

Direct Cortical Control of 3D Neuroprosthetic Devices

Dawn M. Taylor,¹ Stephen I. Helms Tillery,¹
Andrew B. Schwartz^{1,2*}

Three-dimensional (3D) movement of neuroprosthetic devices can be controlled by the activity of cortical neurons when appropriate algorithms are used to decode intended movement in real time. Previous studies assumed that neurons maintain fixed tuning properties, and the studies used subjects who were unaware of the movements predicted by their recorded units. In this study, subjects had real-time visual feedback of their brain-controlled trajectories. Cell tuning properties changed when used for brain-controlled movements. By using control algorithms that track these changes, subjects made long sequences of 3D movements using far fewer cortical units than expected. Daily practice improved movement accuracy and the directional tuning of these units.

Ever since cortical neurons were shown to modulate their activity before movement, researchers have anticipated using these signals to control various prosthetic devices (1, 2). Recent advances in chronic recording electrodes and signal-processing technology now open the possibility of using these cortical signals efficiently in real time (3, 4).

However, many neurons may be needed to predict intended movement accurately enough to make this technology practical. Estimates range from 150 to 600 cells or more being necessary (4, 5), based on open-loop experiments that recreate three-dimensional (3D) arm trajectories from cortical data offline (6). Here, we compare this approach to a closed-loop paradigm in which subjects have visual feedback of the brain-controlled movement. We then incorporate a movement prediction algorithm that tracks learning-induced changes in neural activity patterns.

Rhesus macaques made real and virtual arm movements in a computer-generated, 3D virtual environment by moving a cursor from a central-start position to one of eight targets located radially at the corners of an imaginary cube. The monkeys could not see their actual arm movements, but rather saw two spheres (the stationary “target” and a mobile “cursor”) with motion controlled either by the subject’s hand position (“hand-control”) or by recorded neural activity (“brain-control”) (see supplementary material).

We examined the effect of visual feedback on movements derived from cortical signals by comparing “open-loop” trajectories, created offline from cortical signals recorded during hand-controlled cursor movements, with “closed-loop” trajectories made

by the cursor under real-time brain control. In the closed-loop case, subjects saw the cursor movements created from their cortical signals in real time. In the open-loop case, the trajectories were created offline, after the experiment, from the cortical activity recorded during the movement blocks where the cursor was under hand control (7). Therefore, the subject had no knowledge of these offline brain-predicted trajectories. In both the open- and closed-loop cases, the same cortical decoding algorithm was used to generate trajectories. This decoding algorithm assumed that the cells’ tuning functions remained constant under both conditions.

This experiment was conducted with monkeys “L” and “M” for 32 and 40 days, respectively. In both subjects, about 18 cells were used to create open- and closed-loop trajectories. As expected, with so few cells, the open-loop trajectories were not very accurate. Although these trajectories went toward the correct targets more often than they would have by chance, they usually had at least one of the X, Y, or Z components pointing in the wrong direction.

Closed-loop trajectories ended in the target more often than did open-loop trajectories (Table 1). Both animals improved their closed-loop target hit rate over the course of the experiment, which suggests that the subjects learned to modulate their brain signals more effectively with visual feedback (8).

Many of the cortical units we recorded were stable from day to day. Some were stable for more than 2 years. Other units showed significant changes in their waveforms and movement properties between days (9). The brain-control algorithm was adjusted daily to make use of the current properties of the recorded units. Therefore, subjects had to learn a slightly different brain-to-cursor-movement relation each day. We looked for trends within days that would indicate learning of each new relation. Paired *t* tests showed that subjects initially improved their target hit rate by about 7% from the first to the third block of eight closed-loop movements each day ($P < 0.002$).

Subjects had 10 to 15 s to move the cursor to each target—enough time to use visual feedback to make online error corrections in the closed-loop case. We tested a subject’s ability to make more ballistic brain-controlled movements by continuing the experiment in monkey M for an additional 20 days with an increased brain-controlled cursor gain and the movement time constrained to 800 ms. As in the slow-movement case, the closed-loop trajectories still hit the targets more often than did the open-loop trajectories ($42 \pm 5\%$ versus $12 \pm 5\%$ of targets hit; $P < 0.0001$). Again, there was significant improvement with daily practice (0.9%/day; $P < 0.009$) as well as an initial improvement of about 7% within each day (first to third block; $P < 0.05$) (10). Despite the shorter movement time, visual feedback still allowed the subject to learn from consistent errors in the brain-controlled trajectories.

In these experiments, the movement-prediction algorithms were based on fixed tuning

Table 1. Mean \pm standard deviation of daily statistics from the open- versus closed-loop experiment. Percent time in the correct octant was calculated per movement as the % time the trajectory’s X, Y, and Z components had the same signs as the target [based on coordinate system with (0, 0, 0) at the center start position and each target located equal distance in the \pm X, Y, and Z directions]. Differences between the open- and closed-loop values were significant in all six categories ($P < 0.0001$).

	Monkey		
	L	M	Both
Closed-loop brain-controlled trajectories			
% Targets hit	52 \pm 14	46 \pm 18	49 \pm 17
% Time in correct octant	36 \pm 9	34 \pm 11	35 \pm 11
Open-loop brain-predicted trajectories			
% Targets hit	32 \pm 11	23 \pm 5	27 \pm 9
% Time in correct octant	23 \pm 9	23 \pm 9	23 \pm 9
Miscellaneous			
Cells used	18 \pm 4	18 \pm 3	18 \pm 4
Mean R^2	0.63 \pm 0.07	0.64 \pm 0.09	0.64 \pm 0.08

¹Department of Bioengineering, Arizona State University, Tempe, AZ 85287–6006, USA. ²The Neurosciences Institute, San Diego, CA 92121, USA.

*To whom correspondence should be addressed. E-mail: aschwartz@nsi.edu

properties obtained from neural activity recorded each day during a baseline set of hand-controlled cursor movements. This type of calibration cannot be carried out in movement-impaired patients. We have developed a “coadaptive” movement prediction algorithm that does not require physical limb movements or any a priori knowledge of cell tuning properties. By iteratively refining estimates of cell tuning properties as the subject attempted 3D brain-controlled cursor movements, we were able to track learning-induced changes in cell tuning properties.

We tested this coadaptive method in two healthy macaques by restraining both arms during a brain-control task after first recording each day's baseline hand-controlled movements and calculating each cell's tuning properties. At the end of each day's experiment, tuning functions were also calculated directly from the cortical activity collected during the brain-controlled movements. Figure 1A shows a unit whose directional tuning differed significantly between the brain-controlled movements and the hand-controlled movements made earlier that day. Both well-isolated individual cells and inseparable multi-cell groups showed substantial changes in their preferred directions between the two tasks. On average, the magnitude of these changes increased over the course of the experiment (Fig. 1B), and the direction of these changes varied from cell to cell (Fig. 1C). By the last 2 weeks of the experiment, these individual shifts in preferred direction became consistent from day to day (Fig. 1D).

Across days, the directional tuning of most units improved in the brain-control task versus the hand-control task (Fig. 2, A and B) (11). This increase in tuning quality was due, in part, to an improved fit of the units' firing rates to a cosine tuning equation under brain control (12). Although the control algorithm was designed to accommodate the most common deviations from cosine tuning (i.e., larger increases in rate with movements in the preferred direction than decreases in rate with movements opposite the preferred direction, as can be seen in Fig. 2C), the units still showed changes in their average tuning properties that were closer to a true linear function of the cosine of the angle between movement and preferred direction (Fig. 2D). These changes may have provided more uniform control and stability over the workspace (13). The daily improvement in cosine tuning was mirrored by a steady increase in the accuracy of the brain-controlled movements (Fig. 2E).

This technique shows how we could train immobile patients to make 3D cursor movements by coadapting a prediction algorithm to their changing cell tuning properties. However, for patients to control useful prosthetic devices, they would need to use this prediction algorithm without continued adaptation

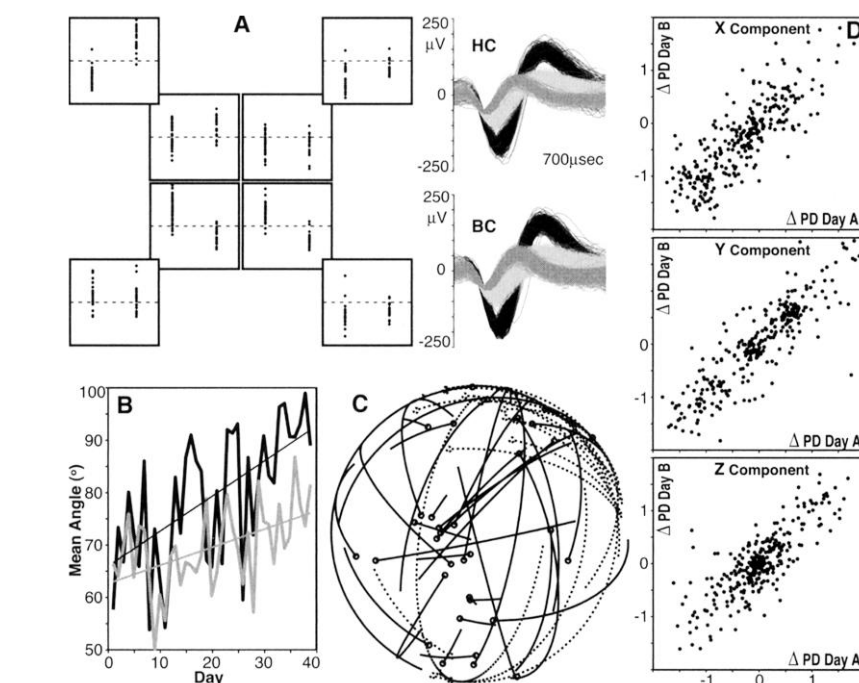


Fig. 1. Changes in cortical activity between hand-control and brain-control tasks in subject M. (A) Cell with a 107° change in tuning direction between the hand-control (HC) and brain-control (BC) tasks (the unit waveform is shown in black). Each dot is the mean firing rate during one movement. HC rates are in the right column and BC rates are in the left column of each square. The eight squares correspond to the eight target directions (center four = distal; outer four = proximal). (B) Daily mean angles (thick lines) between hand- and brain-controlled preferred directions for all cells significantly tuned during both tasks (black = contralateral and gray = ipsilateral units to the arm moved during the hand-control task). The thin diagonal lines are linear fits with slopes significant at $P < 0.006$ (contra) and $P < 0.0001$ (ipsi). (C) Lines connecting hand-controlled preferred directions with brain-controlled preferred directions (circle ends) projected onto a unit sphere (day 28, only cells significantly tuned in both tasks; black = contra.; dotted = ipsi.). (D) Change in the X, Y, and Z components of the preferred direction unit vectors between the hand- and brain-controlled tasks plotted day-against-day for eight random pairs of days (days 27 or later, only units that were significantly tuned in both tasks on both days; 35 ± 3 units per pair of days) (see supplementary material).

of its parameters, and they would want to make a wider variety of movements than the ones practiced during the coadaptive training.

We tested these issues by following several days' coadaptation training with an additional movement task, the constant-parameter prediction algorithm (CPPA) task, which used fixed tuning parameters, added novel target positions, and required 180° changes in movement directions (Fig. 3).

There was no significant difference between the novel and trained target hit rates in either animal, and both monkeys improved their performance with daily practice (Table 2) (14). Movies of these brain-controlled movements are included in the supplementary material.

Our work shows that visual feedback combined with an algorithm that tracks changes in cortical tuning parameters improves the efficacy of cortical activity as a control signal for both fast and slow brain-controlled movements. Switching from the hand-controlled to the brain-controlled task caused global changes in the tuning parameters of the recorded neuronal population.

How the rest of the population shaped and supported these changes is still an open question. The increased consistency of these changes across days combined with the improvement in performance suggests that the learning process settled on an effective set of parameters for the imposed control scheme.

During the first several days of the coadaptive experiment, our monkeys pushed methodically against the arm restraints in the direction they needed the cursor to move. However, this behavior quickly subsided as performance improved. Spot checks of electromyographic (EMG) activity on the well-adapted days showed suppression of EMG activity throughout the brain-control task. This indicates that it is possible to develop effective brain-control modulation patterns in the absence of physical limb movements or normal muscle activation patterns.

Although the healthy, arms-restrained animal model may not address all the issues related to retraining cortex altered from disuse after an injury or illness, recent magnetic

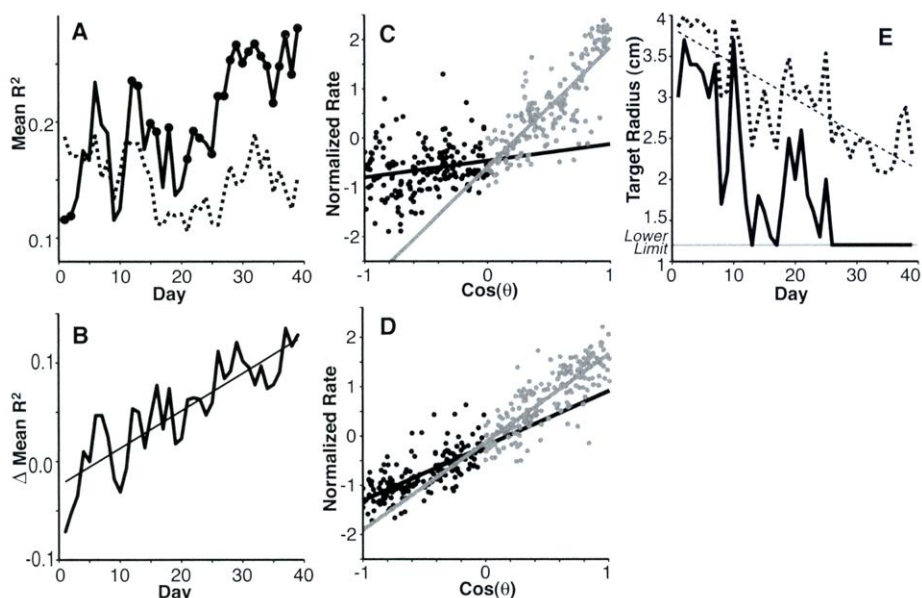
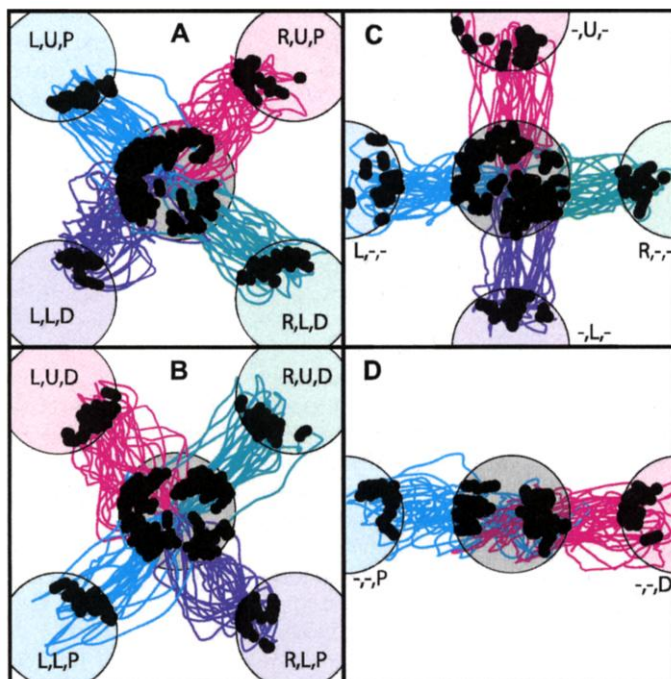


Fig. 2. Changes in cell tuning quality and performance. (A) Daily average R^2 from regressing each cell's mean firing rate per movement against target direction. The dotted line shows hand-control R^2 values. The black line shows brain-control R^2 values. The closed circles indicate days when brain- and hand-control R^2 values were significantly different (paired t test, $P < 0.05$). (B) Difference in R^2 values (brain-control minus hand-control). The thin diagonal line is the linear fit ($P < 0.0001$). (C and D) show average normalized firing rates to each target plotted as a function of the component of the movement in each cell's preferred direction for day 39. The lines are linear fits to all points with $\cos \theta$ above zero (gray) or below zero (black). (C) Hand-controlled task (gray line, $R^2 = 0.65$; black line = 0.04). (D) Brain-controlled task (gray line, $R^2 = 0.67$; black line = 0.54). (E) Daily minimum (solid line) and mean (dotted line) target radii used to maintain a 70% target hit rate. The diagonal dotted line is the linear fit of the daily mean ($P < 0.0001$). The bottom horizontal line shows the minimum target radius allowed (1.2 cm) (see supplementary material).

Fig. 3. Monkey M's brain-controlled trajectories in the CPPA task. Trajectories start from the exact center, go to an outer target (colored circles), and return to the center target (gray circle). Trajectories are color-coded to match their intended targets. The black dots indicate when the intended outer or center target was hit. The three letters by each target indicate Left (L)/Right (R), Upper (U)/Lower (L), and Proximal (P)/Distal (D) target locations. The dashes indicate a middle position. (A and B) are to the eight "trained" targets used in the coadaptive task. (C and D) are to the six "novel" targets.



resonance imaging research indicates that the underlying motor maps are maintained, even after years of paralysis (15). Additionally, there are a few cases where completely immobile or "locked-in" patients have had cortical electrodes implanted and have been

taught to communicate by scrolling through a sequence of letters using the activity of a few motor cortex cells. These human case studies suggest that cortical cells can regain and maintain the level of activity needed to perform prosthetic control tasks—even after

Table 2. Mean \pm standard deviation of daily performance statistics during the coadaptive and CPPA tasks. "# Units recorded" includes "noise" units that were removed during coadaptation. Coadaptive target hit rates were recalculated with targets at various radii based on movements made after the algorithm had converged.

	Monkey	
	M	O
CPPA task		
% Targets hit		
Novel	80 \pm 26	73 \pm 29
Trained	77 \pm 24	62 \pm 30
Center after novel	80 \pm 22	72 \pm 25
Center after trained	82 \pm 19	70 \pm 21
Average movement time (s)		
Novel	1.5 \pm 0.5	2.0 \pm 0.6
Trained	1.5 \pm 0.6	2.6 \pm 0.7
Center after novel	1.3 \pm 0.7	2.0 \pm 1.1
Center after trained	1.6 \pm 0.8	2.0 \pm 0.9
# Days in calculations	12	5
# Units recorded	64 \pm 2	31 \pm 2
# Units used	38 \pm 2	17 \pm 2
Coadaptive task		
% Targets that would have been hit at the following target radii		
1.2 cm	76 \pm 12	20 \pm 22
2.0 cm	86 \pm 8	47 \pm 21
3.0 cm	94 \pm 4	69 \pm 9
4.0 cm	98 \pm 3	81 \pm 8
# Days in calculations	13	14
# Units recorded	64 \pm 2	35 \pm 6
# Units used	39 \pm 2	21 \pm 4

long periods of complete immobility (16). Our results show that neural activity can be reorganized within minutes and, with the proper algorithm, used to achieve brain-controlled virtual movements with nearly the same accuracy, robustness, and speed as normal arm movements.

References and Notes

1. M. Craggs, *Adv. Neurol.* **10**, 91 (1975).
2. For review, see J. R. Wolpaw et al., *IEEE Trans. Rehab. Eng.* **8**, 164 (2000).
3. J. K. Chapin, K. A. Moxon, R. S. Markowitz, M. A. L. Nicholelis, *Nature Neurosci.* **2**, 664 (1999).
4. J. Wessberg et al., *Nature* **408**, 361 (2000).
5. A. P. Georgopoulos, R. E. Kettner, A. B. Schwartz, *J. Neurosci.* **8**, 2931 (1988).
6. These estimates were derived by combining neural activity recorded during repeated trials of similar movements to simulate the recording of larger numbers of neurons during a single movement. The lower estimate (5) used movable single-channel electrodes that recorded a few well-tuned cells each day. The larger estimate (4) used multi-channel fixed arrays where the tuning quality of the recorded cells is more representative of the neural population. With current technology, the number of units that can be recorded simultaneously in animals is between ~30 and 100.
7. Hand- and brain-controlled center-out movements were done in alternating blocks of movements to all eight targets.
8. Closed-loop minus open-loop target hit rate as a function of number of days of practice showed an increase of 1% per day in subject M ($P < 0.0001$) and 0.8% per day in subject L ($P < 0.0004$).
9. Stability based on waveforms and tuning properties (17).

10. Improvement per day is for closed-loop minus open-loop hit rate. P values for differences are based on paired t tests. The mean % of each trajectory in the correct octant (as described in Table 1) = $18 \pm 3\%$ (open-loop) and $32 \pm 3\%$ (closed-loop); $P < 0.0001$.
11. The R^2 values are lower than what's usually seen in the literature, because rates were not averaged together across all movements per target. Brain-controlled R^2 values were calculated at the 420-ms window length because this matched the average window length used in the hand-control calculations.
12. "Tuning equation fit," measured by target-averaged R^2 values, was significantly higher during brain-control than hand-control ($\Delta R^2 = R^2_{\text{(brain-control)}} - R^2_{\text{(hand-control)}} = 0.18 \pm 0.06$), and this difference increased with practice ($d\Delta R^2/\text{day} = 0.004$, $P < 0.0001$).
13. In the earlier arms-free brain-controlled experiments, asymmetries in the tuning functions resulted in widely varying levels of movement accuracy throughout the workspace. Also, becoming more cosine tuned meant that smaller deviations from the mean firing rates were needed to hold the cursor stationary.
14. Monkey M's target hit rate increased by 3% per day ($P < 0.01$), and the mean sequence length increased by 2.5 targets per day ($P < 0.01$). By the last day, the subject was regularly making sequences of 50 to 70 movements without missing.
15. S. Shoham, E. Haggren, E. M. Maynard, R. A. Normann, *Nature* **413**, 793 (2001).
16. P. R. Kennedy, R. A. E. Bakay, M. M. Moore, K. Adams, J. Goldwithe, *IEEE Trans. Rehab. Eng.* **8**, 198 (2000).
17. J. C. Williams, R. L. Rennaker, D. R. Kipke, *Neurocomputing* **26**, 1069 (1999).
18. We thank D. Moran for allowing us to adapt his base virtual reality program; W. Lyn for his recording assistance; and D. Weber for designing the EMG recording system. Supported by a Whitaker Foundation Fellowship, a Philanthropic Education Organization scholarship, and U.S. Public Health Service contract numbers N01-NS-6-2347 and N01-NS-9-2321.

Supporting Online Material

www.sciencemag.org/cgi/content/full/296/5574/1829/DC1

Materials and Methods
Movies S1 to S3

28 January 2002; accepted 12 April 2002

REPORTS

Equilibrium Information from Nonequilibrium Measurements in an Experimental Test of Jarzynski's Equality

Jan Liphardt,^{1,4} Sophie Dumont,² Steven B. Smith,³
Ignacio Tinoco Jr.,^{1,4} Carlos Bustamante^{1,2,3,4*}

Recent advances in statistical mechanical theory can be used to solve a fundamental problem in experimental thermodynamics. In 1997, Jarzynski proved an equality relating the irreversible work to the equilibrium free energy difference, ΔG . This remarkable theoretical result states that it is possible to obtain equilibrium thermodynamic parameters from processes carried out arbitrarily far from equilibrium. We test Jarzynski's equality by mechanically stretching a single molecule of RNA reversibly and irreversibly between two conformations. Application of this equality to the irreversible work trajectories recovers the ΔG profile of the stretching process to within $k_B T/2$ (half the thermal energy) of its best independent estimate, the mean work of reversible stretching. The implementation and test of Jarzynski's equality provides the first example of its use as a bridge between the statistical mechanics of equilibrium and nonequilibrium systems. This work also extends the thermodynamic analysis of single molecule manipulation data beyond the context of equilibrium experiments.

Irreversible processes as diverse as mechanically induced protein unfolding, the fracture of stressed materials, and the sudden formation of crystallization nuclei all involve the time evolution of states far removed from equilibrium. To characterize these nonequilibrium states, it is generally necessary to specify numerous details of the system and its surroundings. By contrast, reversible processes are idealizations in which a system passes only through a succession of equilibrium

states, which can be described completely with only a few variables such as pressure and temperature. Reversible processes are powerful tools in thermodynamics because they make it possible to relate the measured heat and work to the thermodynamic state variables. Yet many processes in nature relax to equilibrium only very slowly, precluding quasi-reversible experiments and thus preventing measurement of the thermodynamic state variables. Solving the problem of recovering thermodynamic variables from irreversible experiments remains one of the unfinished tasks in thermodynamics.

It follows from the laws of thermodynamics, first formulated in the early 19th century, that the increase in Gibbs free energy ΔG and the mean work $\langle w \rangle$ needed to bring about that increase are related by $\Delta G \leq \langle w \rangle$. The equal-

ity holds when a process is carried out reversibly, and the inequality holds otherwise. In 1951, Callen and Welton realized that for any system that remains near equilibrium, the energy dissipated is proportional to the system's fluctuations (*1*). With this fluctuation-dissipation relation, researchers acquired a better estimate of ΔG for irreversible processes: $\Delta G \approx \langle w \rangle - \beta \sigma^2/2$, where σ is the standard deviation of the work distribution and $\beta^{-1} \equiv k_B T$ (where T is absolute temperature and k_B is Boltzmann's constant) (*2–4*). Unfortunately, this ΔG estimate is valid only in the near-equilibrium regime, and so it was thought that free energies could only be obtained for processes remaining close to equilibrium.

This state of affairs changed in 1997, when Jarzynski derived an equality (*5–8*) that relates the free energy difference separating states of a system at positions 0 and z along a reaction coordinate, $\Delta G(z)$, to the work done to irreversibly switch the system between two states,

$$\exp[-\beta \Delta G(z)] = \lim_{N \rightarrow \infty} \langle \exp[-\beta w_i(z, r)] \rangle_N \quad (1)$$

where $\langle \rangle_N$ denotes averaging over N work trajectories, $w_i(z, r)$ represents the work of the i th of N trajectories, and r is the switching rate (*9*). The mechanical work $w_i(z, r)$ required to switch the system between positions 0 and z under the action of a force F is

$$w_i(z, r) \approx \int_0^z F_i(z', r) dz' \quad (2)$$

where $F_i(z', r)$ is the external force applied to the system at position z' with switching rate r (*10*). Equations 1 and 2 state that the free energy change for a reaction can be determined by averaging Boltzmann-weighted work values obtained from repeated irreversible switching of the system (*11, 12*). Unlike most expressions relating equilibrium and nonequilibrium statistical mechanics, Jarzynski's equality holds for systems driven arbitrarily far from equilibrium [for other relations that are valid in the far-from-equilibrium regime, see, e.g. (*13–20*)].

¹Department of Chemistry, ²Biophysics Graduate Group, ³Departments of Physics and Molecular and Cell Biology and Howard Hughes Medical Institute, University of California, Berkeley, CA 94720, USA. ⁴Physical Biosciences Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA.

*To whom correspondence should be addressed: carlos@alice.berkeley.edu