shown in Fig. 4B. The expectation of a high volume of reward improved the overall decode success rate during the brain control segment by 12.2%. Over all the sessions, the increase in the expected value for larger reward volume increased the successful online decode of goals by up to 21% (Fig. 4C) with a median of 8.1% ($n = 32$). The increase in decode performance also occurred when probability (median = 10%, $n = 4$) and reward type (median = 9%, $n = 8$) were varied. Taken together, these results show that cells were better tuned during the preferred reward trials, providing greater information about the target location and thereby improving the decode.

The expected value could also be decoded on a trial-by-trial basis from brain control trials. Offline decodes similar to those used for the goal, produced a mean accuracy of 84.7 ± 8.5% (*17*) (fig. S5). Even more importantly,

expected value (preferred versus nonpreferred for type, magnitude, or probability) and reach goals were simultaneously read out with a mean accuracy of 51.2 ± 8.0% (mean ± SD; chance = 12.5%) (Fig. 4D).

The results of this study show that the goal signal can be used as a source of prosthetic control. At first, it would appear that the retinotopic coding of the plan could be problematic for prosthetic applications when subjects are free to move their eyes. However, the activity within the map of space in PRR is updated with each eye movement to maintain activity for the same locations in extrapersonal space (*9*), and the patterns of eye and hand movements are highly coordinated and stereotyped (*21*). In addition, one brain control session from monkey O conducted with free viewing yielded a four-target decode performance of 80.5%, indicating that intended reaches of animals who are allowed free viewing during reach tasks can be read out with PRR activity.

Recently, the human equivalent of PRR has been identified by means of functional magnetic resonance imaging (*22*). One advantage of using PRR as a target for a cortical prosthetic is the visual nature of the area. Somatosensory feedback regarding the outcome of a movement is often lost with paralysis, but vision generally remains intact. Thus, PRR can receive a very direct visual "error signal" for learning to operate a neural prosthetic in the face of paralysis. PRR is also more anatomically remote from the somatosensory and motor pathways that are damaged in paralysis (*23*). It is possible that PRR will show less degeneration than is seen in other cortical areas that are more closely connected to these pathways.

Consistent with work on cortical plasticity (*24*), the animals learned to improve their performance with time. This plasticity is important for subjects to learn to operate a neural prosthetic. The time course of the plasticity in PRR is in the range of 1 or 2 months, similar to that seen in motor areas for trajectory decoding tasks (*2*, *4*).

Short-term improvements in performance were achieved by manipulating the expected value of reward [see also (*25*)]. The expected value of the probability of reward, the size of the reward, and the type of reward were decoded from the activity in the brain control experiments. These signals in PRR have not been previously observed, and parallel a similar finding of expected value in the nearby lateral intraparietal area (*26*) as well as other cortical and subcortical areas (*20*, *27*, *28*). This activity does not appear to be linked to attention, given that PRR is active selectively for reach plans independent of attention (*10*), and also did not show an enhancement of activity when aversive outcomes were expected.

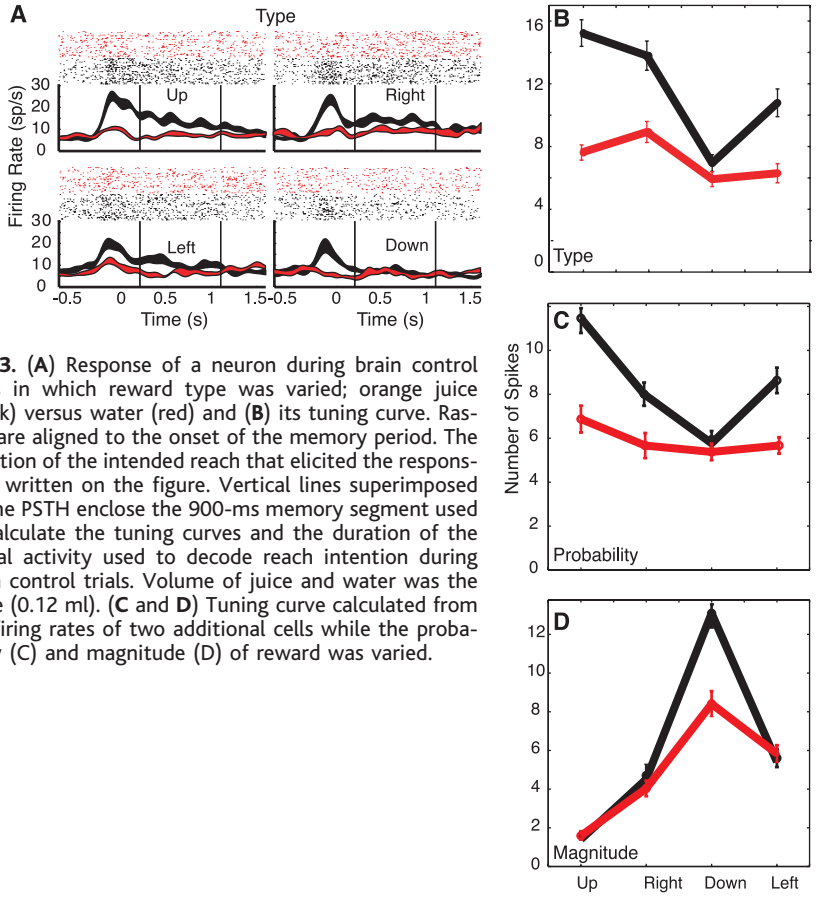

Fig. 3. (**A**) Response of a neuron during brain control trials in which reward type was varied; orange juice (black) versus water (red) and (**B**) its tuning curve. Rasters are aligned to the onset of the memory period. The direction of the intended reach that elicited the responses is written on the figure. Vertical lines superimposed on the PSTH enclose the 900-ms memory segment used to calculate the tuning curves and the duration of the neural activity used to decode reach intention during brain control trials. Volume of juice and water was the same (0.12 ml). (**C** and **D**) Tuning curve calculated from the firing rates of two additional cells while the probability (C) and magnitude (D) of reward was varied.

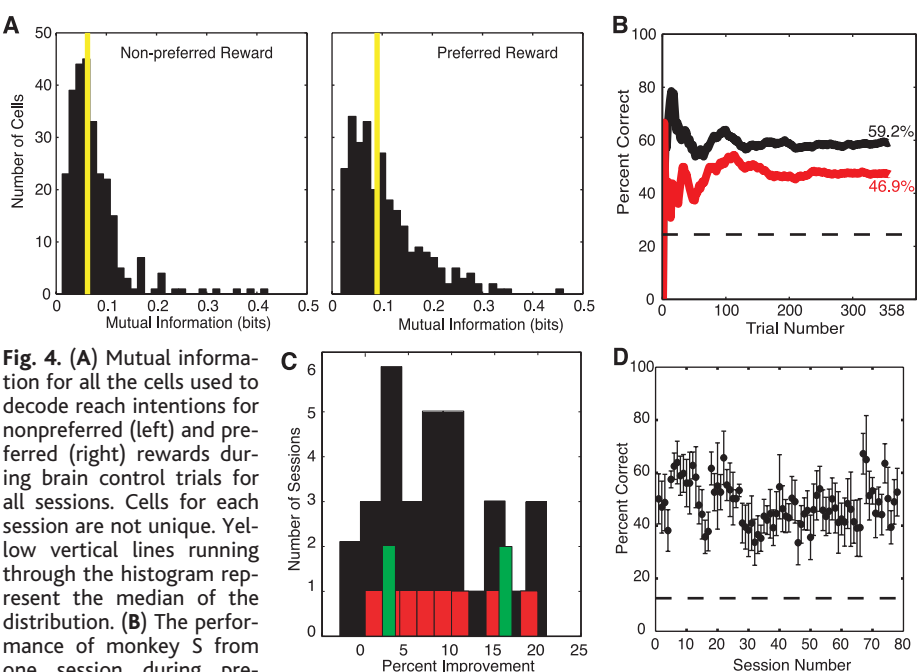

Fig. 4. (**A**) Mutual information for all the cells used to decode reach intentions for nonpreferred (left) and preferred (right) rewards during brain control trials for all sessions. Cells for each session are not unique. Yellow vertical lines running through the histogram represent the median of the distribution. (**B**) The performance of monkey S from one session during preferred (black) and nonpreferred (red) reward conditions. Dashed line represents chance. Decode performance for the two reward conditions is indicated on the plot. (**C**) Improvement in decode between preferred and nonpreferred reward. Black, variable magnitude (high volume, 0.12 ml; low volume, 0.05 ml); red, variable type (juice versus water, volume = 0.12 ml); green, variable probability (high probability = 80%, low probability = 40%). Total number of sessions is 44 (32 reward magnitude, 4 reward probability, and 8 reward type). (**D**) Offline simultaneous decode of four directions and expected value (dashed line shows chance). Error bars show mean ± SD and were obtained by cross-validation (leaving 30 trials out per iteration).

The correlation of increased activity with increased expected reward is substantiated by behavioral data that showed a decrease in reaction times for the preferred rewards (fig. S4). Expected value is a necessary component of the neural system that mediates decision-making (*26*, *29*). On the other hand, it is possible that we are seeing motivational effects that are a direct consequence of expected value (*30*). Further experiments will be required to distinguish between these two explanations.

The decoding of intended goals is an important feature for a cognitive-based prosthetic. Once these goals are decoded, then smart external devices can perform the lower level computations necessary to obtain the goals. For instance, a smart robot can take the desired action and can then compute the trajectory. This cognitive approach is very versatile because the same cognitive, abstract commands can be used to operate a number of devices. The decoding of expected value also has a number of practical applications, particularly for patients that are locked in and cannot speak or move. These signals can directly indicate, online and in parallel with their goals, the preferences of the subject and their motivational level and mood (Fig. 4D). Thus, they could be used to assess the general demeanor of the patient without constantly querying the individual (similar to assessing body language). These signals could also be rapidly manipulated to expedite the learning that patients must undergo to use an external device. Moreover, this research suggests that all kinds of cognitive signals can be decoded from patients. For instance, recording thoughts from speech areas could alleviate the use of more cumbersome letter boards and time-consuming spelling programs, or recordings from emotion centers could provide an online indication of a patient's emotional state.

The cognitive-based prosthetic concept is not restricted for use to a particular brain area, as can be seen by the finding that PRR and PMd activity could both provide goal information. However, some areas will no doubt be better than others depending on the cognitive control signals that are required. Future applications of cognitive-based prosthetics will likely record from multiple cortical areas to derive a number of variables. Moreover, online trajectory information can also be considered as a cognitive variable that can be decoded along with other cognitive variables.

References and Notes

1. M. D. Serruya, N. G. Hatsopoulos, L. Paninski, M. R. Fellows, J. P. Donoghue, *Nature* **416**, 141 (2002).
2. D. M. Taylor, S. I. H. Tillery, A. B. Schwartz, *Science* **296**, 1829 (2002).
3. J. Wessberg *et al.*, *Nature* **408**, 361 (2000).
4. J. M. Carmena *et al.*, *PLoS Biol.* **1**, E42 (2003).
5. F. A. Mussa-Ivaldi, L. E. Miller, *Trends Neurosci.* **26**, 329 (2003).
6. D. J. Felleman, D. C. Van Essen, *Cereb. Cortex* **1**, 1 (1991).
7. P. B. Johnson, S. Ferraina, L. Bianchi, R. Caminiti, *Cereb. Cortex* **6**, 102 (1996).
8. S. P. Wise, D. Boussaoud, P. B. Johnson, R. Caminiti, *Annu. Rev. Neurosci.* **20**, 25 (1997).
9. V. B. Mountcastle, J. C. Lynch, A. Georgopoulos, H. Sakata, C. Acuna, *J. Neurophysiol.* **38**, 871 (1975).
10. L. H. Snyder, A. P. Batista, R. A. Andersen, *Nature* **386**, 167 (1997).
11. A. P. Batista, C. A. Buneo, L. H. Snyder, R. A. Andersen, *Science* **285**, 257 (1999).
12. C. A. Buneo, M. R. Jarvis, A. P. Batista, R. A. Andersen, *Nature* **416** (2002).
13. D. J. Crammond, J. F. Kalaska, *J. Neurophysiol.* **71**, 1281 (1994).
14. D. Boussaoud, F. Bremmer, *Exp. Brain Res.* **128**, 170 (1999).
15. S. Kakei, D. S. Hoffman, P. L. Strick, *Neurosci. Res.* **46**, 1 (2003).
16. H. Scherberger *et al.*, *J. Neurosci. Methods* **130**, 1 (2003).
17. Materials and methods are available as supporting material on *Science* Online.
18. K. V. Shenoy *et al.*, *Neuroreport* **14**, 591 (2003).
19. L. Tremblay, W. Schultz, *Nature* **398**, 704 (1999).
20. W. Schultz, *Nature Rev. Neurosci.* **1**, 199 (2000).
21. D. P. Carey, *Curr. Biol.* **10**, R416 (2000).
22. J. D. Connolly, R. A. Andersen, M. A. Goodale, *Exp. Brain Res.* **153**, 140 (2003).
23. J. A. Turner *et al.*, *IEEE Trans. Neural Syst. Rehabil. Eng.* **9**, 154 (2001).
24. D. V. Buonomano, M. M. Merzenich, *Annu. Rev. Neurosci.* **21**, 149 (1998).
25. E. E. Fetz, *Science* **163**, 955 (1969).
26. M. L. Platt, P. W. Glimcher, *Nature* **400**, 233 (1999).
27. M. Watanabe *et al.*, *Exp. Brain Res.* **140**, 511 (2001).
28. M. I. Leon, M. N. Shadlen, *Neuron* **24**, 415 (1999).
29. J. D. Roitman, M. N. Shadlen, *J. Neurosci.* **22**, 9475 (2002).
30. M. R. Roesch, C. R. Olson, *J. Neurophysiol.* **90**, 1766 (2003).

31. We thank J. Burdick, D. Meeker, B. Pesaran, D. Rizzuto, and S. Cao for helpful discussion during the course of this study; G. Mulliken and R. Battacharyya for help with data collection; B. Grieve, K. Pejsa, and L. Martel for help with animal handling and training; K. Bernheim for help with the magnetic resonance imaging; J. Baer and C. Lindsell for veterinary assistance; I. Fineman for surgical help; V. Shcherbatyuk for computer support; and T. Yao for administrative support. We thank the Defense Advanced Research Projects Agency (DARPA), the National Eye Institute (NEI), the Office of Naval Research (ONR), the James G. Boswell Foundation NSF, the Sloan-Swartz Center for Theoretical Neurobiology at the California Institute of Technology, and the Christopher Reeve Paralysis Foundation for supporting this research. B.D.C. was supported by a long-term fellowship from the Human Frontier Science Program.

Supporting Online Material

www.sciencemag.org/cgi/content/full/305/5681/258/DC1
Materials and Methods
Figs. S1 to S6
References

16 March 2004; accepted 7 June 2004

Supplementary Material

Methods

General

Single cell and multiunit signals were recorded by a multichannel recording system (Plexon Inc, Texas) from 96 paralyne coated tungsten or platinum/iridium electrodes (impedance ≈ 300 kΩ) (Microprobe Inc. Maryland) implanted in the medial intraparietal area (MIP), a subdivision of the parietal reach region (PRR), and area 5 (*1*) of three rhesus monkeys trained to perform a memory reach task. One monkey (monkey S) also had 64 electrodes implanted in the dorsal premotor area (PMd) in a separate surgery. Each session consisted of a *reach segment* and a *brain control segment*. Trials in both segments were initiated in the same way: after the monkeys acquired a central red fixation point with the eyes and touched a central green target, a peripheral cue was flashed indicating the location of one out of four, five, six, or eight reach targets (Figure 1a) (cue epoch). Reach targets were uniformly distributed around the central fixation point. As soon as the fixation point and central green target were acquired, hand and eye movements were restricted by a real time behavioural controller (LabVIEW, National Instruments). Eye position was monitored using a scleral search coil (CNC Engineering, monkeys S and O), or an infrared reflection system (ISCAN, monkey C) while hand position was monitored using an acoustic touch screen (ELO Touch). In order to successfully complete a trial, the monkeys were not allowed to move their eyes. In addition, the reaching hand had to be in contact with the centrally located green target at all times except after the GO signal which appeared during the reach segment of the session. After the offset of the cue, a delay of 1.5 ± 0.3 seconds ensued. During the reach segment, the green central target was extinguished after the memory period

indicating to the animal to reach to the remembered target location (motor epoch). After reaching to the location of the extinguished cue, the monkeys had to maintain contact with the screen for 350ms. If successful, the cue was illuminated and the monkeys had to maintain contact with the visible target for an additional 300ms before they were rewarded. Any break in eye fixation during the trial aborted it.

In the brain control trials the intended reach location was decoded from a 900 ms interval of the delay period starting 200 ms after cue offset. Unless otherwise noted, all brain control analysis and tuning curves of cells presented in this study are based on this 900 ms interval. If the correct position was decoded, the cue location was illuminated with a larger cue and the monkeys received a reward. The monkeys were not allowed to reach or break fixation until after the reward had been delivered. No feedback was given to the monkeys when the wrong target location was decoded. Instead, the green central target was extinguished indicating to the monkeys to initiate a reach. Therefore the monkeys had to continue with a failed decode trial as if it was a reach trial. The adaptive database was not updated after the failed decode trials (see below).

Variable Reward

Only a single aspect of the reward (magnitude, probability or type) was varied during a given session. The size of the cue indicated the magnitude or type of reward, or the probability of obtaining a reward. In order to control for the effects of cue size on the firing rate, the association between cue size and reward was varied on different sessions. The mapping of cue size to reward condition had no effect on the representation of expected value. The magnitude of the reward was 0.05ml and 0.12 ml for low and high volume respectively. When probability was varied, a constant volume reward (0.12 ml)

was delivered either 40% or 80% of the time upon successful completion of the trial. Hence, the monkeys were not rewarded on all trials but had to complete all the trials presented. When reward size or type was varied, reward probability was fixed at 100%. Reward type (orange juice vs. water) was delivered from two sipper tubes that were calibrated to dispense equal volumes of liquid. The sipper tubes were placed near the monkey's mouth with the location of the tube altered on different days. No effect of juice tube placement on the firing rate was found.

Database

During the reach trials, the activity of all the cells was recorded and a database containing the firing rates was constructed. Once enough trials were collected (30 reaches for each target location except the PMd recordings which used 20 reaches per target location), the brain control segment of the task was initiated. The goal of this segment was to have the monkeys position a cursor on the screen with their thoughts. A selection of single and multiunit activity was then chosen from the database predicated on their tuning properties assessed using an ANOVA on the spiking activity during a 900 ms interval of the memory period beginning 200 ms after cue offset. Many more neurons than those chosen were recorded from the arrays. For example, many cells exhibited strong visual and motor related activity. This activity was also tuned and can easily be used to decode target location with a high success rate (Figure S1). However, our goal is to decode intentions represented by cognitive signals and not responses directly elicited by stimuli or neural activity associated with overt movements. Therefore, only those neurons that showed clear tuning in the memory period as assessed by the ANOVA were chosen for the decode (monkey S (parietal): range = [5, 13], median = 6; monkey C:

range = [4, 48] median = 6.5; monkey O: range = [6, 10], median = 7.5; monkey S (premotor): 3 sessions using 8, 15 and 16 neurons. median = 15). (For the first 5 sessions of monkey C, we did not utilize the results of the ANOVA but instead used 48 channels for the decode).

Decode Algorithm

The movement intention that was encoded by the neural activity in the memory period for each trial in the brain control task was then decoded using a Bayesian algorithm on a family of Haar wavelet coefficients (*2*). Bayes rule is defined by $P(s \mid r) = \frac{P(r \mid s)P(s)}{P(r)}$ where r is the response and s is the target direction. $P(s|r)$ was calculated for all directions and the direction decoded taken to be the maximum of all $P(s|r)$. 100 wavelet coefficients were calculated by projecting the spike train recorded during 900 ms of the memory period onto a family of Haar wavelet functions. In this way, temporal features in the spike train that cannot be described by the number of spikes in the memory period (equivalent to firing rate) were exploited (*2*). Haar wavelets are defined by (*3*):

$$\Psi(t) = \begin{cases} 1 & if\ 0 \le t \le \frac{1}{2} \\ -1 & if\ \frac{1}{2} \le t < 1 \\ 0 & otherwise \end{cases}$$

where dilations and translations of Ψ generate the orthonormal basis:

$$\Psi_{j,n}(t) = \frac{1}{\sqrt{2^j}} \Psi\left(\frac{t - 2^j n}{2^j} \right)$$

where j and n are integers that represent the frequency content of the wavelet. Note that the zeroth wavelet coefficient ($j,n,t = 0$) is simply the number of spikes in the 900 ms

portion of the memory epoch used in the decode since the wavelet being projected onto it is the step function. The Haar wavelets improved the Bayesian decode by taking advantage of the temporal features of the spike train in the memory period. Although we calculated a large number of coefficients, only a few (usually less than 5) had relevant information. The optimal coefficients can be calculated by applying sorting algorithms to the coefficients based on information theory (2).

Offline decode on 10 sessions using a Bayesian algorithm with wavelets yielded a performance that was on average 6.6 ± 2.9 % better than offline decode that did not use the wavelets (range = [-0.4 9.1]). The number of spikes in the memory period (zeroth wavelet coefficient) yielded the greatest amount of information about the intended goal. The first wavelet coefficient also yielded tuned information useful for decode. The significance of this coefficient implies that the delay period had a different rate at the first and second half of the memory period that was useful for decoding.

The brain control session shown in the left panel of Figure 1D is based on a database composed of 20 reaches / direction. However, we used 50 reaches / direction to build the database for the offline decode. Not only did the decode performance improve using a greater number of neurons, but it also improved by using a greater number of trials in the database (87 % for 8 targets using all 16 neurons; Figure 1D right panel). However, for the 4 target decode, which is the main experimental condition used in this study, 30 reaches per direction was optimal as indicated by offline simulations.

Off-line decode results suggest that we can also improve the performance using larger training sets with a Fisher linear discriminant (FLD) algorithm. Using data obtained during brain control trials to run offline decodes, FLD improved the decode by 8.7 ± 6.2 % (mean ± standard deviation). However, we decided not to use this algorithm

on-line since the number of trials needed in the training set that would yield a decode performance better than the Bayesian algorithm approached 100 reaches / direction. This would substantially reduce the number of decode trials and was not even possible for some 6 target sessions. The use of a database with a small number of trials is more advantageous for neural prosthetics since patients do not need to be burdened with prolonged training sessions.

Adaptive vs. Frozen databases

Most sessions were run using the adaptive database but, on occasion, the frozen database was used (189 adaptive, 10 frozen). The adaptive database was simply a fixed (30 trials per direction) database moving in time. The database was continuously updated with new successful trials while old trials were removed using a first-in-first-out rule. This way, the database always contained the latest 30 successful trials per direction. The frozen database was composed of the trials collected during the reach segment of the session. Thereafter, the database was not updated but was frozen in time. Offline analysis indicated that no advantage was gained by using either approach (Figure S2). Mean success rate for monkey S for all sessions achieved using the adaptive decode (mean ± standard deviation) is 43.3 ± 11.1 % while the success rate using the frozen database is 42.1 ± 11.7 %. The two distributions are not statistically different.

Mutual Information

Mutual information is a useful measure as it quantifies the degree of tuning of a neuron as opposed to a statistical p – value which merely provides a probability of

whether a neuron is tuned or not (*4, 5*). The information carried by neurons was calculated using (*6*):

$$I(r,s) = \sum_{r,s} P(r,s) \log \frac{P(r,s)}{P(r)P(s)}$$

where *s* is the target direction and *r* is the response and the *log* is base 2. For brain control and reach trials, the mutual information was calculated on an equal number of trials. The joint distribution *P(r,s)* was estimated using a 2-D histogram between the stimulus and the response. The number of directions in a particular session dictated the number of stimulus bins in the histogram. Eight bins were used for the response which places the histogram outside the sparse region (*5*). The marginal distribution of the 2-D histogram was then used as an estimate of the probabilities *P(r)* and *P(s)*.

Learning statistics.

Figure 2B in the main text shows the mean mutual information from reach and brain control for each of 68 consecutive sessions for monkey S. During the first 20 sessions, the information about target location is high during the reach segment (when the database was being built) and much lower during the brain control segment for the same cells. The mean of the difference in the mutual information between the reach and the brain control segments for the first 20 sessions was 0.11 ± .002 bits (mean ± standard error) and was significantly different from zero (t-test, $P < 0.01$). The difference for the last 20 sessions was 0.028 ± .004 bits, also significantly different from zero ($P < 0.01$). However, the difference during the first 20 sessions is significantly greater than the difference from last 20 sessions ($P < 0.01$). Therefore, the information carried by cells recorded during the same session increased more during the brain control segment than

during the reach segment over the course of 68 sessions. This effect can also be seen by considering the rate of information increase within the reach and brain control segments. The regression slope for the mutual information during the reach segment was 0.0023 ± 0.0003 bits/session while the slope of the best fit regression line for the mutual information during brain control was 0.0031 ± 0.0003 bits/session. Both these slopes are significantly different from zero (t-test, $P < 0.01$). However, the rate of increase of the information during the brain control segments is greater than the rate of increase during the reach sessions (regression of the difference between reach trials and brain control trials is -0.0018 ± 0.0004 bits/session which is significantly less than zero $P < 0.01$). The same effect was shown by monkey C with less difference (regression of the difference between reach trials and brain control trials is -0.0008 ± 0.0003 bits/session which is significantly less than zero, $P < 0.01$).

The slope of the performance as a function of session number is 0.48 ± 0.25 percentage points / session for the last 10 sessions which is statistically greater than 0 ($p<0.02$) (Figure 2A). This positive slope implies that the performance may have continued to increase if more sessions were performed.

Electromyography

Percutaneous EMGs were recorded from the anterior deltoid (Figure S3), posterior deltoid, the rhombus and the lower trapezius of monkey C over 5 separate sessions during reach trials. EMG's were low-pass filtered (cutoff 1000hz), sampled at 2500 Hz and rectified before alignment. If the neural activity of the memory period was related to the direct activation of muscles, then increased EMG should be observed when the monkey is planning a movement in a muscle's preferred direction. Likewise, if the

increased direction tuning during preferred rewards is related to muscle activation then there should be an increase in EMG direction tuning for the preferred rewards. For all individual muscles tested, there was no statistically significant EMG directional tuning in the delay period during brain control trials for either the low or high reward condition. For the anterior deltoid example shown in Figure S3, the EMGs during the memory period increased by up to 4% when the preferred reward was indicated but this increase was not directionally tuned. Changing the reward had no effect on the activity of the rhombus, posterior deltoid or the trapezius. Thus, the directionally tuned increase in neural activity recorded in the high reward condition during the memory period was not associated with a significant directionally tuned increase in limb EMG.

Reaction Time

Reward manipulation also affected behavioral performance. The reaction time from the GO signal to the beginning of the movement of all trials during the reach segment of sessions using variable rewards was calculated. The expectation of preferred reward decreased the mean reach reaction time from 320 ± 1.51 ms to 309 ± 1.35 ms (mean ± standard error) (Figure S4). Reaction time is significantly smaller for the preferred reward condition ($P < 0.01$). This enhanced motor performance is consistent with increased motivation.

Expected Value Decode

We can decode the expected value of the reward. Figure S5A depicts an off-line decode of expected value of reward type using a frozen database for one brain control session. This binary decode was run independent of reach direction. For the same cells

used to decode direction, we can correctly identify whether the reward on the current trial is orange juice or water over 85% of the time (Figure S5A). Repeating this analysis over all the sessions for monkeys S and C and O, we can decode the expected value with an overall mean of 80.4 ± 8.6 %. For reach trials, decode performance was 74.5 ± 5.2 %. For brain control trials, decode performance was 84.7 ± 8.5 % (mean ± standard deviation) (Figure S5B).

Decode Interval Length

Offline decode on brain control trials indicated that memory period intervals as low as 100 ms can yield decode rates that are significantly greater than chance (Figure S6). There is however a steady increase in the performance as the interval size increased. No conclusion can be made on whether asymptotic behavior was reached since performance continued to increase as the limit of the memory period was reached. Time intervals extracted from the beginning of the memory period yielded better feedback performance than the same sized intervals from latter portions of the memory period (based on offline analysis on 5 sessions with high performance from monkey S. Brain control performance for 100 ms obtained from the beginning of the memory period was 51.2 ± 4.7 % while 100 ms obtained from the end of the memory period yielded a performance of 38.9 ± 7.2 % (chance 25%)).

Figure S1

Cumulative percent of correctly decoded trials using 700 ms of the motor burst (-100 ms to 600 ms after the GO signal in reach trials) of 4 parietal neurons during reach trials for 1 session.

Figure S2

Offline decode results performed with an adaptive (red) and frozen (black) database for all the sessions in consecutive order for monkey S. No statistical difference exists between the 2 populations.

Figure S3.

Percutaneous EMG recorded from the anterior deltoid of monkey C during reach trials. Black: high reward. Red: small reward. Plots are aligned to the onset of the cue. Reach directions are indicated on the plot. EMGs were smoothed with a moving window of 10 trials.

Figure S4.

Reach reaction time for preferred (n = 6671 reaches) and nonpreferred (n = 7180 reaches) conditions for monkeys S and C. Bars are SE.

Figure S5.

A) Decode result of expected value from a single brain control session and B) all the sessions where expected value of reward was manipulated. Error bars are standard deviation obtained by crossvalidation (leaving 30 trials out per

iteration). Sessions are not in consecutive order. The first 36 sessions are reach sessions (red) and the last 44 sessions are brain control sessions (black). Dashed line is chance.

Figure S6.

Offline decode on 16 sessions from monkey S using various time interval lengths of the memory period. Note that the time on the x-axis is not continuous but represents the length of the memory period that yielded the corresponding feedback performance. All intervals shown start 200 ms after the offset of the cue and last for the duration indicated on the x-axis. For example, the corresponding y -value at the interval marked 0.2 seconds corresponds to 200ms of the memory period starting 200 ms after the onset of the memory period (201-400ms of the memory period).

Figure S1

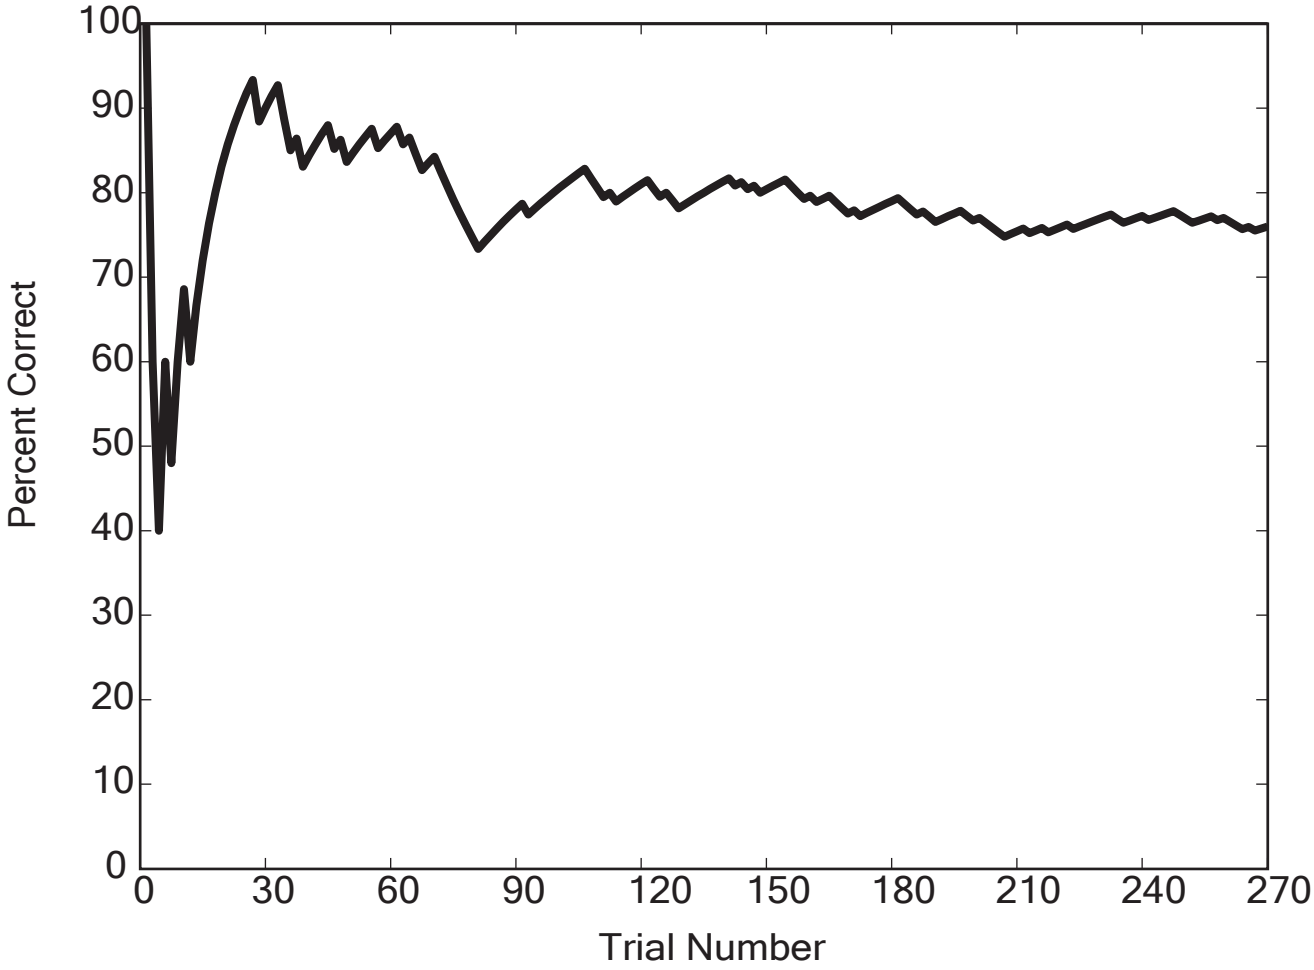
Percent Correct
100
90
80
70
60
50
40
30
20
10
0
0
30
60
90
120
150
180
210
240
270
Trial Number

Figure S2

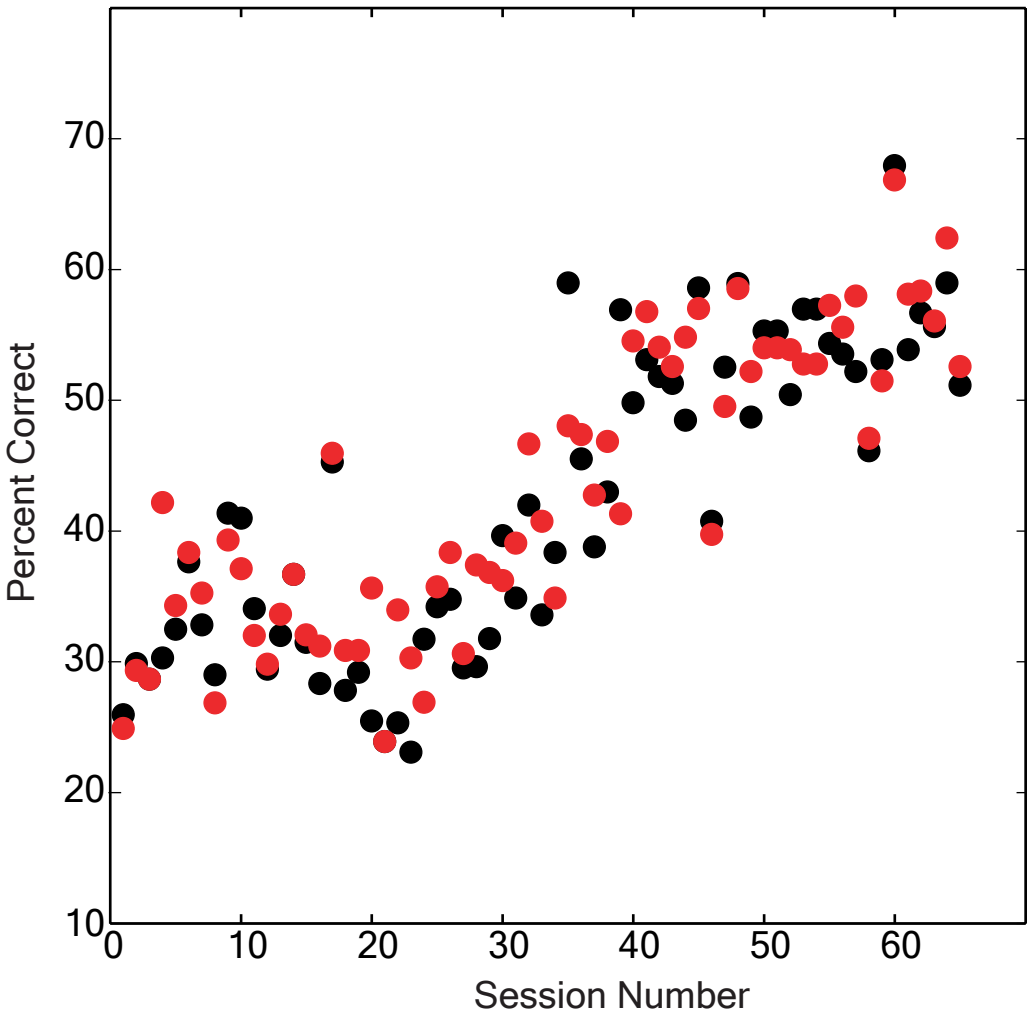
Percent Correct
Session Number
10
20
30
40
50
60
70
0
10
20
30
40
50
60

Figure S3

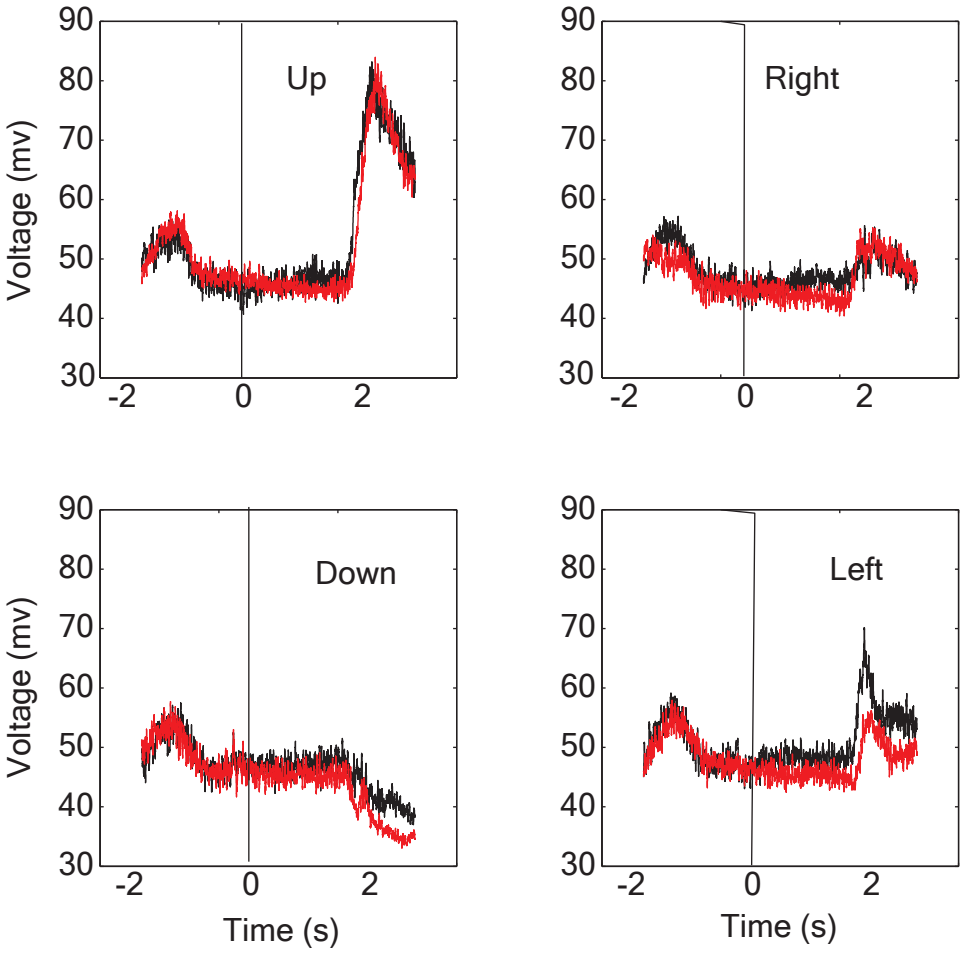
Up
Right
Down
Left
Voltage (mv)
Voltage (mv)
Time (s)
Time (s)
90
80
70
60
50
40
30
-2
0
2

Figure S4

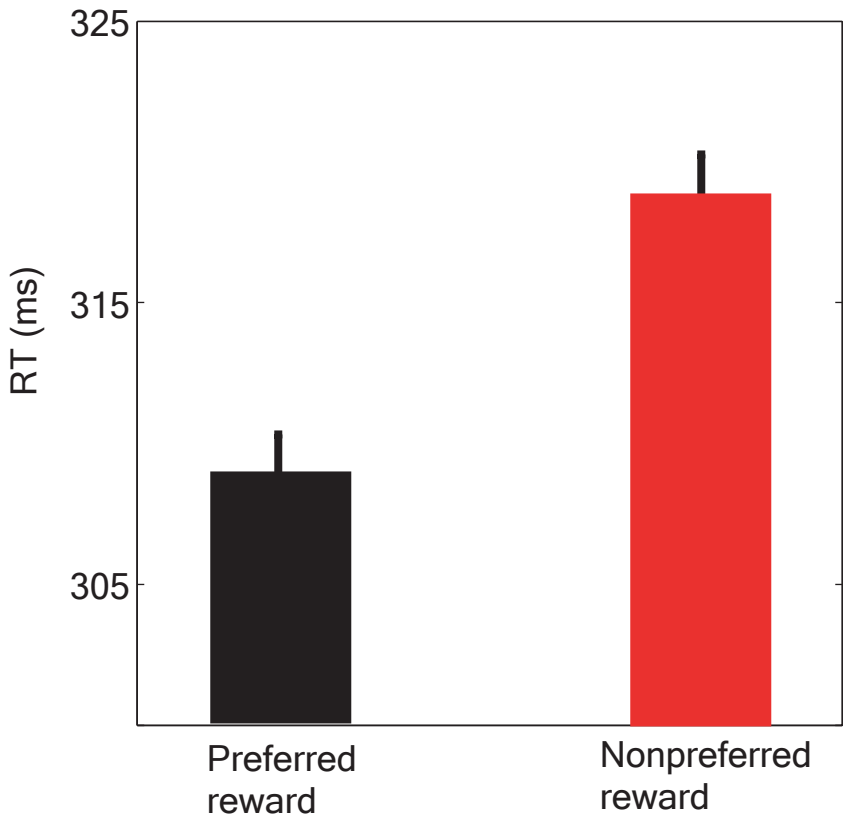

Figure S5

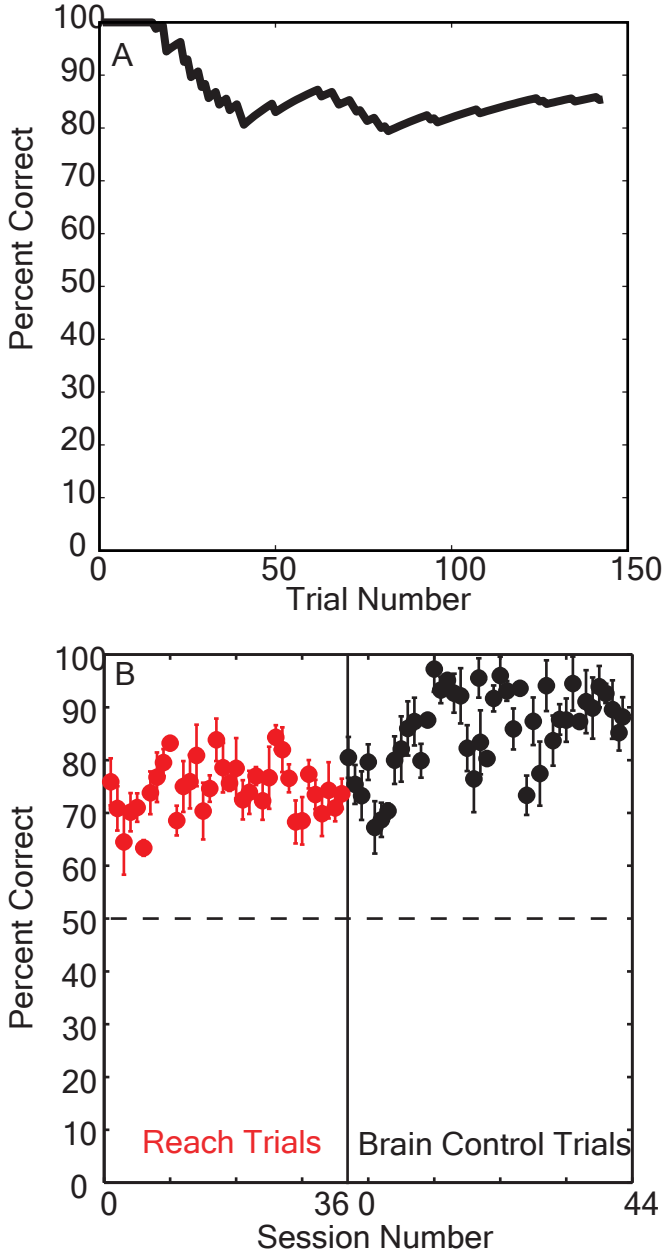
A
Percent Correct
Trial Number
B
Reach Trials
Brain Control Trials
Session Number

Figure S6

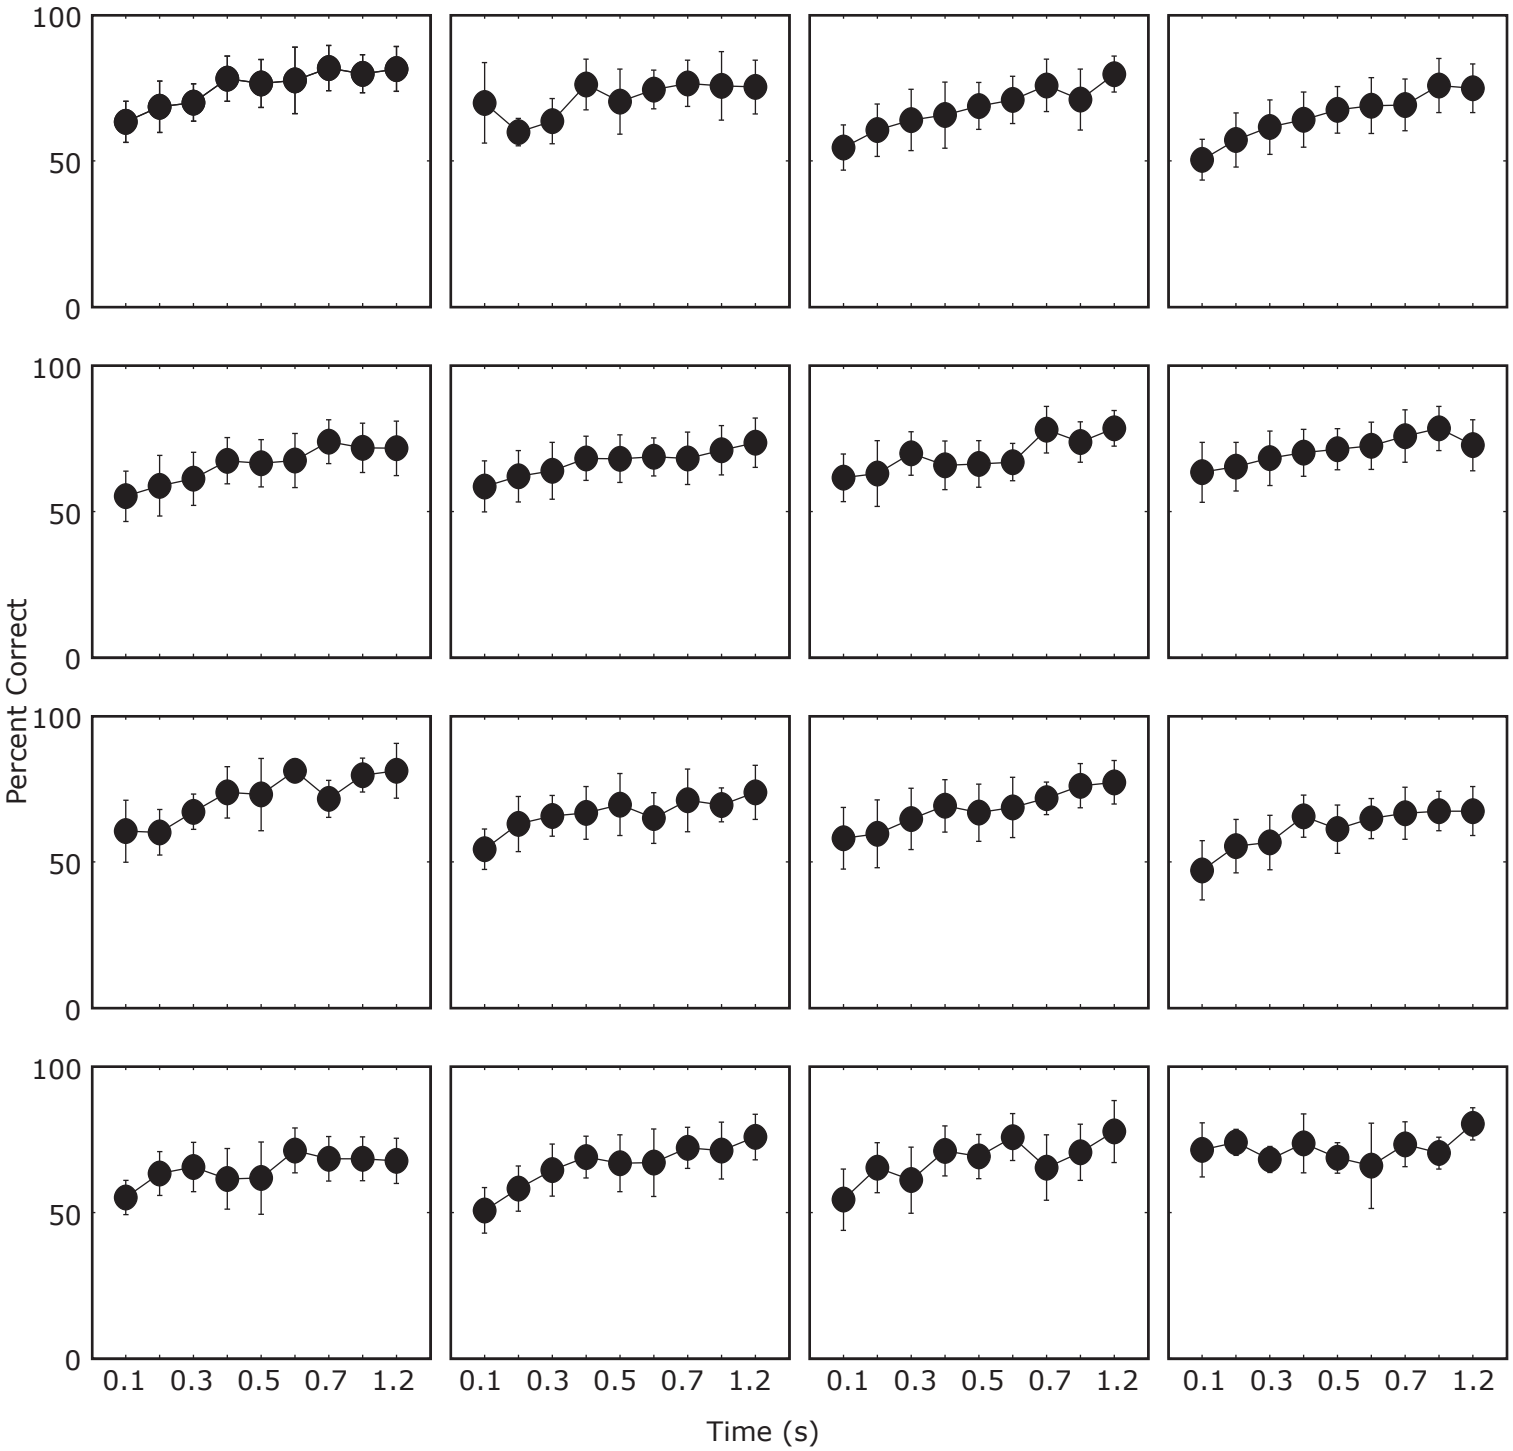

References

1. H. Scherberger *et al.*, *J Neurosci Methods* **130**, 1 (2003).
2. S. Cao, PhD Thesis. *Department of Electrical Engineering* (2003).
3. S. Mallat, *A Wavelet Tour of Signal Processing* (Academic Press, 2001).
4. A. Grunewald, J. F. Linden, R. A. Andersen, *J Neurophysiol* **82**, 330 (1999).
5. S. P. Strong, R. Koberle, R. R. D. van Steveninck, W. Bialek, *Phys Rev Let* **80**, 197 (1998).
6. T. M. Cover, J. A. Thomas, *Elements of Information Theory* (John Wiley and Sons, 1991).