

Supplementary Methods

Progression of Experiments. After a first round of experiments with monkey P, a set of necessary improvements to the paradigm were identified. These improvements were implemented, and a second round of experiments performed with monkey A. The improved experiments could not be performed with monkey P because recordings from the cortical implant had faded by the time of the second round of experiments. The improvements were:

- 1) Replacement of the robotic arm to improve mechanical and control properties
- 2) Introduction of the presentation device to:
 - a. record the target location, and
 - b. remove the tendency of the human presenter to help the loading by moving their hand to meet the gripper.
- 3) Implementation of direct cortical gripper control
- 4) Improvement of the *assisted* control paradigm to enable calibration and training for gripper control

Extraction Algorithm. The algorithm used to extract an arm control signal from the real-time stream of neural data was a version of the population vector algorithm (PVA)¹⁻³. Given a population of N units ($i = \{1, 2, 3, \dots, N\}$, each being either a single unit or multi-unit cluster) that fired $c_i[n]$ spikes during discrete time step n (i.e. between time t_{n-1} and

time t_n , where $\Delta t = t_n - t_{n-1} = 30 \text{ ms}$), an “instantaneous” firing rate, $f_i[n]$, for each unit was calculated:

$$f_i[n] = \frac{c_i[n]}{\Delta t}. \quad (\text{eq. 1})$$

The firing rate was smoothed using a finite impulse response (FIR) filter, $h[k]$:

$$s_i[n] = \sum_{k=0}^{W-1} f_i[n-k]h[k], \quad (\text{eq. 2})$$

where W was the number of filter coefficients (see the Robotic Arms and Control Software section below for actual values of used coefficients). The smoothed firing rate, $s_i[n]$, was normalized using each unit’s baseline rate, b_i , and modulation depth, m_i :

$$r_i[n] = \frac{s_i[n] - b_i}{m_i}. \quad (\text{eq. 3})$$

The population vector, $\vec{u}[n]$, was obtained as the vector sum of preferred directions, $\vec{p}_i$, weighted by the normalized firing rates, $r_i[n]$:

$$\vec{u}[n] = \frac{N_D}{N} \sum_{i=1}^N r_i[n] \vec{p}_i, \quad (\text{eq. 4})$$

where N_D was the number of dimensions in $\vec{p}_i$ and $\vec{u}[n]$. Scaling by $1/N$, in eq. 4, kept the population vector in a normalized range and scaling by N_D kept its magnitude from decreasing as N_D was increased over the course of training (see Monkey Training below). In the ultimate self-feeding task, the preferred direction vectors and the population vector had components for each dimension that was being extracted, i.e. $\vec{p}_i = \{p_{xi}, p_{yi}, p_{zi}, p_{gi}\}$,

$\vec{u}[n] = \{u_x[n], u_y[n], u_z[n], u_g[n]\}$ and $N_D = 4$. The first three components of $\vec{u}[n]$ were interpreted as endpoint velocity,

$$\vec{v}_e[n] = k_{es} \vec{u}_{xyz}[n] + \vec{k}_{ed}, \quad (\text{eq. 5})$$

where $\vec{u}_{xyz}[n] = \{u_x[n], u_y[n], u_z[n]\}$, k_{es} was a speed constant to convert the endpoint components of the population vector from a normalized range to a physical velocity (typically, 100 to 250 mm/s) and $\vec{k}_{ed}$ was a constant drift correction term (typically, the x-component was -15 to -40 mm/s, and the other components zero). For monkey A, on some days, the magnitude of $\vec{v}_e[n]$ was scaled by a piecewise polynomial non-linear speed gain function to allow faster reaching while not sacrificing stability at low speeds (Supplementary Fig. 1). The last component of $\vec{u}[n]$ was interpreted as the velocity of the gripper aperture,

$$v_g[n] = k_{gs} u_g[n] + k_{gd}, \quad (\text{eq. 6})$$

where k_{gs} was a speed constant to convert the gripper component of the population vector from a normalized range to a suitable command value (typically, 4 to 6 s⁻¹), and k_{gd} was a constant drift correction term (typically 0.5 to 0.7 s⁻¹). The extracted velocities were integrated to obtain command position for endpoint,

$$\vec{p}_e[n] = \vec{p}_e[n-1] + \vec{v}_e[n] \Delta t, \quad (\text{eq. 7})$$

and gripper aperture (a unitless quantity where 0 means fully closed and 1 means fully open),

$$a_g[n] = a_g[n-1] + v_g[n] \Delta t. \quad (\text{eq. 8})$$

$\bar{p}_e[n]$ and $a_g[n]$ were sent as commands to the robot control software (except when gripper control had not yet been implemented, $\bar{p}_e[n]$ was sent without $a_g[n]$). $\bar{p}_e[n]$ was interpreted as a vector in an arbitrary cartesian coordinate system with an origin and orientation that were fixed relative to the robot arm's base. The monkey was positioned next to the arm so that its mouth was at the coordinate system origin when the head was pointing directly forward.

The drift correction terms in equations 5 and 6 were necessary because an offset in endpoint velocity and gripper aperture velocity is caused by estimation error in baseline firing rate parameters, b_i . The estimation error is due to asymmetry in the task (monkey is more motivated on retrieval movements than reaching movements), deviation from the cosine tuning model that is implicitly assumed by the calibration model below (actual firing rates do not modulate equally above and below baseline rate) and, noise in firing rates.

Calibration Model: As can be seen from the equations above, the extraction algorithm relied on the parameters b_i (baseline firing rate), m_i (modulation depth) and $\bar{p}_i$ (preferred direction), collectively called *tuning parameters*, that describe how each unit modulates its firing rate with arm velocity. These parameters had to be calibrated before prosthetic control could be performed. In this section, Greek symbols are used (as opposed to Latin ones in the previous section), to help distinguish the variables used in parameter calibration from variables used in read-time command extraction. For example, φ in parameter estimation refers to an average firing rate over a whole movement period as defined below, but f in real-time extraction refers to an “instantaneous” firing rate. Previous work^{4,5} has shown that the firing rate, φ , of a unit in the proximal arm area of

the primary motor cortex during natural reaching in 3D space can be approximated by the model,

$$\varphi = \beta_0 + \beta_x v_x + \beta_y v_y + \beta_z v_z, \quad (\text{eq. 9})$$

where $\vec{v} = \{v_x, v_y, v_z\}$ is arm endpoint velocity, β_0 is the baseline rate and

$\vec{\beta} = \{\beta_x, \beta_y, \beta_z\}$ is a vector in the unit's preferred direction with the modulation depth as its magnitude. This equation has a form suitable for linear regression, allowing the tuning parameters to be estimated easily using data collected during natural arm movement.

However, if this technology is to be used for paralyzed persons or amputees, natural arm movement cannot be used. Furthermore, previous work³ shows that tuning parameters estimated from natural arm movement are not optimal for brain-controlled movement. In this study, a variation of eq. 9 was used, that worked without natural arm movement and also added a component for the gripper:

$$\varphi_{ij} = \beta_{0i} + \beta_{xi} \delta_{xj} + \beta_{yi} \delta_{yj} + \beta_{zi} \delta_{zj} + \beta_{gi} \delta_{gj}. \quad (\text{eq. 10})$$

This model had the same form as eq. 9, but the key distinction was that rather than using the velocity of the natural arm ($\vec{v}$ of eq. 9), it used a target displacement vector

$\vec{\delta}_j = \{\delta_{xj}, \delta_{yj}, \delta_{zj}, \delta_{gj}\}$, that represented the normalized displacement from the prosthetic arm's initial state, $\vec{\zeta}_{0j} = \{\zeta_{0xj}, \zeta_{0yj}, \zeta_{0zj}, \zeta_{0gj}\}$, to the target state, $\vec{\zeta}_{Tj} = \{\zeta_{Txj}, \zeta_{Tyj}, \zeta_{Tzj}, \zeta_{Tgj}\}$, during the j -th segment of movement:

$$\vec{\delta}_j = \left\{ \frac{\Delta_{xj}}{D}, \frac{\Delta_{yj}}{D}, \frac{\Delta_{zj}}{D}, \Delta_{gj} \right\}, \text{ where} \quad (\text{eq. 11})$$

$$\vec{\Delta}_j = \{\Delta_{xj}, \Delta_{yj}, \Delta_{zj}, \Delta_{gj}\} = \vec{\zeta}_{Tj} - \vec{\zeta}_{0j}, \text{ and} \quad (\text{eq. 12})$$

D was a normalization constant to rescale the magnitude of the x , y and z components so that all components would have a normalized range of roughly -1 to 1 (the gripper component was in that range without rescaling). The value of D was chosen arbitrarily as 220.3 mm (a fixed value representing the approximate distance from mouth at which targets were presented). In monkey P's experiments $\vec{\delta}_j$ was defined as a unit vector in the same direction as $\vec{\Delta}_j$. φ_{ij} was the firing rate of unit i , averaged over the j -th segment of movement. Firing rates, φ_{ij} , and target movement vectors, $\vec{\delta}_j$, collected over a number of movement segments over a number of trials (see Calibration Procedure sections below), were input into multiple linear least-squares regression to estimate the b -coefficients of eq. 10 (the regression is performed independently for each unit, i). Finally, the tuning parameters used by the extraction algorithm were obtained from the b -coefficients:

$$b_i = \beta_{0i}, \quad (\text{eq. 13})$$

$$m_i = |\vec{\beta}_i|, \text{ where } \vec{\beta}_i = \{\beta_{xi}, \beta_{yi}, \beta_{zi}, \beta_{gi}\}, \text{ and} \quad (\text{eq. 14})$$

$$\vec{p}_i = \frac{\vec{\beta}_i}{|\vec{\beta}_i|}. \quad (\text{eq. 15})$$

Calibration Procedure A (final version): To calibrate the tuning parameters, an iterative process was used, where initial estimates were based on observation-related activity (Wahnoun et al. ⁶ and Fig. 4). The monkey watched the arm automatically perform 4 successful trials consisting of reaching, loading and retrieval of a food piece from each of 4 locations in random order (lower-left, lower-right, upper-left and upper-

right). The mean firing rate, ϕ_{ij} , for each unit, i , and values for $\bar{\zeta}_{0j}$ and $\bar{\zeta}_{Tj}$, were collected for each segment of movement, $j = \{1, 2, 3, \dots, J\}$, where J refers to the number of movement segments collected per *iteration* of the calibration procedure (there is one *iteration* per *repetition* of the task, see the Assisted Control Paradigm section). The initial arm state, $\bar{\zeta}_{0j}$, was defined as the actual arm state at the beginning of the j -th movement segment. Target arm state, $\bar{\zeta}_{Tj}$, had a pre-defined value for each task period (Supplementary Table 1). There were 6 segments per trial, one per task period (*Move A*, *Home A*, *Loading*, *Move B*, *Home B*, and *Unloading*), i.e. $J = 6 \times 4 = 24$ for 4 trials. Tuning parameters were estimated from the collected data as described by equations 11-16. These initial parameters were then used by the EM to provide the monkey with partial control during the next iteration. During each iteration, another 4 trials worth of data were collected, with one successful movement cycle to each of the four locations (data from unsuccessful trials was not used). At the end of each iteration k , the cumulative data set, $j = \{1, 2, 3, \dots, Jk\}$, was used to refine the tuning parameter estimates, the EM was updated with the new parameters to provide the monkey with better control, and the proportion of automated control decreased. Units with a modulation depth less than a cutoff, $m_i < M_c$, and units with their r^2 -value (from regression) less than 0.1, were excluded from the population vector (M_c was typically 4Hz). A total of 4 iterations of the calibration procedure were typically performed at the beginning of a daily session, and the final estimated tuning parameters were used by the EM for the remainder of the day.

Calibration Procedure B (initial version): A different calibration procedure was used when gripper control had not yet been implemented in the extraction algorithm. This

procedure was used for all of monkey P's experiments, and for the first phase of monkey A's experiments. The key difference was that in procedure A, data for estimating tuning parameters were collected during (at least partially) automated movement of the arm, but in procedure B, the movement was not automated while the data were collected. The fact, that successful calibration was achieved with procedure B, meant that observation of successful movement was not required for the subject to produce directionally modulated activity (Supplementary Fig. 2a). In procedure B, the initial tuning parameters were set to arbitrary initial values ($\vec{p}_i$ was random, b_i was set to 10 Hz, and m_i to 50 Hz for all units). Data for tuning parameter estimation were collected during *Move A* and *Move B* while the arm was controlled completely by the EM's output. Because of the arbitrary initial settings of the tuning parameters, during the first iteration, the movement velocity was unrelated to the animal's intention, but directionally modulated activity was assumed to be present. *Move A* and *Move B* ended after a brief timeout (0.5-1 s), and proceeded to *Home A* or *Home B*, respectively. Arm movement during *Home A*, *Loading*, *Home B* and *Unloading* was completely automated. A trial was labelled "successful", as long as the animal appeared to be paying attention during both *Move A* and *Move B*. At the end of each iteration of the procedure, tuning parameters were estimated based on equations 10-15 (with the gripper component removed from each respective vector), and the new parameters were applied to real-time control during the next iteration. After each iteration of calibration, the monkey's control improved and the preferred directions quickly converged on their final values (Supplementary Fig. 2b). Units with a modulation depth less than a cutoff, $m_i < M_c$, were excluded from the population vector (M_c was typically 4Hz). Unlike in procedure A, an r^2 cutoff was not used here. Each *iteration* consisted of

3 successful trials where each trial was a full movement cycle to one of 4 fixed target locations (the *iterations* in this version of the calibration procedure were not aligned with *repetitions* of the task, see the Assisted Control Paradigm section). Food target presentation location, $\vec{\tau}_j$, in this version of the procedure, was defined as the nominal location (because the actual location was not known, since the presentation device had not yet been implemented).

Neurophysiological Recordings. Intracortical microelectrodes were implanted in the proximal arm region of the primary motor cortex. Spike signals were acquired using a 96-channel Plexon MAP system (Plexon Inc., Dallas, TX, USA). Monkey P had 4 microwire arrays in each hemisphere. The arrays consisted of 16 teflon-coated tungsten wires, each with a diameter of $50\mu m$, arranged in a 2×8 grid with $300\mu m$ spacing. All 64 channels from the right hemisphere and 32 of the 64 from the left were connected for recording at any one time. Monkey A was implanted with a Utah array (Cyberkinetics, Inc., Foxborough, MA, U.S.A.) in the right hemisphere, consisting of a 10×10 grid of electrodes with $400\mu m$ spacing and a shank length of $1.5mm$. Out of the 100 electrodes on the Utah array, 96 were wired for recording and the remaining 4 were unconnected. The number of units typically isolated each day was 20-50 for monkey P (mostly from the right hemisphere, the left hemisphere typically yielded only a few or no channels with spiking activity that could be isolated). Of the 20-50 isolated units for monkey P, 10-30 were typically used for control. For monkey A, 150-180 units were isolated from the right hemisphere with 60-120 used for control. Spikes were sorted using the box-sorting and

PCA methods in Plexon's SortClient software (a part of their Rasputin package). Most of the sorted units were multi-unit clusters and some were single units.

Assisted Control Paradigm. During training and calibration, a paradigm of *assisted* control was used, whereby automated control was mixed with the monkey's cortical control. During training, the purpose of this assistance was to shape the monkey's behaviour by operant conditioning. By gradually decreasing automated assistance over a series of days, it was possible to keep each task period at an appropriate level of difficulty, so that the monkey would stay motivated and improve. During calibration, the purpose of the assistance was to provide a behavioural template for the monkey to produce modulated neural activity. To apply the correct type and amount of assistance at each task period (Fig. 1b), the robot control software kept track of task periods in real time (Supplementary Table 2). At each 30 ms time step when the Robot Control Module (RCM, see below) received a command from the Extraction Module (EM, see above), it applied three types of assistance (Supplementary Fig. 3) that were combined in configurable proportions per task state. Assistance was applied separately to $\vec{p}_e[n]$, the movement component of the monkey brain control command, and $a_g[n]$, the gripper component. Deviation gain was applied to $\vec{p}_e[n]$, resulting in $\vec{p}_{dg}[n]$ whereby movement perpendicular to the target direction was weighted by a "deviation gain" between 0 and 1. Target direction was defined as the instantaneous direction from the current endpoint position to the current target (the target was at the mouth for retrieval moves). This results in partially assisted 3D control where it was more difficult to go in the wrong direction than in the correct direction. A deviation gain of 0 would result in the endpoint

being able to move only along a line between the mouth and the target whereas a deviation gain of 1 would result in full 3D control.

Attraction assistance was applied to obtain the effect of “attracting” the endpoint toward target by mixing $\vec{p}_{dg}[n]$ with a vector toward the target $\vec{\tau}[n]$ (eqs. 16-18).

$$p_{bc}[n] = p_{final}[n-1] + (p_{dg}[n] - p_{final}[n-1]) * MovementBCGain \quad (\text{eq. 16})$$

$$AttractionVector[n] = \frac{(\vec{\tau}[n] - p_{bc}[n])}{|\vec{\tau}[n] - p_{bc}[n]|} * AttractionSpeed * \Delta t \quad (\text{eq. 17})$$

$$p_{final}[n] = p_{bc}[n] + AttractionVector[n] * (1 - MovementBCGain) \quad (\text{eq. 18})$$

$\vec{p}_{final}[n]$ was used to move the robot arm. The values of *MovementBCGain* ranged from 0 (full automatic control) to 1 (full monkey control). *AttractionSpeed* was essentially a configurable constant, but as the monkey moved the arm closer to the target, *AttractionSpeed* became slower, to prevent overshooting the target (Supplementary Fig. 4).

Gripper assistance consisted of 2 steps: calculating an assisted gripper command $ag[n]$, and combining it with the extraction module’s gripper command $a_g[n]$ from eq. 8.

$$ag[n] = ag[n-1] + GripperAssistVelocity * \Delta t \quad (\text{eq. 19})$$

$$a_{final}[n] = a_g[n] * GripperBCGain + ag[n] * (1 - GripperBCGain) \quad (\text{eq. 20})$$

GripperAssistVelocity had a magnitude of *GripperAssistSpeed* and a sign determined by the desired action for the gripper (opening or closing). *GripperAssistSpeed* was usually constant within each session based on pre-set configuration, with a typical value of 3 s^{-1} .

$a_{final}[n]$ was used to control the gripper.

Targets in the *assisted* training and calibration tasks were presented at 4 fixed positions for the 3D and 4D task (Fig. 3c and d), 2 fixed positions for the 2D training phase, and one fixed position for the 1D training phase (see below). At the beginning of each trial, one of the fixed targets was chosen by the behavioural control software and displayed on a screen visible to the trainer, but not to the monkey (the screen was behind the monkey and slightly to the monkey's left). The trainer then presented the food at the approximate corresponding spatial location in the workspace. The fact, that the actual presented location did not exactly match the ideal target location, did not matter because the only purpose of using the computer program to select targets, was to keep a balanced distribution of presentations at the categorical locations (lower-left, lower-right, upper-left or upper-right). As in previous work from our group³, trials were grouped into *repetitions*, so that at the beginning of each repetition, all targets were placed on a "to be attempted" list. For each trial, a target was randomly chosen from the list. If the trial was successful, the target was removed from the list. If the trial was unsuccessful, it remained on the list. At the beginning of the next trial, a new target was randomly chosen from the list. After three unsuccessful attempts at a given target, the target was removed from the list to keep the monkey motivated. When no targets remained on the "to be attempted" list, the repetition was over and a new one began. For example, during 4D or 3D control, this meant that a repetition could consist of 4 to 12 trials (4 if each target successful on first try, 12 if three attempts made at each target, or some number in between). This procedure helped to make sure that there would be one successful trial per target location per repetition, to keep the data balanced during a Calibration Procedure. It also ensured that if there was a particular target that the monkey consistently failed on during training,

then that target would end up being presented more times than the other targets, giving the monkey more practice at that target. (During monkey P's experiments, the presentation device had not yet been implemented, so the exact location of presentation was not known. However, the human presenter learned to be accurate because in the *assisted* task, the endpoint was automatically homed in on the ideal target location, giving the presenter feedback on the accuracy of their presentation.)

Continuous Self-feeding. This was the ultimate task paradigm where the arm was completely controlled by the monkey without any automated assistance, continuously throughout the session, without a break even during the *Inter-trial* state. Automated assistance (see above) was disabled during continuous self-feeding. The purpose of this paradigm was to demonstrate the feasibility of using a cortically controlled prosthetic arm for actual feeding. This task was difficult because of the positional accuracy required for successful loading (Supplementary Fig. 5). The real-time robot control software did not keep track of task periods during this paradigm because the task was unstructured – anything was allowed, including multiple approaches to the target, multiple attempts at loading etc, making it difficult to fit within a pre-defined series of task periods. As such, the task periods in Figure 1 are for descriptive purposes only and their durations were determined as approximate values based on Supplementary Video 1. The success rates reported for this paradigm were determined from a record of button presses by an observer or, in some cases, from a video record.

Monkey Training. Control of the prosthetic arm relied on modulation of neuronal firing rates with desired movement direction. Unlike human patients, who could be verbally instructed on how to use the arm, monkeys needed an extra initial training step to condition them to use the arm. They were initially trained to control the arm using a one-dimensional (1D) joystick, the movement of which was mapped to forward-backward displacement of the arm's endpoint. (For monkey A, in addition to controlling endpoint movement, a pressure sensor in the joystick handle was mapped to gripper closing.) This training step was equivalent to the Hand-Control training in previous experiments³ where the hand was optically tracked to provide cursor control. In this study, the joystick was used in place of tracking the monkey's arm, because it was easier to implement. Within a few days, the monkeys learned to make successful reaching and retrieval movements to feed themselves with the robot arm and gripper under joystick control, after which they graduated to cortical control.

3D cortical control was attempted immediately following joystick control, but this was unsuccessful, presumably because the change in task difficulty was too sudden. Therefore, intermediate training stages were created, using the *assisted* task paradigm:

- 1) *1D Cortical Control with Attraction Assist*: movement of the endpoint is restricted by the control algorithm to the *x*-dimension (depth) with the *y*- and *z*-components set to zero, allowing the robot's hand to move straight ahead away from the mouth and back. Targets are presented directly ahead of the monkey's mouth. This task is the same as the 1D joystick task, except the robotic arm is controlled by cortical activity and the monkey's arm is restrained.

- 2) *2D Cortical Control with Attraction Assist*: movement of the endpoint is restricted to the vertical *xy*-plane or the horizontal *xz*-plane. Targets are presented in the plane.
- 3) *3D Cortical Control with Deviation Gain and Attraction Assist*: movement assistance provided as described in the Assisted Control Paradigm section above, gripper control automated.
- 4) *4D Cortical Control with Attraction Assist*: movement and gripper assistance provided as described in the Assisted Control Paradigm section above.

Monkey P trained on assisted cortical control for 3 days using the 1D task, 11 days using the 2D task and 19 days using the 3D task. After the Deviation Gain reached 1, i.e. full 3D control during *Move A* and *Move B*, monkey P continued to train on 3D *assisted* control with *Attraction Assist* in *Home A*, *Loading*, *Home B* and *Unloading* for 30 days before performing continuous self-feeding. During this final stage of *assisted* training, the success rate increased over the first 10 days and stayed at a consistent high level until day 17 while the homing radius was fixed at 50 *mm* (Supplementary Fig. 6). Thereafter the homing radius was gradually decreased to as low as 10 *mm* to train the monkey to make more precise reaches. The success rate (fraction of successful trials out of attempted trials) fluctuated as the homing radius was changed, but a measure of performance, calculated as

$$Performance = \frac{Success\ Rate}{Homing\ Radius}, \quad (\text{eq. 21})$$

shows an upward trend throughout. This measure takes into account the increasing difficulty as the homing radius is decreased. After this training, monkey P performed the 3D continuous self-feeding task.

Monkey A trained for 2 days on 1D cortical control, skipped the 2D training, trained for 8 days on 3D control with *Deviation Gain* and *Attraction Assist*. After *Deviation Gain* reached 1, the monkey continued to train on the 3D task where the monkey had full control during *Move A* and *Move B*, but *Attraction Assist* was applied during *Home A*, *Loading*, *Home B* and *Unloading*. The radius of the homing regions was decreased over the next 6 days and then the monkey performed continuous 3D self-feeding for 7 days. After cortical gripper control was implemented, monkey A trained on 4D cortical control with *Attraction Assist* for 36 days while the amount of movement assistance was decreased until movement was fully controlled by the monkey, and then gripper assistance was decreased until fully controlled by the monkey. After this training, the monkey performed 2 days of the 4D continuous self-feeding task.

Robotic Arms and Control Software. There were two separate arms used over the course of the experiments. Monkey P used a custom-made anthropomorphic arm from Keshen Prosthetics (Shanghai, China). Monkey A used a standard WAM arm with a shortened upper-arm link from Barrett Technology Inc. (Cambridge, MA, U.S.A.). The two arms were functionally equivalent for the purpose of this study. The general shape and degrees of freedom of both arms resembled that of a human arm. Both arms used DC motors embedded in the arm to actuate four axes: shoulder flexion, shoulder abduction, shoulder rotation and elbow flexion. The motors were servo-controlled in joint angular position PID mode using feedback from optical encoders (a National Instruments FW-7344 controller was used for the Keshen arm, whereas the Barrett arm came with its own Linux-based controller). Command updates were sent to the controller from a computer

system every 30 ms. These command updates were computed from the monkey's cortical activity (see "Extraction Algorithm" above) in the form of a Cartesian endpoint position, which was converted to joint angular positions by the Robot Control Software (see below). The Keshen arm was replaced with the Barrett arm for better mechanical stability. The gear-driven mechanism of the Keshen arm was subject to play between the gears, resulting in free movement of the joints, even when the motor was not moving. This resulted in undesirable oscillatory deviations from the command position. The Barrett arm, on the other hand, is cable-driven, resulting in minimal play between the motors and the output shaft. The Barrett arm was able to follow the command position accurately (Supplementary Fig. 7). The Barrett arm's maximal speed at the endpoint was 2000 mm/s (WAM Arm's User Guide).

In addition to the four proximal joints, each arm was fitted with a motorized two-fingered gripper with a custom-made controller. The two fingers were mechanically linked so that a single motor moved both simultaneously, providing a single DOF of finger aperture control. Thus, the total DOF of the robotic system was 5 (4 for the arm and 1 for the gripper), but only 4 DOF were independently controlled using cortical signals (3 for the arm endpoint position and 1 for the gripper).

Since the control signal for the arms was based on inherently noisy instantaneous firing rates, a smoothing filter (eq. 2) was used in the extraction algorithm to produce a reasonably smooth control signal. The filter coefficients that were used were changed from time to time for each arm: a 5-sample filter, $h[k] = [.2, .2, .2, .2, .2]$ was typically used for monkey A and an 11-sample filter

$h[k] = [.013, .039, .078, .123, .159, .173, .159, .123, .078, .039, .013]$ for monkey P. The 5-

sample filter was originally used with monkey P, but was switched to the 11-sample filter to achieve smoother movements. The main results from monkey P (i.e. continuous self-feeding) are from sessions where the 11-sample filter was used. Monkey A initially used the 11-sample filter, but was switched to the 5-sample filter to reduce control delay. The main results from monkey A (i.e. continuous self-feeding) were from sessions where the 5-sample filter was used. For monkey A, the 5-sample filter provided sufficient smoothing because the population vector was less noisy due to the much higher number of recorded units compared to monkey P (see Neurophysiological Recordings).

Control Delay. An important characteristic of a real-time control system is the delay between input and output, i.e. how long does it take before a change in the input signal is reflected in the output. The control delay can be sub-divided into system delay and memory delay. System delay is how long it takes to acquire a sample of the input signal, compute the output, and effect the output. Memory delay is a result of memory states in the control algorithm (i.e. the smoothing filters in the EM and the robot controller). System delay was ~ 60 ms, consisting of spike counting delay (15 ms), software system delay (~ 15 ms, measured using pulses at input and output that were timed by hardware) and mechanical delay (~ 30 ms). Memory delay was 90 ms, consisting of EM filtering delay (60 ms for the 5-sample filter) and WAM command filtering delay (30 ms). Therefore the total control delay was $\sim 60 + 90 = \sim 150$ ms.

Food Presentation Device. In order to get accurate measurements of the food target location in 3D space, food targets were placed on the tip of a rigid device that had infra-

red emitting optical tracking markers on it. The markers were tracked using an Optotrak 3020 system (Northern Digital Inc, Waterloo, Ontario, Canada). The tip location was calculated from the marker locations using trigonometry.

Robot Control Software. The robot control module (RCM), a custom software module in charge of communicating with the robot controller, received a command from the Extraction Module (EM, see Extraction Algorithm above) every 30 *ms*. The RCM served the following functions:

- 1) Apply automatic assistance and mix it with the cortical command as described in the Assisted Control Paradigm section above;
- 2) During training, allow the human operator to override gripper control using button presses on a control pad;
- 3) Apply workspace limits (see below);
- 4) Calculate joint angular command for the robot controller from the endpoint command.

Workspace limits for monkey P were -1 mm (backward) to 201 mm (forward) for the *x*-dimension, -81 mm (lower) to 71 mm (upper) for the *y*-dimension, and -81 mm (left) to 71 mm (right) for the *z*-dimension. For monkey A, the limits were -20 mm (backward) to 210 mm (forward) for the *x*-dimension, -150 mm (lower) to 210 mm (upper) for the *y*-dimension, and -150 mm (left) to 150 mm (right) for the *z*-dimension.

Joint angles were calculated using an inverse kinematics algorithm. There were 3 degrees of freedom (DOF) to the Cartesian endpoint position (*x*, *y* and *z*) while the robot arm had 4 DOF (angular position of each joint). In order to constrain the extra DOF, the

concept of swivel angle⁷ was used. Swivel angle specifies how high the elbow is raised, defined as the angle between two planes: 1) the plane passing through the arm's endpoint, shoulder and elbow, and 2) the vertical plane passing through the endpoint and the shoulder. Kang et al.⁸ have described an algorithm that uses an energy minimization approach for finding a swivel angle, resulting in natural arm movements. For computational simplicity, and based on the observation that swivel angles calculated using the Kang et al. algorithm did not vary much within our limited workspace, we used a version of inverse kinematics with the swivel angle set to a constant 30 degrees, resulting in fairly natural-looking arm movements.

As a special case in a limited number of trials, for the continuous self-feeding task by monkey P, gripper control was implemented as a dependent degree of freedom controlled by the endpoint movement command signal based on a displacement threshold. The idea was to open the gripper when it moved forward to prepare it for gripping a target, and to close it whenever it was stabilized (designed on the assumption that the subject would stabilize it at the target). Whenever the total *x*-displacement (i.e. forward movement) within the latest 600ms exceeded 50mm, the gripper was opened. Whenever the path length in the *x*-dimension within the last 600ms was below 20mm, the gripper was closed. An additional closing criterion was based on backward movement, so that the gripper would close when the subject retracted the arm back toward its body without having loaded anything into the gripper. Whenever the *x*-displacement within the last 600ms exceeded – 20mm, the gripper was closed. This gripper control algorithm was not used for monkey A, because monkey A used its cortical activity to control the gripper directly as an independent 4-th dimension.

Supplementary Data

To give further details on Monkey A's performance, the success rates were broken down by task state and food type (Supplementary Table 3). Monkey A used 116 out of 185 sorted units on the first day of continuous self-feeding and 94 out of 175 on the second day. The decreased number of units from the first to the second day may have been a possible reason for the drop in success rate from 66% to 58%. On the first day, the non-linear speed gain was used, but on the second day it was not (see Extraction Algorithm above).

Monkey P had far fewer units (13-24 used out of 21-45 isolated) than monkey A, and yet was able to perform well on an easier self-feeding task. This task differed from that of monkey A in two respects: 1) the gripper, instead of being controlled directly by cortical activity, was controlled by cortical activity indirectly by virtue of being linked to the cortically controlled movement of the arm's endpoint (see Robot Control Software above); 2) the food was presented by hand instead of being placed on a presentation device (there is a natural tendency for the person presenting the food to make the loading easier by helping to position the fruit accurately between the gripper fingers, thus eliminating the difficult *Home A* and *Loading* periods). Monkey P performed 1-3 continuous control sessions per day, with an average success rate of 78% over a total of 1064 trials over 13 days. The highest daily success rate was 93% on day 2. On day 9, the monkey achieved 36 successful trials in a row. The number of units isolated and the number of units used for control for monkey P also varied over days (Supplementary Fig. 8a) but there was no clear correlation between the number of used units and success rate (Supplementary Fig. 8b). Monkey P's distribution of preferred directions (PD-s) of units used for robot control was non-uniform, yet good control was achieved. The distribution

varied from day to day but there was usually at least one half of space that was sparsely populated (e.g. in the distribution in Supplementary Figure 8c, all but two PD-s are in the lower half of space). Theoretical studies have shown that non-uniformity of the PD distribution decreases the prediction accuracy of PVA², but it was not known until now how much of an effect it has in practice. The finding that good control can still be achieved is important for applicability of this technology to human prosthetic use, because the likelihood of getting a uniform distribution is low in practice. Movement consistency during the 4-fixed-target training task for monkey P (Supplementary Fig. 9) is comparable to that of monkey A (Fig. 3c and d).

When the extraction algorithm was extended to include gripper control, a design choice was made to treat gripper aperture velocity as a fourth dimension in the model driven by all units in the population. An alternative would have been to build a separate one-dimensional model driven by a subset of units. The choice to include gripper as a fourth dimension in a single model was predicated on the hypothesis that units would exhibit both endpoint-tuning and gripper-tuning in different relative amounts per unit. This hypothesis was found to be true (Fig. 2a and g) and, as a result, gripper control was independent of endpoint control (Supplementary Fig. 10). Variation in relative amounts of gripper vs. endpoint modulation could have been partly due to the fact that the 10×10 electrode array with roughly a 4×4mm spatial span could have covered a range of more proximal-arm related and more distal-arm related parts of the motor map. However, units that were primarily gripper-related were found on the same electrode sites as units that were primarily endpoint-related (Supplementary Fig. 11).

To show that monkey A was able to reach the targets without help, it was important to quantify the amount and direction of movement of the target. During continuous self-feeding with monkey P, the human trainer sometimes (subconsciously)

moved the target slightly to meet the gripper (this can be seen in Supplementary Video 2). In the second set of experiments with monkey A, targets were presented using a presentation device rather than by hand, and the presentation device was optically tracked so that the location of its tip (where the food target was placed) was known. During a typical trial, the target was moved quickly to a presentation position, stayed steady during *Move A* and *Home A*, moved around slightly during *Loading* due to interaction with the arm, and was then moved away. To show that the target was not moved toward the arm endpoint to help the monkey, target movement across all trials in one of the continuous self-feeding sessions was quantified based on the following procedure: 1) The continuous data record was manually segmented into trials based on a synchronized video record; 2) A subset of trials was chosen where task periods up to and including *Loading* were successful; 3) For each trial, a time period (*Target Test* period) corresponding roughly to *Move A* and *Home A* was identified in the data (starting when the change in target position between consecutive 30 ms samples fell below 0.5 mm as it was moved into a stable presentation position, and ending when either minimal distance over the whole trial between target and arm endpoint was achieved or when the gripper first touched the target, whichever came first); 4) Target position increment vectors for each 30 ms sample during the *Target Test* period were projected onto a line passing through the target and arm endpoint (projections with a positive magnitude would point toward the arm endpoint, and projections with a negative magnitude would point away); 5) The projected increments were summed to obtain total displacement of target along the target-to-arm-endpoint line. Total displacement toward target was 1.9 ± 3.1 mm (mean $\pm$ std), i.e. not significantly helping the monkey.

Supplementary References

- ¹ A. P. Georgopoulos, R. E. Kettner, and A. B. Schwartz, "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 (1988).
- ² A. B. Schwartz, D. M. Taylor, and S. I. Helms-Tillery, "Extraction algorithms for cortical control of arm prosthetics," *Curr Opin Neurobiol* **11**, 701 (2001).
- ³ D. M. Taylor, S. I. Helms Tillery, and A. B. Schwartz, "Direct cortical control of 3D neuroprosthetic devices," *Science* **296**, 1829 (2002).
- ⁴ A. B. Schwartz, R. E. Kettner, and A. P. Georgopoulos, "Primate motor cortex and free arm movements to visual targets in three-dimensional space. I. Relations between single cell discharge and direction of movement," *J Neurosci* **8**, 2913 (1988).
- ⁵ G. A. Reina, D. W. Moran, and A. B. Schwartz, "On the relationship between joint angular velocity and motor cortical discharge during reaching," *J Neurophysiol* **85**(6), 2576 (2001).
- ⁶ R. Wahnoun, J. He, and S. I. Helms Tillery, "Selection and parameterization of cortical neurons for neuroprosthetic control," *J Neural Eng* **3**(2), 162 (2006).
- ⁷ D. Tolani and N. I. Badler, "Real-time inverse kinematics of the human arm," *Presence Teleoper Virtual Environ* **5**(4), 393 (1996).
- ⁸ T. Kang, J. P. He, and S. I. Helms Tillery, "Determining natural arm configuration along a reaching trajectory," *Exp Brain Res* **167**(3), 352 (2005).

Supplementary Tables

Task Period	$\vec{\zeta}_{Tej} =$	$\zeta_{Tgj} =$
Move A	$\vec{\tau}_j$	1
Home A	$\vec{\tau}_j$	1
Loading	$\vec{\tau}_j$	0
Move B	$\vec{\mu}$	0
Home B	$\vec{\mu}$	0
Unloading	$\vec{\mu}$	0

Supplementary Table 1. Pre-defined values of target arm state, $\vec{\zeta}_{Tj}$. $\vec{\zeta}_{Tej} = \{\zeta_{Txj}, \zeta_{Tyj}, \zeta_{Tzj}\}$ refers to the endpoint component of $\vec{\zeta}_{Tj}$. ζ_{Tgj} is the gripper component of $\vec{\zeta}_{Tj}$. $\vec{\tau}_j$ is the actual location of the presented food target (based on optical tracking of the presentation device) at the beginning of movement segment, j . $\vec{\mu}$ is the nominal location of the monkey’s mouth. A gripper value of 1 represents maximal aperture, and 0 represents a closed gripper.

Task Period	Condition for transitioning to the next period	Condition for failing a trial (and transitioning to Inter-trial period)
<i>Inter-trial</i>	“Continue” button pressed	None
<i>Presentation</i>	“Continue” button pressed	None
<i>Move A</i>	Endpoint position gets within “Homing Radius” of the target or “Continue” button pressed	Timed out or “Fail” button pressed
<i>Home A</i>	Endpoint gets within “Loading Radius” of the target	Timed out or “Fail” button pressed
<i>Loading</i>	Gripper command value gets below a “Closed Threshold”	Timed out, exited “Loading Radius” or “Fail” button pressed
<i>Move B</i>	Endpoint position gets within “Homing Radius” of the mouth or “Continue” button pressed	Timed out or “Fail” button pressed
<i>Home B</i>	Endpoint gets within “Unloading Radius” of the mouth	Timed out or “Fail” button pressed
<i>Unloading</i>	Time spent in <i>Unloading</i> exceeds a timeout value	“Fail” button pressed

Supplementary Table 2. Conditions for transitioning between task periods during the *assisted* task. “Continue” and “Fail” are buttons on a control pad operated by the trainer.


	All food-types combined	Marsh-mallow	Grape half	Marshmallow + Blueberry	Blueberry	Marshmallow + Grape half	Whole grape
No. of Presentations	330	187	69	65	6	2	1
No. of Attempts	298	171	63	56	5	2	1
No. of Successes	182	111	36	33	1	1	0
Complete Success %	61	65	57	59	20	50	0
Move A Success %	98	98	100	96	100	100	100
Home A Success %	89	87	98	84	100	100	100
Loading Success %	83	81	89	80	80	100	0
Move B Success %	76	78	73	73	80	100	0
Home B Success %	73	74	70	73	60	50	0

Supplementary Table 3. Monkey A’s success statistics combined over the two sessions of the continuous self-feeding task, broken down by food type and task period. The 61% *Complete Success* rate (bold) is the one reported in the main text as the overall success rate. It refers to the percentage of attempted trials where the monkey succeeded in getting the food into its mouth. The success rates listed per task period indicate the percentage of attempted trials where all periods at least up to and including that period were successful. Columns where two food types are listed together (e.g. Marshmallow + Blueberry) indicate trials where both were loaded onto the presentation device simultaneously and trials were considered successful if the monkey was successful with at least one of the two.


Trial No.	Delay (<i>no. of frames</i>)	Delay (<i>ms</i>)
1	17	567
2	26	868
3	0	0
4	20	667
5	23	767
6	8	267
7	19	634

Supplementary Table 4. Delay between monkey A's right hand movement and closing of the gripper. These measurements were obtained by going through Supplementary Video 1 frame by frame, for each of the 7 trials, visually identifying the frame when the monkey started extending its wrist, and then counting the number of frames until the gripper started to close. The video has a frame rate of 29.97 frames per second, resulting in roughly a 33 *ms* delay per frame.


Supplementary Figures




Supplementary Figure 1. Non-linear gain function designed to suppress low speeds and amplify high speeds.




Supplementary Figure 2. Data from a calibration session for monkey P. **a**, Directional tuning of a single unit (031a) during 6-7 repetitions of endpoint movements in 8 directions. With a preferred direction of (0.76,0.29,-0.57), this unit fired maximally in the forward-up-left direction (F,U,L) while reaching to the upper left target, and fired the least in the backward-down-right direction (B,D,R) while retrieving from the same target. Each row of lines in the raster plot represents a single trial (trial number shown to the left of each line). The 6-7 trials to each target comprise 9 iterations of calibration. **b**, Angular difference between the final preferred direction vector and the preferred direction at a given calibration iteration as a function of iteration number, for unit 031a (dotted line), and average (solid line) and standard deviation (bars) over all units. The data at iteration 0 correspond to the initial random preferred direction.




Supplementary Figure 3. Schematic diagram to explain order of operations in applying automated assistance during the *assisted* control paradigm. $\vec{p}_e[n]$ is the endpoint command output by the extraction module. $\vec{p}_{final}[n]$ is the final cartesian endpoint command that gets directly converted to joint angular robot command issued to the robot controller.




Supplementary Figure 4. Schematic graph to explain how the amount of attraction assistance was calculated as a function of endpoint distance from target during the *assisted* control paradigm. The *AttractionSpeed* that was applied, was essentially a constant value, except at close proximity to the target, the speed was attenuated linearly to zero to avoid overshooting the target. *AttractionMaxSpeed* and *AttractionLimitDistance* were configurable parameters with typical values of 50-150 mm/sec and 10 mm respectively.




Supplementary Figure 5. Quantification of positioning accuracy required for successful loading determined in an offline test without the monkey. Arm and gripper were automated using the *assisted* control paradigm with a BCGain value of 0 (i.e. control purely automated by *Attraction Assist* and *Gripper Assist* whereby the endpoint moves to target in a direct line, then the gripper closes, and endpoint moves back toward the “mouth” location). Marshmallows were presented on the tip of the presentation device and the automated system was able to load the marshmallow successfully in 13 out of 20 trials (65% success rate). The 20-trial test was then repeated at each of 24 different offsets where the automated system was told the wrong target position, offset from the true position by -15, -10, -5, 5, 10 and 15 mm in each of *x*, *y* and *z* dimensions. Percentage success rate drops below half the maximal value at an offset of 5 mm in negative *z* (aiming too low) and negative *x* (aiming too near, i.e. not far enough forward) and at an offset of 10 mm in most directions.




Supplementary Figure 6. Performance of monkey P during the 30 days of *assisted* 3D training (with no assistance during *Move A* and *Move B*). Performance is defined as Success Rate divided by Homing Radius. Success Rate drops when task difficulty is increased by decreasing the Homing Radius from day 18 onwards (arrow), but the Performance measure continues to show an upward trend.




Supplementary Figure 7. 3-dimensional spatial plot of endpoint trajectories of the Barrett arm (green) and the Keshen arm (blue) during a test when the command position (not shown) was moved at a constant speed (180 mm/s) from the “mouth” location to each of four target locations and back. Trajectories were obtained by placing an infra-red emitting marker on the end of each arm, running the arms simultaneously, and tracking the markers using an Optotrak 3020 system (Northern Digital Inc., Waterloo, Ontario, Canada). The Barrett arm follows command closely as shown by the straight trajectory lines (the slight curvature of the lines is caused by an offset of the optical marker from the arm endpoint because the marker could not be placed exactly on the endpoint). The Keshen arm shows significant oscillations around the commanded straight-line paths.




Supplementary Figure 8. Data from the 13 days of continuous self-feeding by monkey P. **a**, Number of units isolated (blue) and number of units used for control (red). **b**, Success rate. The monkey appeared unmotivated on days 7 and 13 when sudden drops in success rate can be seen. Sessions that consisted of 5 or fewer trials, and sessions in which the p-value from the regression in the calibration task (averaged over all units used in control) was more than 0.1, were excluded from this plot (because these criteria indicate that the monkey was not motivated). **c**, Distribution of preferred directions of the 15 units used on one of the 13 days.




Supplementary Figure 9. Reaching and retrieval trajectory consistency for monkey P during the training phase when the *assisted* control paradigm was used but monkey had full control during at least *Move A* and *Move B*. Solid colour lines indicate average endpoint trajectories for reaching (**a**) and retrieval (**b**) movements and the semi-transparent coloured regions represent the standard deviation of the trajectories. Convergence points of the straight grey lines represent locations of the four targets and the mouth, the four lines from the mouth to the targets being the ideal trajectories. The grey spheres represent the regions around the targets where the task switched from the *Move A* or *Move B* period to *Home A* or *Home B* respectively. The trajectories, their standard deviations, and the target region radii were averaged over all training sessions where the monkey had full 3D control during *Move A* and *Move B*.




Supplementary Figure 10. Scatter plot of endpoint speed vs. gripper aperture velocity as output by the extraction algorithm during a continuous self-feeding session. The session lasted about 40 *min* and included 70,003 data points (one data point per 30 *ms*). 5,000 of those points were randomly picked for plotting to avoid “over-crowding” the plot.



Supplementary Figure 11. Variability of gripper modulation across units on the same electrode (channel) on the first day of continuous self-feeding. The 116 units used for control were recorded from 63 channels. There were 24 channels with one unit each, 27 with two units each, 10 with three units each and 2 with four units each. Gripper modulation was calculated as the absolute value of the gripper component in the preferred direction vector. Since preferred direction vectors have unit length, a gripper modulation of 1 means that a unit is purely gripper-tuned and not-at-all endpoint movement tuned, whereas a value of 0 means that the unit is completely endpoint movement tuned and not-at-all gripper tuned, and a value between 0 and 1 means the unit is partly tuned to both. To get a measure of how different the units on a given channel were in terms of gripper vs. endpoint modulation, the gripper modulation values for all units on each channel (for the 39 channels that had two or more units each) were summarized by calculating their range, resulting in a single number per channel. This shows that individual units on a single electrode could be gripper-tuned and endpoint-tuned.



Supplementary Figure 12. Proportional gripper control. Endpoint position components (x, y and z) and gripper command are shown as a function of time during three trials from a session in monkey A's training period when the monkey had started attempting full 4-dimensional control but was not yet proficient at it. During this training period, proportional gripper control was seen, i.e. the monkey sometimes opened or closed the gripper part-way or varied the rate of opening or closing, as seen in this figure. Three separate trials were concatenated in time (with gaps between them) for plotting purposes.



Supplementary Figure 13. Keeping the gripper closed during retrieval. 3D endpoint trajectories during 4 trials from a session during monkey A's training when the monkey had started attempting full 4-dimensional control but was not yet proficient at it. During this training period, the monkey frequently kept the gripper closed all the way back to the mouth during the retrieval movement. This figure has the same format as Figure 2f, showing gripper aperture as continuously varying colour from blue (closed) to purple (part-way open) to red (open).

Supplementary Video Legends

Supplementary Video 1. Continuous self-feeding by monkey A showing 7 consecutive successful trials. The monkey's cortical control is 4-dimensional, including 3 dimensions of endpoint control plus gripper control. Part of the animal's head was obscured in the video frames (for this and the other videos).

Supplementary Video 2. Continuous self-feeding by monkey P showing 6 consecutive trials (5 successful). Monkey's cortical control is 3-dimensional, i.e. endpoint control. The gripper is controlled as a dimension dependent on endpoint movement: it opens when the arm moves forward and closes when the arm is held stable or moved backward. On the fourth reach, the food is dropped but the monkey immediately stops the arm, waits for a new target, and makes a reaching movement directly to the new target.

Supplementary Video 3. Target tracking. Monkey A reaches toward a presented food target during a continuous self-feeding session, but then the target is suddenly moved to a location where a direct move to target would knock the food off the presentation device. The monkey then moves the arm endpoint in a curved path to avoid the collision, and successfully obtains the food.

Supplementary Video 4. Finger licking behaviour. Monkey A licks the gripper fingers during a continuous self-feeding session. When a target is presented, the monkey starts reaching toward the target but then notices that there is more to be licked on the gripper, ignores the target, moves the arm back to the mouth to lick the fingers more, and then finally reaches toward the target. Being outside the task requirements, this emergent behaviour is an example of embodied prosthetic control.

Supplementary Video 5. Using the arm to push food into the mouth. Monkey A reaches out, grips and retrieves a marshmallow during a continuous self-feeding session. Upon unloading, the marshmallow ends up barely between the animal's lips, about to fall out. At that point, the monkey is unable to get the food into its mouth without a helping "hand", so it uses the robotic arm to push the food into its mouth. The video is shown first at normal speed and then replayed using slow motion.