

LETTERS

Cortical control of a prosthetic arm for self-feeding

Meel Velliste¹, Sagi Perel^{2,3}, M. Chance Spalding^{2,3}, Andrew S. Whitford^{2,3} & Andrew B. Schwartz¹⁻⁶

Arm movement is well represented in populations of neurons recorded from the motor cortex¹⁻⁷. Cortical activity patterns have been used in the new field of brain-machine interfaces⁸⁻¹¹ to show how cursors on computer displays can be moved in two- and three-dimensional space¹²⁻²². Although the ability to move a cursor can be useful in its own right, this technology could be applied to restore arm and hand function for amputees and paralysed persons. However, the use of cortical signals to control a multi-jointed prosthetic device for direct real-time interaction with the physical environment ('embodiment') has not been demonstrated. Here we describe a system that permits embodied prosthetic control; we show how monkeys (*Macaca mulatta*) use their motor cortical activity to control a mechanized arm replica in a self-feeding task. In addition to the three dimensions of movement, the subjects' cortical signals also proportionally controlled a gripper on the end of the arm. Owing to the physical interaction between the monkey, the robotic arm and objects in the workspace, this new task presented a higher level of difficulty than previous virtual (cursor-control) experiments. Apart from an example of simple one-dimensional control²³, previous experiments have lacked physical interaction even in cases where a robotic arm^{16,19,24} or hand²⁰ was included in the control loop, because the subjects did not use it to interact with physical objects—an interaction that cannot be fully simulated. This demonstration of multi-degree-of-freedom embodied prosthetic control paves the way towards the development of dexterous prosthetic devices that could ultimately achieve arm and hand function at a near-natural level.

Two monkeys were implanted with intracortical microelectrode arrays in their primary motor cortices. Each monkey used the signals to control a robotic arm to feed itself. The robotic arms used in these experiments had five degrees of freedom: three at the shoulder, one at the elbow and one at the hand. Like a human arm, they permitted shoulder flexion/extension, shoulder abduction/adduction, internal/external rotation of the shoulder and flexion/extension of the elbow. The hand consisted of a motorized gripper with the movement of its two 'fingers' linked, providing proportional control of the distance between them. Monkeys were first trained to operate the arm using a joystick (Supplementary Methods). Their own arms were then restrained and the prosthetic arm was controlled with populations of single- and multi-unit spiking activity from the motor cortex. The neural activity was differentially modulated when food was presented at different target locations in front of the monkey. Based on previous work²⁴, we used this modulation to represent velocity of the prosthetic arm's endpoint (a point between the fingertips of the hand/gripper) as an expression of the intention to move^{2,3}. The recorded signal was also used by the subject to open and close the gripper as it grasped and moved the food to the mouth. The endpoint velocity and gripper command were extracted from the instantaneous firing rates of simultaneously recorded units using a real-time extraction algorithm.

Many algorithms of varying complexity have been developed in open-loop^{7,25-27} or closed-loop experiments¹²⁻²⁴, but here we show that a simple algorithm functioned well in this application. The population vector algorithm²⁸ (PVA) used here was similar to algorithms used in some cursor-control experiments^{15,21}. It relies on the directional tuning of each unit, characterized by a single preferred direction in which the unit fires maximally. The real-time population vector is essentially a vector sum of the preferred directions of the units in the recorded population, weighted by the instantaneous firing rates of the units, and was taken here to represent four dimensions—velocity of the endpoint in an arbitrary extrinsic three-dimensional cartesian coordinate frame, and aperture velocity between gripper fingers (fourth dimension). The endpoint velocity was integrated to obtain endpoint position, and converted to a joint-angular command position, for each of the robot's four degrees of freedom, using inverse kinematics. Degree-of-freedom redundancy was solved by constraining elbow elevation in a way that resulted in natural-looking movements (Supplementary Methods). As the monkey's cortical command signal was decoded in small time-increments (30 ms), the control was effectively continuous and the animal was able to continuously change the speed and direction of arm movement and gripper aperture. Details of the control algorithm are in Supplementary Methods.

To demonstrate fully embodied control (Fig. 1), monkeys learned a continuous self-feeding task involving real-time physical interaction between the arm, a food target, a presentation device (designed to record the target's three-dimensional location) and their mouth. Unlike short control windows used in previous studies, each monkey controlled the arm and gripper continuously during an entire session (not only during reaching and retrieval movements but also during loading/unloading and between trials). The task was challenging owing to the positional accuracy required (about 5–10 mm from the target centre position at the time of gripper closing). The required accuracy for retrieval was much lower because the monkey could move its head to meet the gripper. Supplementary Video 1 shows monkey A performing seven consecutive successful trials of continuous self-feeding. It can be seen from the video that the monkey was still chewing on the previous piece of food while reaching for the next one. It can also be seen that the monkey was able to move its head and eyes naturally without affecting control of the prosthetic arm. Example signals from the last four trials of the video show the correspondence between the spike signals of the 116 units used for control during that session and the resulting arm and gripper movement (Fig. 2).

Monkey A performed 2 days of the continuous self-feeding task with a combined success rate of 61% (67 successes out of 101 attempted trials on the first day, and 115 out of 197 on the second day). To put this success rate in perspective, a task of comparable difficulty to a previous virtual cursor control study from our group¹⁵ would be to simply move the prosthetic arm's endpoint near the target

¹Department of Neurobiology, School of Medicine, E1440 BST, Lothrop Street, University of Pittsburgh, Pittsburgh, Pennsylvania 15213, USA. ²Department of Bioengineering, 749 Benedum Hall, University of Pittsburgh, Pittsburgh, Pennsylvania 15261, USA. ³Center for the Neural Basis of Cognition, University of Pittsburgh and Carnegie Mellon University, Pittsburgh, Pennsylvania 15213, USA. ⁴Department of Physical Medicine and Rehabilitation, University of Pittsburgh, Pittsburgh, Pennsylvania 15213, USA. ⁵McGowan Institute for Regenerative Medicine, University of Pittsburgh, Pittsburgh, Pennsylvania 15219, USA. ⁶Robotics Institute, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213, USA.

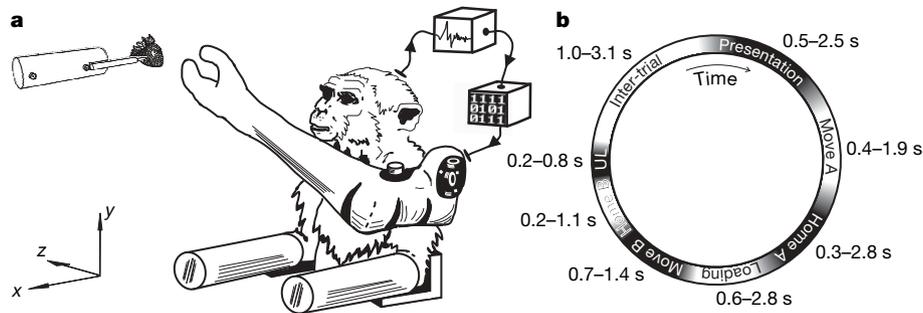


Figure 1 | Behavioural paradigm. **a**, Embodied control setup. Each monkey had its arms restrained (inserted up to the elbow in horizontal tubes, shown at bottom of image), and a prosthetic arm positioned next to its shoulder. Spiking activity was processed (boxes at top right) and used to control the three-dimensional arm velocity and the gripper aperture velocity in real time. Food targets were presented (top left) at arbitrary positions.

b, Timeline of trial periods during the continuous self-feeding task. Each trial started with presentation of a food piece, and a successful trial ended with the monkey unloading (UL) the food from the gripper into its mouth (see Methods). Owing to the continuous nature of the task, there were no clear boundaries between the task periods.

(that is, complete the Move A period only, without being required to home in, load, retrieve and unload). (The Move A period is defined in Methods, and shown within the timeline in Fig. 1b.) Monkeys in that previous study had a success rate of 80%, whereas our monkey A successfully completed the Move A period in 98% of attempted trials (Supplementary Table 3). Distance of the targets in this task

(184 ± 31 mm, mean \pm s.d.) was also greater than that in the previous study. Monkey P performed a version of the continuous self-feeding task (Supplementary Video 2) with an average success rate of 78% (1,064 trials over 13 days), typically using just 15–25 cortical units for control. Monkey P’s success rate was higher than monkey A’s because the task was easier (see Supplementary Methods).

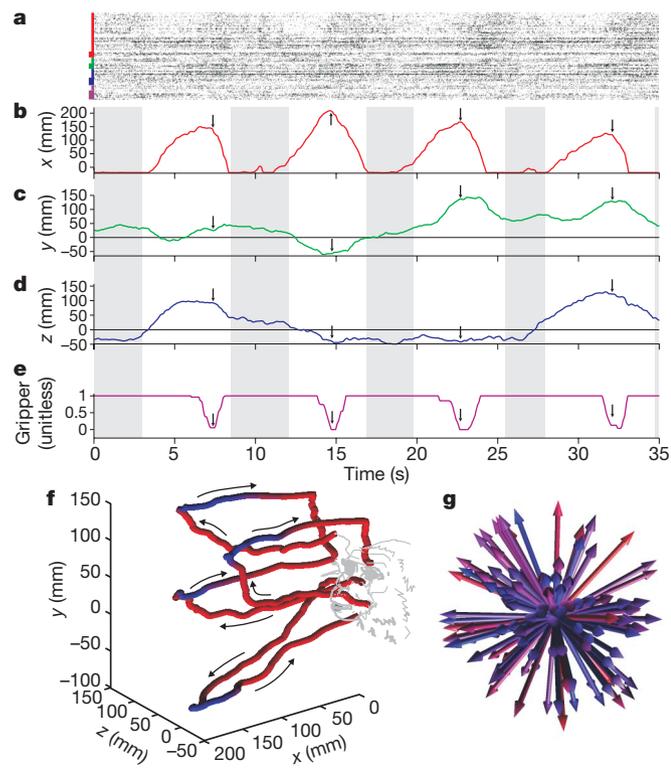


Figure 2 | Unfiltered kinematic and spike data. **a**, Spike rasters of 116 units used for control. Rows represent spike occurrences for each unit, grouped by major tuning component (red, x; green, y; blue, z; purple, gripper). Groups are further sorted by negative major tuning component (thin bar) versus positive (thick bar). **b–d**, The x, y, and z components, respectively, of robot endpoint position. Grey background indicates inter-trial intervals. Arrows indicate gripper closing at target. **e**, Gripper command aperture (0, closed; 1, open). **f**, Spatial trajectories for the same four trials. Colour indicates gripper aperture (blue, closed; purple, half-closed; red, open). Arrows indicate movement direction. **g**, Distribution of the four-dimensional preferred directions of the 116 units used. Arrow direction indicates x, y, z components, colour indicates gripper component (blue, negative; purple, zero; red, positive).

The fact that the gripper opens and closes fully each time (Fig. 2e) indicates good performance, because full opening is advantageous on approach to target and full closing is required for loading. The fact that the task requirements allow the monkey to drive the gripper aperture to both limits makes this fourth dimension easier to control than the x, y and z dimensions. However, the monkey is capable of partially opening or closing the gripper, as shown by data from an earlier training session (Supplementary Fig. 12).

Figure 2f reveals a surprising point: after gripping the food and pulling it off the presentation device, the monkey gradually opened the gripper on the way back to the mouth (Move B) and the gripper was typically fully open before it reached the mouth. One might expect the food to have dropped when the gripper was opened, but this was not always the case because marshmallows, and even grape halves to some extent, tended to stick to the gripper fingers. In an earlier training session, the monkey kept the gripper closed all the way back to the mouth (Supplementary Fig. 13). Over the course of training, the monkey must have learned that keeping the gripper closed was unnecessary, illustrating the importance of working within a physical environment.

We assume that an arm that moves naturally with a bell-shaped speed profile^{29,30} will be easier to control than one that moves in an unfamiliar way. Monkey A’s individual-trial profiles (Fig. 3a) show a large bell-shaped peak for retrieval movements. Reaching movements consist of multiple smaller bell-shaped peaks indicative of corrective movements. The speed profiles shared qualitative characteristics with natural movements, but the duration of prosthetic movements (3–5 s for monkey A, including reaching, loading and retrieval) is not yet down to the same level as natural movements (1–2 s). The corrective movements and long movement duration are consistent with extensive use of visual feedback in this task.

The animal controlled the exact path of the arm to achieve the correct approach direction to position the gripper in the precise location needed to grasp the food. This was demonstrated by the curved path taken to avoid knocking the food piece off the presentation device (Fig. 3b and Supplementary Video 3). It is also important that there be no apparent control delay—that is, lag between the desire to move and the movement of the prosthetic. The delay between spike signals and movement of the robotic arm was approximately 150 ms (Supplementary Methods). This is not very different from the control delay of a natural arm⁶. An example of lag-free control can be seen in Supplementary Video 2, where the food

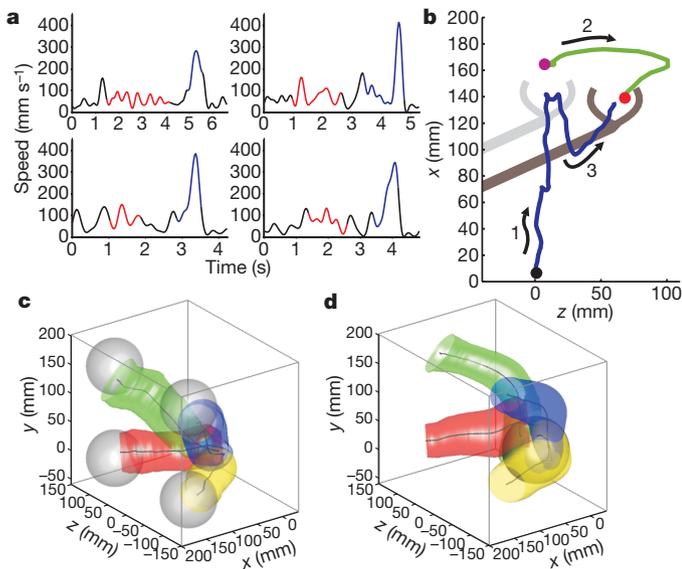


Figure 3 | Movement quality. **a**, Speed profiles from four trials. Time zero marks the beginning of forward arm movement. Reaching (red) begins when the target is in position and ends when the gripper touches the target or minimal distance between target and endpoint is achieved (whichever comes first). Retrieval (from food off the presentation device to mouth contact) is blue, and the graph ends with food in the monkey's mouth (obtained from video record). **b**, Target tracking. Endpoint trajectory (blue, arrow 1) from an initial position (black dot) towards an initial target (purple dot). When the gripper was about to arrive (light grey sketch) at the initial target, the target was shifted (green trajectory, arrow 2) to a new position (red dot). The monkey then moved the arm in a curved path (arrow 3) to avoid knocking the food off the presentation device, positioning the gripper (dark grey sketch) to grasp the food. This trial is also shown in Supplementary Video 3. **c, d**, Endpoint trajectory variability (monkey A) for reaching (**c**, Move A period) and retrieval (**d**, Move B). Semi-transparent coloured regions represent trajectory standard deviation (over all sessions) around average trajectories (grey lines) to each target. Grey spheres (radius 46 mm, averaged over all sessions) represent regions where training assistance was applied.

dropped out of the gripper unexpectedly during a retrieval movement and the animal immediately stopped moving the arm.

Some displays of embodiment would never be seen in a virtual environment. For example, the monkey moved the arm to lick the gripper fingers while ignoring a presented food target (Supplementary Video 4), and sometimes used the gripper fingers to give a second push to the food when unloading (Supplementary Video 5). These behaviours were not task requirements, but emerged as well

capabilities were learned, demonstrating how the monkey used the robot arm as a surrogate for its own.

The monkeys' arms were restrained in these experiments to prevent them from grabbing the food directly with their own hands. The restraints did not prevent them from making small wrist and hand movements. In particular, monkey A can be seen making characteristic movements with its right hand (Supplementary Video 1): extending the wrist and fingers while closing the prosthetic gripper, then rotating its wrist and flexing the fingers while retrieving the food with the prosthetic arm. It could be argued that these movements might facilitate prosthetic control. However, there are several reasons we find this unlikely. First, the electrode array was implanted in the right hemisphere (the same side as the monkey's own moving hand), while predominant motor cortical output projects to the opposite side of the body. Second, the monkey's hand movement was only loosely coupled to prosthetic control. For example, the temporal correspondence between wrist extension and gripper closing varied between zero and almost a full second (Supplementary Table 4). Third, movement is not required for brain-controlled tasks, as monkeys in other studies made no movement with their arms^{14,15,17,22} and paralysed humans have well modulated motor cortical activity capable of driving prosthetic devices^{12,13,20}. The arm and hand movements seen here may be vestigial, remnants of the joystick task carried out during initial training.

As an intermediate training step towards continuous self-feeding, after the monkeys learned to operate the device with a joystick, they performed an assisted brain-controlled task where the monkey's control was mixed with automated control. The types and amounts of assistance were configurable in each task period. For example, during the Home A and Loading periods (defined in Methods), the training program partially guided the endpoint towards the target by adding a vector pointing towards the target to the endpoint velocity. Gripper opening was partially aided during Move A and Home A by adding a positive value to aperture velocity, and closing was aided during Loading by adding a negative value. Monkey P also used another type of assistance, where the amount of deviation from a straight line towards the target was limited by a gain factor. The relative proportion of all types of automated assistance in the overall control signal was reduced over several weeks until both the arm endpoint movement and gripper were controlled purely by the monkey's cortical command. Full details of assisted control are in Supplementary Methods. Targets during the training period were presented at four discrete locations. This allowed a measure of trajectory consistency to be computed over repeated trials (Fig. 3c and d). Like natural arm movements, the reaching and retrieval movements of the prosthetic arm show some variability, but are generally consistent between trials.

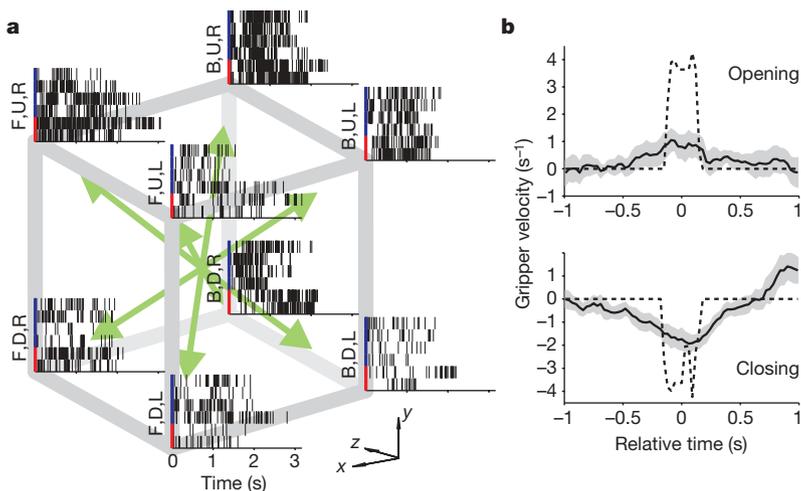


Figure 4 | Unit modulation. **a**, Spike rasters of a single unit during six movements in each of eight directions. This unit (with $\{x,y,z\}$ components of its preferred direction, $PD = \{-0.52, 0.21, 0.47\}$) fired maximally in the backward-up-right direction (B,U,R) while retrieving from the lower left target, and fired least in the forward-down-left direction (F,D,L) while reaching to the same target. The modulation was consistent during (blue side bars) and after calibration (red side bars). **b**, Gripper modulation. Aperture command velocity (dotted line) and off-line predicted aperture velocity from neural data (solid line, ± 2 standard errors) during automatic gripper control, showing that the monkey's cortical population is modulated for observed gripper movement.

PVA, the extraction algorithm used, is dependent on accurate estimates of the recorded units' tuning properties. At the beginning of each day, the tuning properties were estimated in a calibration procedure that did not require the monkey to move its arm. Because motor cortical units modulate their firing rates when the subject watches automatic task performance²¹, the assisted task (the same as in the description of training above) was used for calibration. During the first iteration of four trials (one successful trial per target location), the monkey watched the automated performance of reach, grip and retrieval and then received the food. A trial was cancelled if the monkey did not appear to pay attention. Modulation evident during the first iteration was used to get an initial estimate of each unit's tuning properties (Supplementary Methods). During the next iteration, these initial estimates were used by the extraction algorithm to generate a signal that was mixed with the automated control. Tuning parameters were re-estimated at the end of each iteration while gradually decreasing the automated contribution until both arm movement and the gripper were fully controlled by the monkey's cortical activity. An example of the modulation during and after calibration is shown in Fig. 4a. In addition to endpoint movement, in the current study we also used observation-related activity for gripper control (Fig. 4b) in several phases of the training procedure, culminating in its skilled use (Fig. 2e and f).

With this study, we have expanded the capabilities of prosthetic devices through the use of observation-based training and closed-loop cortical control, allowing the use of this four-dimensional anthropomorphic arm in everyday tasks. These concepts can be incorporated into future designs of prostheses for dexterous movement.

METHODS SUMMARY

The timeline of each trial was divided into functional periods (Fig. 1b). A trial began with a piece of food being placed on the presentation device and the device moved to a location within the monkey's workspace to provide a reaching target (Presentation). The monkey often started moving the arm forward slowly before the presentation was complete. When the target was in place, the monkey started a directed reaching movement while simultaneously opening the gripper (Move A). Upon approach, the animal made small homing adjustments to get the endpoint aligned with the target (Home A), and then closed the gripper while actively stabilizing the endpoint position (Loading). If loading was successful, the monkey made a retrieval movement back towards the mouth while keeping the gripper closed (Move B), then made small adjustments to home in on the mouth (Home B) and stabilized the endpoint while using its mouth to unload the food from the gripper (Unloading). A trial was considered successful if the monkey managed to retrieve and eat the presented food. Each trial was followed by an inter-trial period while a new piece of food was prepared for presentation (Inter-trial). During continuous self-feeding, these task periods had no meaning during the execution of the task, but rather were imposed afterwards for purposes of data analysis. In contrast, during training and calibration, a real-time software module kept track of the task periods based on button-presses by a human operator and based on distance of arm endpoint from the tip of the food target presentation device. During training, this real-time delineation of task periods was used so that automated assistance could be applied differently during each task period depending on what aspect of the task the monkey was having difficulty with. During calibration, the delineation of task periods was used so that firing rates collected during each task period could be regressed against appropriate behavioural correlates. Further details on training and calibration are given in Supplementary Methods. Figures 2f and 3c, d are parallel-projection 3D plots. Figure 2g is in perspective-projection.

Received 14 November 2007; accepted 4 April 2008.

Published online 28 May 2008.

- Georgopoulos, A. P., Kalaska, J. F., Crutcher, M. D., Caminiti, R. & Massey, J. T. in *Dynamic Aspects of Neocortical Function* (eds Edelman, G. M., Gall, W. E. & Cowan, W. M.) 501–524 (Wiley & Sons, New York, 1984).
- Georgopoulos, A. P., Kettner, R. E. & Schwartz, A. B. Primate motor cortex and free arm movements to visual targets in three-dimensional space. II. Coding of the direction of movement by a neuronal population. *J. Neurosci.* **8**, 2928–2937 (1988).

- Schwartz, A. B. Direct cortical representation of drawing. *Science* **265**, 540–542 (1994).
- Schwartz, A. B. & Moran, D. W. Motor cortical activity during drawing movements: Population representation during lemniscate tracing. *J. Neurophysiol.* **82**, 2705–2718 (1999).
- Moran, D. W. & Schwartz, A. B. Motor cortical activity during drawing movements: Population representation during spiral tracing. *J. Neurophysiol.* **82**, 2693–2704 (1999).
- Moran, D. W. & Schwartz, A. B. Motor cortical representation of speed and direction during reaching. *J. Neurophysiol.* **82**, 2676–2692 (1999).
- Wessberg, J. *et al.* Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. *Nature* **408**, 361–365 (2000).
- Wolpaw, J. R., Birbaumer, N., McFarland, D. J., Pfurtscheller, G. & Vaughan, T. M. Brain-computer interfaces for communication and control. *Clin. Neurophysiol.* **113**, 767–791 (2002).
- Schwartz, A. B. Cortical neural prosthetics. *Annu. Rev. Neurosci.* **27**, 487–507 (2004).
- Leuthardt, E. C., Schalk, G., Moran, D. & Ojemann, J. G. The emerging world of motor neuroprosthetics: A neurosurgical perspective. *Neurosurgery* **59**, 1–14 (2006).
- Schwartz, A. B., Cui, X. T., Weber, D. J. & Moran, D. W. Brain-controlled interfaces: Movement restoration with neural prosthetics. *Neuron* **52**, 205–220 (2006).
- Kennedy, P. R. & Bakay, R. A. E. Restoration of neural output from a paralyzed patient by a direct brain connection. *Neuroreport* **9**, 1707–1711 (1998).
- Kennedy, P. R., Bakay, R. A., Moore, M. M., Adams, K. & Goldwaite, J. Direct control of a computer from the human central nervous system. *IEEE Trans. Rehabil. Eng.* **8**, 198–202 (2000).
- Serruya, M. D., Hatsopoulos, N. G., Paninski, L., Fellows, M. R. & Donoghue, J. P. Instant neural control of a movement signal. *Nature* **416**, 141–142 (2002).
- Taylor, D. M., Helms Tillery, S. I. & Schwartz, A. B. Direct cortical control of 3D neuroprosthetic devices. *Science* **296**, 1829–1832 (2002).
- Carmena, J. M. *et al.* Learning to control a brain-machine interface for reaching and grasping by primates. *PLoS Biol.* **1**, 193–208 (2003).
- Musallam, S., Corneil, B. D., Greger, B., Scherberger, H. & Andersen, R. A. Cognitive control signals for neural prosthetics. *Science* **305**, 258–262 (2004).
- Wolpaw, J. R. & McFarland, D. J. Control of a two-dimensional movement signal by a noninvasive brain-computer interface in humans. *Proc. Natl Acad. Sci. USA* **101**, 17849–17854 (2004).
- Lebedev, M. A. *et al.* Cortical ensemble adaptation to represent velocity of an artificial actuator controlled by a brain-machine interface. *J. Neurosci.* **25**, 4681–4693 (2005).
- Hochberg, L. R. *et al.* Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature* **442**, 164–171 (2006).
- Wahnoun, R., He, J. & Helms Tillery, S. I. Selection and parameterization of cortical neurons for neuroprosthetic control. *J. Neural Eng.* **3**, 162–171 (2006).
- Santhanam, G., Ryu, S. I., Yu, B. M., Afshar, A. & Shenoy, K. V. A high-performance brain-computer interface. *Nature* **442**, 195–198 (2006).
- Chapin, J. K., Moxon, K. A., Markowitz, R. S. & Nicolelis, M. A. L. Real-time control of a robot arm using simultaneously recorded neurons in the motor cortex. *Nature Neurosci.* **2**, 664–670 (1999).
- Helms Tillery, S. I., Taylor, D. M. & Schwartz, A. B. The general utility of a neuroprosthetic device under direct cortical control. *Proc. 25th Annu. Int. Conf. IEEE EMBS* **3**, 2043–2046 (2003).
- Brockwell, A. E., Rojas, A. L. & Kass, R. E. Recursive bayesian decoding of motor cortical signals by particle filtering. *J. Neurophysiol.* **91**, 1899–1907 (2004).
- Sanchez, J. C., Erdogmus, D., Nicolelis, M. A. L., Wessberg, J. & Principe, J. C. Interpreting spatial and temporal neural activity through a recurrent neural network brain-machine interface. *IEEE Trans. Neural Syst. Rehabil. Eng.* **13**, 213–219 (2005).
- Yu, B. M. *et al.* Mixture of trajectory models for neural decoding of goal-directed movements. *J. Neurophysiol.* **97**, 3763–3780 (2007).
- Schwartz, A. B., Taylor, D. M. & Helms Tillery, S. I. Extraction algorithms for cortical control of arm prosthetics. *Curr. Opin. Neurobiol.* **11**, 701–707 (2001).
- Morasso, P. Spatial control of arm movements. *Exp. Brain Res.* **42**, 223–227 (1981).
- Soechting, J. F. Effect of target size on spatial and temporal characteristics of a pointing movement in man. *Exp. Brain Res.* **54**, 121–132 (1984).

Supplementary Information is linked to the online version of the paper at www.nature.com/nature.

Acknowledgements We thank S. Clanton and M. Wu for help with software and hardware development, S. Chase for discussions and I. Albrecht for the illustration in Fig. 1a. Support contributed by NIH-NINDS-N01-2-2346 and NIH NS050256.

Author Information Reprints and permissions information is available at www.nature.com/reprints. The authors declare competing financial interests: details accompany the full-text HTML version of the paper at www.nature.com/nature. Correspondence and requests for materials should be addressed to A.B.S. (abs21@pitt.edu).