Distributed cortical adaptation during learning of a brain–computer interface task

Jeremiah D. Wander[a,1], Timothy Blakely[a], Kai J. Miller[b,2], Kurt E. Weaver[c], Lise A. Johnson[d], Jared D. Olson[e], Eberhard E. Fetz[a,b,f], Rajesh P. N. Rao[a,b,g], and Jeffrey G. Ojemann[b,c,d]

Departments of [a]Bioengineering, [c]Radiology, [d]Neurological Surgery, [e]Rehabilitation Medicine, [f]Physiology and Biophysics, and [g]Computer Science and Engineering and [b]Program in Neurobiology and Behavior, University of Washington, Seattle, WA 98195

The majority of subjects who attempt to learn control of a brain–computer interface (BCI) can do so with adequate training. Much like when one learns to type or ride a bicycle, BCI users report transitioning from a deliberate, cognitively focused mindset to near automatic control as training progresses. What are the neural correlates of this process of BCI skill acquisition? Seven subjects were implanted with electrocorticography (ECoG) electrodes and had multiple opportunities to practice a 1D BCI task. As subjects became proficient, strong initial task-related activation was followed by lessening of activation in prefrontal cortex, premotor cortex, and posterior parietal cortex, areas that have previously been implicated in the cognitive phase of motor sequence learning and abstract task learning. These results demonstrate that, although the use of a BCI only requires modulation of a local population of neurons, a distributed network of cortical areas is involved in the acquisition of BCI proficiency.

brain–machine interface | motor learning | plasticity | electrophysiology | high gamma

Over the last 50 years, it has been demonstrated that, when given feedback, the brain can learn to volitionally modulate the activity of single neurons (1–5) and populations of neurons (6–9). This modulation can occur in the absence of overt movement (6, 9) or even when overt movement is not possible (3, 10). The use of these signals to control external devices for restoration of lost function or as potential feedback signals for rehabilitation holds great promise for the future; however, our understanding of the underlying processes that drive this activity modulation is incomplete.

The ability to voluntarily modulate neural activity to control a brain–computer interface (BCI) appears to be a learned skill, similar to learning to ride a bike or swing a golf club. Investigators have repeatedly documented that task performance typically increases over the course of practice (11–15), which is indicative of a learning process taking place in the brain. In many cases, human BCI users have anecdotally reported transitioning from a very deliberate, cognitive approach to nearly automatic execution. Whereas they were initially focused on the cognitive task (e.g., motor imagery) that drives neural activity and in turn controls the end effector, after training, they report a goal-directed approach focused directly on the end effector itself.

BCI use is interesting in that it shares characteristics with both concrete motor tasks and abstract cognitive tasks (16). Users of BCIs are manipulating their physical environment, but doing so without necessarily moving their bodies, thus depriving their brains of the multiple modalities of sensory feedback that have been demonstrated to be vital to learning a motor skill (17).

The transition from a cognitive to an automatic phase, combined with the fact that the majority of single-unit and field-potential BCIs are driven by signals recorded from motor areas, suggests the hypothesis that BCI skill learning shares many commonalities with motor sequence learning. Indeed, in most cases, large initial gains in performance are followed by gradually diminishing improvements, congruent with Fitts' well-known

work on motor performance during skill learning (18). Following results on neural activity patterns during motor skill learning (19, 20), one would therefore expect significant involvement of motor, prefrontal, and parietal cortical areas, as well as the striatum and cerebellum during the cognitive phase of BCI skill acquisition.

To investigate the role of distributed cortical networks in BCI skill acquisition, we recorded population-scale neural activity simultaneously from various locations across the cortex using electrocorticography (ECoG) while subjects performed a BCI task. The task required subjects to modulate a subrange (~70–100 Hz) of the high-gamma band (HG; 70–200 Hz) in ECoG surface potentials to control the position of a cursor on a computer monitor. HG activity was used because it has been previously shown to correspond to an increase in the firing rate and/or coherence of underlying local neural populations (21). We investigated two hypotheses: (i) that motor learning networks across the cortex participate in BCI learning even though the task only requires modulation of a localized neural population and (ii) that activity in these distributed cortical networks changes with increasing BCI task proficiency.

Results

Seven human subjects (two women; mean age, 25.6 y; range, 18–32 y) with intractable epilepsy were implanted for long-term (4–10 d) monitoring. During the monitoring period, these subjects participated in multiple experimental recording sessions that took place over the course of 1–3 d. The first session began with a series of two basic functional screening tasks, in which subjects performed gross motor movement or imagined motor movement in response to a visual cue. Functional screening was conducted to determine the electrodes that exhibited statistically significant increase in HG activity during actual movement or imagined movement compared with rest (9, 22). The controlling electrode that was used for online control was selected from these electrodes based on the magnitude of observed power change and anatomical relevance.

The BCI task was the 1D, two-target right-justified box task (23). In this task, the subject is presented one of two targets along the right side of a monitor and a cursor that starts on the left side of the monitor and travels to the right at a fixed horizontal velocity. During each trial, the subject modulates HG activity to control the vertical velocity of the cursor such that the cursor comes in contact with the target. HG activity recorded from the controlling electrode is mapped to vertical cursor

velocity using a simple linear decoder that was fixed in the first recording session, such that execution of motor imagery or overt motor movement causes the cursor to travel up and remaining at rest causes the cursor to travel down. The task consists of four periods: rest, targeting, feedback, and reward. Because the horizontal velocity of the cursor is fixed, the length of the feedback period for each trial is also fixed to 3 s. In subsequent discussion of this task, targets requiring an increase of HG activity are referred to as "up targets" and targets requiring a decrease of HG activity are referred to as "down targets."

Behavioral Performance. Overall performance, depicted in Fig. 1, increased significantly as subjects gained experience with the task ($P = 0.035$; right-sided t-test, $n = 7$). Behavioral performance typically began above chance levels and saturated quickly. Cursor trajectories could potentially be used to infer nonsaturating performance metrics but would require an assumption be made regarding optimal trajectory (see *SI Results* for further discussion). Thus, we used activity patterns at the controlling electrode to assess learning-related changes and define a transition from a learned to an unlearned state.
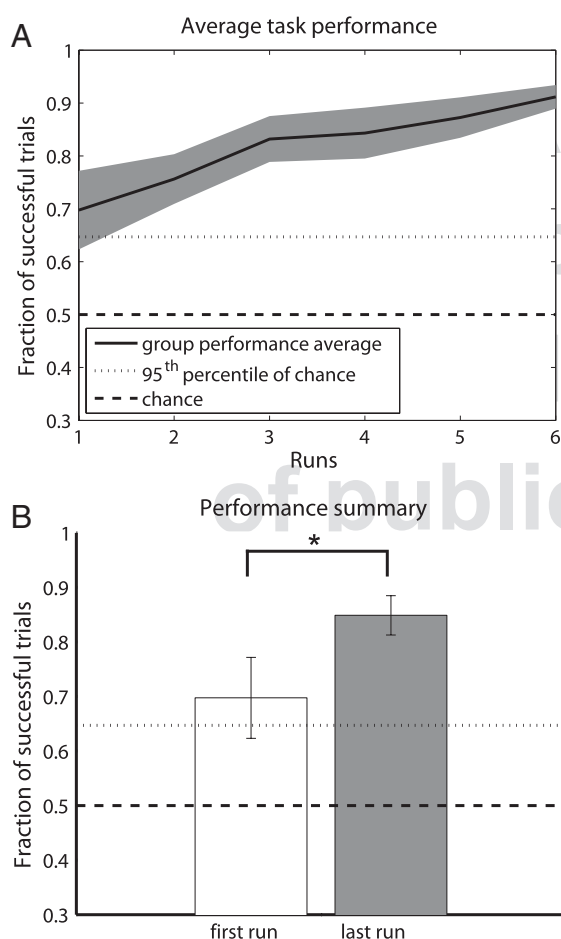
Fig. 1. Behavioral performance. (*A*) Mean task performance for all users. SEM shown in gray around trend line. Chance performance and the 95th percentile of chance performance are shown as horizontal dashed and dotted lines, respectively. Note that mean performance was above chance beginning with the first run. The number of runs performed by each user varied: all seven users contributed to data for the first three runs, six users contributed through the first five runs, and four contributed to the sixth run. See Table S1 for additional details. (*B*) Comparison of performance for the first and last runs performed by each subject. Significant differences (*P* < 0.05) denoted with an asterisk.

Volitional Modulation of Activity at the Controlling Electrode. In six of seven subjects (S1–S6), we found a statistically significant increase in HG activity during the feedback period of up targets compared with all rest periods (right-sided two-sample t test, Bonferonni corrected, $27 < N_1 < 98$, $51 < N_2 < 187$; $P < 0.0001$; see Table S1 for details). The seventh subject (S7, $N_1 = 45$, $N_2 = 85$, $P = 0.271$) did not demonstrate a significant difference between up targets and rest; however, in this subject, we found a statistically significant decrease in HG activity between down targets and rest (left-sided two-sample t test, Bonferonni corrected, $N_1 = 45$, $N_2 = 85$, $P = 0.0028$), suggesting that, although activity suppression below baseline is generally not used as a control strategy, it may have been the strategy used by this subject. These same activity decreases during down targets were not observed in the other six subjects (left-sided two-sample t test, Bonferonni corrected, $27 < N_1 < 100$, $51 < N_2 < 187$; $0.075 < P < 1$).

Task-Modulated Activity Throughout Cortex. We recorded ECoG data from 652 electrodes across the seven subjects. Of these electrodes, 83 were excluded from analyses because they contained nonphysiologic artifacts (resulting from poor contact, placement over scar tissue, etc). In the remaining 569, we found 152 electrodes showing statistically significant increases in HG activity between feedback during up targets compared with rest (right-sided two-sample t test, Bonferroni corrected, $27 < N_1 < 98$, $51 < N_2 < 187$; $P < 8.787 \times 10^{-5}$). Further, of that same 569 electrodes, we found 125 electrodes showing statistically significant increases in HG activity between all targets and rest (right-sided two-sample t test, Bonferroni corrected, $54 < N_1 < 198$, $51 < N_2 < 187$; $P < 8.787 \times 10^{-5}$). Electrodes showing significant activity increase were concentrated in cortical areas previously known to be associated with motor learning: primary motor cortex (MC), primary somatosensory cortex (SC), dorsolateral prefrontal cortex (PFC), and dorsal/ventral premotor cortices (PMd/PMv) (24, 25). Additionally, electrodes showing increases were found in posterior parietal cortex (PPC), an area associated with sensorimotor tasks involving visual feedback (26). Additional task-modulated electrodes were found in other cortical areas, such as the temporal parietal junction and the inferior temporal gyrus, although to a lesser extent. Task-modulated activity in supplementary motor area (SMA) was noted in the only subject with SMA coverage, but this finding was not included in our analyses as it could not be verified in multiple subjects. Fig. 2*A* illustrates the spatial distribution of electrodes showing significant activation for up targets relative to rest.

Determination of Unlearned vs. Learned States. Motor skill learning has been characterized as having multiple distinct learning phases: cognitive, associative, and autonomous (18). These phases loosely correspond to understanding what actions a certain skill requires, optimization of performance at a skill through repeated action, and development of an automatic capacity to carry out the skill, respectively. Recent work has developed a neural model describing the involvement of various cortical and subcortical structures during motor learning (19, 27). This model postulates involvement of prefrontal and premotor as well as parietal areas during the initial cognitive phase of motor sequence learning.

With this motivation in mind, we sought to capture differences in HG activity at the controlling electrode that changed with this transition from an "unlearned" to a "learned" state. The methodology for this is described in detail in *SI Materials and Methods*. In brief, we used a data-driven approach, assuming that as a subject transitions from an unlearned state to a learned state, there would be a corresponding change in the subject's ability to differentially modulate activity at the controlling electrode for up and down targets. To determine the approximate trial when this transition was taking place, the difference of HG activation for
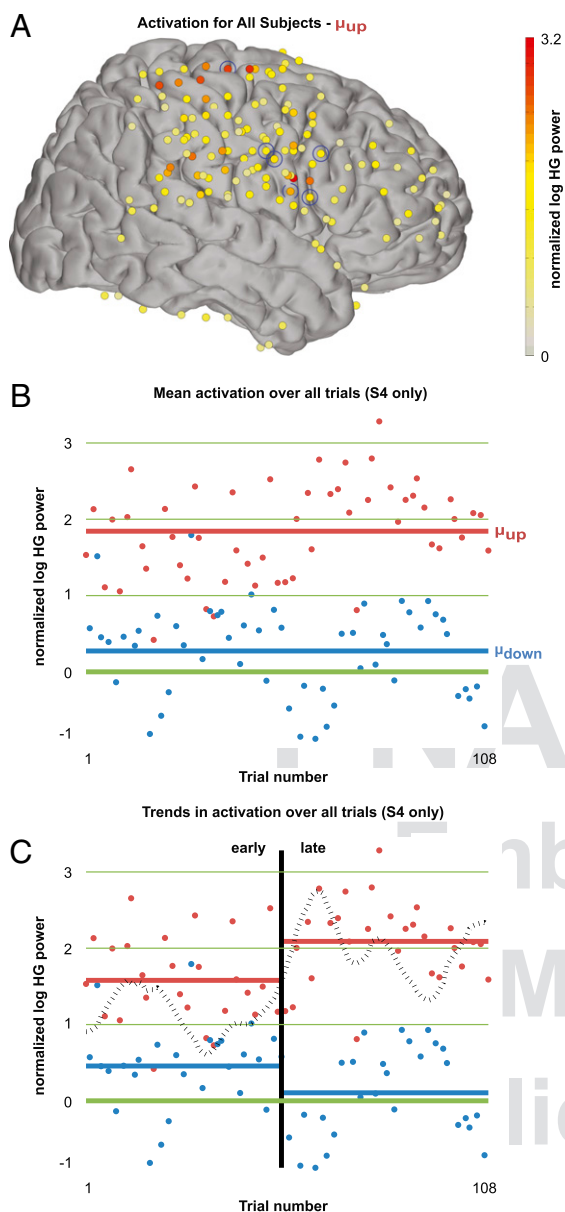
Fig. 2. Cortexwide activity during BCI use and modeling of early vs. late activity patterns. (*A*) Activation during up targets for all lateral electrodes for all subjects (left coverage projected to right hemisphere) shown on the Talairach brain. Note widespread cortical activation including frontal, middle-parietal, and posterior-parietal areas. Controlling electrodes are circled in blue. (*B*) Activation for an example subject (S4) during each feedback period normalized against log HG power during rest. Each dot represents one trial; up and down targets are shown in red and blue, respectively. Thick red and blue lines represent mean activation for all trials, respectively. (*C*) Early-late trial division shown for subject S4. Separability of high-gamma activity during up and down trials was used as a measurement of task proficiency. A separabililty measure (black dotted line) was modeled as two Gaussian distributions and fit to the data such that the distance between the distributions was maximized.

up and down targets was estimated and divided into unlearned and learned trials, occurring before and after an assumed transition trial, respectively. With the location of this assumed transition trial being the single free parameter, we then fit a model of two Gaussian distributions to the unlearned and learned trials such that the distance (see *SI Materials and Methods*) between these two distributions was maximized. In comparing unlearned and learned distributions, distance value magnitudes close to

one implied that the majority of the variance in the joint distribution can be explained by the difference in the means of the two subdistributions. The sign of the resulting distance measure demonstrates the direction of shift in the means of the two subdistributions, where a positive value implies that the mean of the learned distribution is greater than the mean of the unlearned distribution. Within the context of this task, a large, positive distance value for a given transition trial implies that the subject's ability to differentially modulate activity during up targets compared with down targets was much greater after the transition trial compared with before. Furthermore, a relatively large distance value implies that one of the solution strategies used was to increase the difference in activity at the controlling electrode in up targets relative to down targets. This possibility is compared with the alternative that the user only reduces the variability of their control signal as they gain experience. If this was the case, we would have observed relatively low distances between the unlearned and learned distributions.

In all seven subjects, relatively large distances between the distributions of power separation for learned and unlearned states demonstrated that the applied model effectively separated those two states ($0.2039 <$ distance < 0.8689). Accordingly, the differences in distributions of power separation for unlearned and learned states were highly significant (both-sided two-sample *t* test, Bonferonni corrected, $P < 0.0001$; see Table S1 for details).

Cortexwide Changes in Activation from Unlearned to Learned States. HG activation of all noncontrolling electrodes was visualized on a trial-by-trial basis to observe activity dynamics at these electrodes. Fig. 3 shows example trial-by-trial plots for electrodes located throughout motor-learning-associated cortical areas (see Fig. S1 for similar plots separated by subject). MC/SC and PMd/PMv electrodes demonstrated activation primarily during up targets, and PFC and PPC demonstrated activation during both up and down targets. As can be seen in Fig. S1, changes in activation in these electrodes over the course of many trials were often, although not always, well aligned with the transition from unlearned to learned state as defined solely by HG power at the controlling electrode. It is notable that in some subjects, transitions in HG activity patterns were approximately temporally aligned with breaks in experimental sessions, suggesting that offline learning was taking place during these periods. This result is complimentary to previous findings of increased sleep spindle density local to the controlling electrode correlated with training on an ECoG BCI (28).

Dynamics of activity in remote electrodes were quantitatively evaluated by comparing mean activity during feedback in these electrodes in the unlearned and learned states. Sixty-seven electrodes showed a significant change in HG activity during all targets from the unlearned to learned states (both-sided two-sample *t* test, Bonferroni corrected, $P < 8.787 \times 10^{-5}$, $20 < N_1 < 75$, $15 < N_2 < 166$). A large portion of these changes corresponded to significant lessening in activation in the frontal cortex, superior parietal cortex, and posterior parietal cortex, as depicted in Fig. 4*A*. A smaller portion corresponded to significant increases of activation in areas surrounding the controlling electrode. To determine anatomically relevant patterns of activity dynamics, electrodes were assigned to cortical areas using the human motor area template (HMAT) atlas (29) and the Talairach Daemon (30, 31). Significant lessening of activation was found in 31 electrodes corresponding to labels of PMd, PMv, PFC, and PPC, as shown in Fig. 4*B*.

Discussion

Our results demonstrate that when learning a BCI cursor-control task based on signals recorded from a single electrode over motor cortex, there is a distributed network of cortical areas that is involved in acquisition of skill in the task. This network includes, but
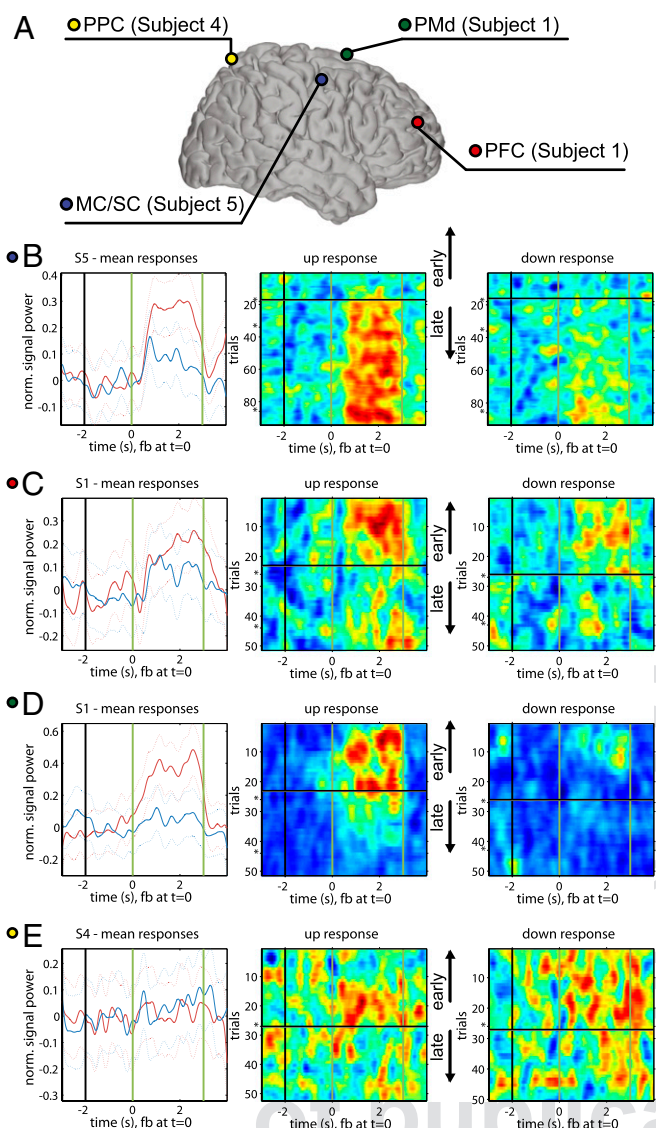
Fig. 3. Time-by-trial high-gamma activation. (*A*) Distribution of electrodes shown in *B–E* on the Talairach brain. The subject from which each electrode was taken is noted; not all subjects had coverage in all areas (Fig. S2). (*B–E*) Subplot 1: average HG activation in a given electrode for all trials separated by up and down targets. Subject is specified in the subplot title. Phases of the task (ordered from L to R: inter-stimulus interval, target presentation, feedback, reward) are separated by vertical bars. Dotted line represents SEM. Subplots 2 and 3: trial-by-trial HG activation for all trials, separated by up (subplot 2) and down (subplot 3) targets. Trial count is shown on the vertical axis. Breaks in the experimental session of more than 8 h are denoted with an asterisk (*). Time, as described for subplot 1, is shown on the horizontal axis. The black horizontal bar represents the transition trial, and the early and late trials relative to this point are denoted with black arrows. (*B*) MC/SC electrode shows continued increase in activation for up targets, congruent with changes at the controlling electrode. (*C*) PFC electrode shows activation for both up and down targets, decreasing near the point of model separation. (*D*) PMd electrode shows activation for up targets, decreasing near the point of model separation. (*E*) PPC electrode shows subtle activation during up and down targets, decreasing near the point of model separation.

is not limited to, surrounding sensorimotor and visuomotor areas: PMd/PMv, PFC, and PPC. We further demonstrated that as a user develops proficiency with the BCI, activity in parts of this network (PMd, PFC, and PPC) decreases, a finding that is indicative of the mental shift from cognitive to automatic task execution often anecdotally reported in human BCI studies.

Controlling a BCI with volitional modulation of activity from motor areas is intriguing in that it is arguably a nonmotor or at least quasi-motor task. One might expect that the same cortical (PMd/PMv, PFC, PPC) (24, 25) and subcortical areas (basal ganglia, cerebellum) (32) that are involved in motor tasks would also be involved in the control of a BCI. However, the two tasks are significantly different in that, with the exception of visual feedback, the afferent activity that is critical to successful acquisition of traditional motor tasks is not present during the use of a BCI. Thus, motor-based BCI may provide an opportunity to study these networks in the absence of or with manipulated sensory feedback.

In this study, we sought to investigate changes taking place in motor networks as users developed both their control strategy and execution proficiency. To do so, we used a data-driven approach that relied solely on the signal that each subject was using to control the BCI. We found that dorsal premotor, prefrontal, and posterior parietal cortices exhibited decreased task-modulated activity as users transitioned from a naive to more experienced
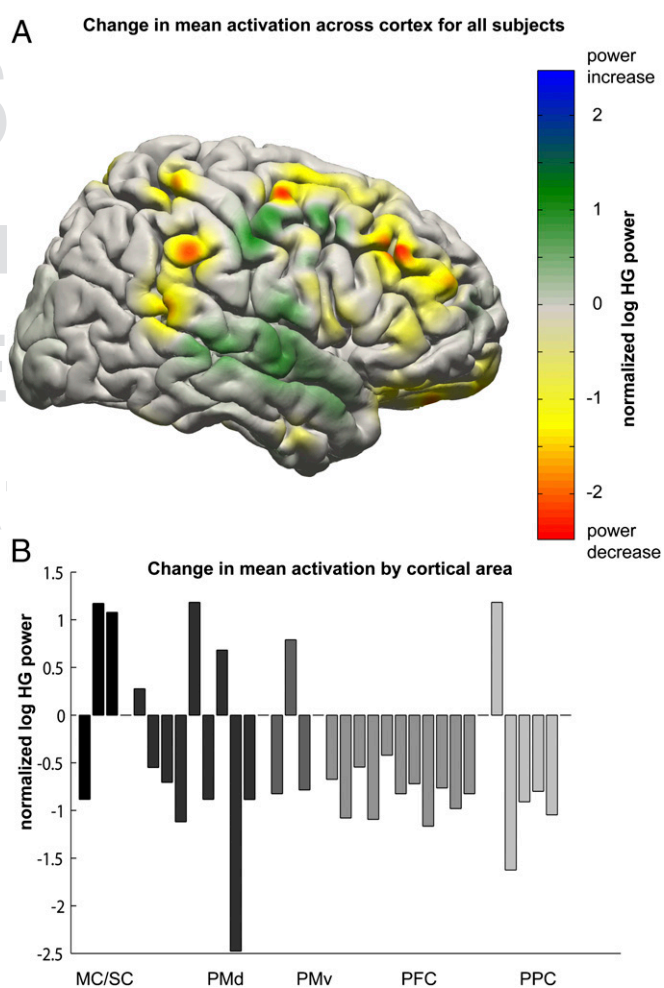

Fig. 4. Cortexwide activation changes. (*A*) Spatial distribution of change in mean activation comparing early to late trials for all targets for all subjects. Activations for individual electrodes are normalized against rest periods from the same electrode for a given run. Activation change values are blurred using a 12-cm FWHM Gaussian filter. Frontal areas and posterior parietal areas exhibited lessening in task-related activation over the course of BCI use. (*B*) Change in mean activation for all electrodes showing significant change from early to late trials, classified into approximate cortical areas.

state. This transition occurred quickly and was contemporaneous with changes taking place in the control signal.

It is interesting to note the patterns of distributed activity observed across multiple subjects. The time course of activation in premotor electrodes parallels tightly that of corresponding MC/SC electrodes. Further, premotor areas were primarily more activated for up targets than down targets, instances in which the controlling electrodes were required to exhibit the same activity. These observations suggest a direct relationship between premotor areas and the controlling electrode. This pattern can be contrasted with activity patterns observed in prefrontal and posterior parietal areas, which exhibited fairly equivalent activation for both up and down targets. These activity patterns are a reflection of effort or engagement on the part of the subject, and correspondingly decrease as the subject gains task automaticity.

In contrast to ECoG BCIs, noninvasive, EEG-based BCIs typically harness power changes occurring in lower frequencies to achieve control. When EEG subjects are performing motor imagery, the typical frequency band chosen is in the mu-beta range (12–30 Hz), because power changes in this band have been demonstrated to be anticorrelated with movement and motor imagery (33). When using the methods described above to assess whether similar changes were taking place in the mu-beta range as subjects develop experience with an ECoG BCI, we found that, although activity in the mu-beta band was strongly task modulated, it did not undergo the same spatiotemporal changes we observed in HG activity. It is important to note, however, that as the cursor was not being driven by mu-beta activity changes, the impact of this observation to noninvasive BCIs will require further investigation. Details can be found in the *SI Results*.

A logical and necessary extension of these findings is to investigate the roles that subcortical networks play in this same learning process. Previous studies have demonstrated the vital and differential roles of the basal ganglia and cerebellum during motor sequence learning and motor adaptation (27, 32), but the involvement of these subcortical networks in the process of BCI skill acquisition remains an open question, with implications for both fundamental neuroscience and the incremental improvement of BCI frameworks. Recent work performed by Koralek and colleagues (5) has demonstrated that the striatum is involved in and critical to development of proficiency with a BCI in a rat model. This finding was notable, given that effective use of the BCI in that study did not explicitly require recruitment of the motor system outside of motor cortex, yet task performance was degraded in subjects with impaired cortico-striatal interaction. Investigation of the role of these structures in humans will have to be left to other recording modalities, as ECoG provides information only regarding activity near the cortical surface and is subject to spatial undersampling based on the distribution of electrodes.

The distributed dynamics in cortical activity that we demonstrated here have significant implications in development of coadaptive BCI frameworks. Investigators have recognized the need for BCI architectures that accommodate the dynamic nature of the neural signals used as inputs (34, 35), but our findings suggest that a number of the most commonly used methods [e.g., common spatial patterns (36)] may require updating to handle both the spatial and temporal variability in input signals.

The ability to detect correlates of cognitive load during BCI task learning and execution holds great potential for expanding the current limitations of BCIs. Continuous control of a BCI has been demonstrated in two dimensions using ECoG (13) and EEG (14), with degrading performance as users attempt control in more dimensions. Real-time monitoring of cognitive load during BCI skill acquisition would allow for titrated increases in task complexity, potentially facilitating an increase in the total number of dimensions that could be simultaneously controlled, thus allowing operation of more complex devices.

Our results also demonstrate the potential that BCI holds as a technique for probing neural systems in vivo (37). By applying specific task requirements but allowing users to use native learning strategies, we observed that distributed cortical networks are involved in the cognitive phase of BCI skill acquisition, but the degree of involvement of these networks lessens as users transition to automatic execution of the task. Future work will allow us to probe whether activity in this network is renewed when task dynamics are perturbed and the user is required to adapt to novel task conditions.

Materials and Methods

Subjects. Seven subjects with intractable epilepsy were implanted with platinum subdural ECoG electrodes for the purpose of seizure focus localization at Harborview Medical Center and Children's Hospital (Seattle, WA). Research was conducted under the oversite of Seattle Children's Hospital Institutional review board FWA00002443, protocol number 12193. The specific location of the grids was determined based on clinical indication. Electrode grids were constructed of 3-mm-diameter platinum pads spaced at 1 cm center-to-center and embedded in silastic (AdTech). Electrodes were arranged in 8 × 2 or 8 × 8 grids or 1 × 4, 1 × 6, or 1 × 8 strips. Subjects provided informed consent in accordance with the Institutional Review Board's direction, and patient data were anonymized in accordance with Health Insurance Portability and Accountability Act mandate. None of the subjects had previous experience with BCI. Subjects were being observed for seizure focus localization over the course of 4–10 d, although participation in research studies typically did not begin until the third postoperative day.

Recordings. Experimental recordings were performed at the patient's bedside without interrupting the clinical recording systems. Synamps2 (Neuroscan) and g.USBamps (GugerTec) sampled at 1,000 and 1,200 Hz, respectively, were used for recording. Cortical potentials referenced against a scalp electrode and were digitized and processed using the BCI2000 software suite (38), which provided real-time feedback to the user.

BCI Task. BCI paradigms were driven by spectral power changes in a portion of the HG frequency band of a single electrode determined to be modulated by motor imagery by an initial screening task as is described in *SI Materials and Methods*. Only a subset of the HG range was used during online control (~75–100 Hz) for computational tractability and to eliminate the need for real-time notch filtering to reduce line noise harmonics. HG activity was chosen as the control feature as it has been previously postulated that HG activity is a correlate of underlying population level firing rates or coherence in firing (21). All subjects performed the standard right justified box task (23), depicted in Fig. S3. Subjects were instructed to conduct or imagine the movement associated with the chosen controlling electrode when they were presented with a target filling the top half of the right most side of the screen (an up target) and were instructed to rest when they were presented with a target filling the bottom half of the right most side of the screen (a down target). The cursor's vertical velocity was updated every 40 ms and controlled by changes in HG activity at the controlling electrode as calculated by an autoregressive filter using the previous 500 ms of data. This time-variant estimate of HG activity was normalized against 6 s of stored data for the current target type and then mapped to cursor velocity. The normalizer was typically adapting (collecting reference data and updating normalization parameters) only during the first run (18 trials); however, in cases where nonstationarity of the signals showed obvious bias, the normalizer was allowed to recalibrate. Subjects participated in the experiment over the course of multiple days. Duration of the recording sessions was dictated by the subjects' willingness and capability to participate.

Anatomical Labeling. The process by which electrodes were associated with anatomical labels is discussed in *SI Materials and Methods*; in brief, preoperative MRI was coregistered with postoperative CT scans to allow for localization of electrodes with respect to each subject's brain. These electrode locations were then mapped to Talairach coordinates. Anatomical labels were estimated using the HMAT (29) and the Talairach Daemon (30, 31). The HMAT atlas estimates functional areas corresponding to coordinates in Talairach space; as it was constructed from the meta-analysis of 126 motor-based functional MRI studies, it does not include posterior parietal cortex or prefrontal cortex. These areas were identified as consisting of Brodmann areas 7/40 and 8/9/46, respectively. All offline data analyses were performed using custom Matlab software.

Preprocessing. For each implanted grid or strip, signals were common average rereferenced by subtracting the average signal recorded at all electrodes within that grid or strip. This rereferencing was done to eliminate any common noise introduced by activity recorded at the reference electrode. Rereferenced signals were band pass filtered for the HG range (70–200 Hz) using a fourth-order Butterworth filter, and a time-variant estimate of band power was calculated using the square of the magnitude of the Hilbert transform. Data were then log transformed to become approximately normally distributed such that the assumptions made by statistical tests used in further analyses would be met. To account for changes in recording characteristics from session to session, the resulting log-power estimate was z-normalized with respect to the rest samples only per channel and session. Throughout the remainder of this section, the signal resulting from the preprocessing steps will be referred to as HG activation.

Quantification of Average Responses. Trials were divided into four periods: rest, targeting, feedback, and reward. These periods were 1, 2, 3, and 1 s, respectively, resulting in a trial length of 7 s. The HG activation for each of these periods was calculated by averaging all HG activation samples within a given period. Fig. 2B shows mean HG activations during feedback on a trial-by-trial basis for a single subject (S4). Average activation for a given electrode across all trials of a given type was calculated as the mean HG activation during all feedback periods for trials of a given type minus the mean HG activation during all rest periods.

Estimation of HG Activation Separability and Learning States. As was described briefly in *Results*, we sought to quantify differences in activation observed at the controlling electrode that were indicative of a transition from an unlearned to a learned state. The method for this is described at length in *SI Materials and Methods*; briefly, we developed a running estimate of each subject's ability to separate HG activation in up targets as opposed to down targets. This estimate was divided into two distributions such that the distance between those distributions was maximized. The trial that divides these two distributions was used as the transition trial for the comparison of data from the unlearned to the learned state.

Generation of Time Course Activations. The average time course of HG activation for all trials of a given type, like those shown in Fig. 3 B–E, were constructed by averaging the time course of HG activation for a given target and smoothing with a Gaussian kernel [250-ms full width at half maximum (FWHM)]. SEs were calculated before smoothing and were also smoothed using a Gaussian kernel (250-ms FWHM). Trial-by-trial progressions of the time course of HG activation were generated by aligning all trials into an $M \times T$ matrix, where M was the number of trials of a given type and T was the length of a trial in samples. This matrix was then smoothed simultaneously in both dimensions using a 2D Gaussian kernel (300-ms and seven-trial FWHM). Fig. S1 shows average and trial-by-trial time course activations for all subjects.

Comparison of Early and Late Activations. Shifts in activation from early to late learning states were calculated in much the same way that average responses were calculated, except that instead of comparing activation during the feedback period to the activation during rest, activation during the feedback period of early trials was compared with activation during the feedback period of late trials. These changes were summarized by subtracting average early activation from average late activation such that a negative change as is shown in Fig. 4B for the majority of PMd/PMv, PFC, and PPC electrodes is indicative of a decrease in activation from early to late trials.

1. Fetz EE (1969) Operant conditioning of cortical unit activity. *Science* 163(3870): 955–958.
2. Fetz EE, Baker MA (1973) Operantly conditioned patterns on precentral unit activity and correlated responses in adjacent cells and contralateral muscles. *J Neurophysiol* 36(2):179–204.
3. Kennedy PR, Bakay RA (1998) Restoration of neural output from a paralyzed patient by a direct brain connection. *Neuroreport* 9(8):1707–1711.
4. Chapin JK, Moxon KA, Markowitz RS, Nicolelis MA (1999) Real-time control of a robot arm using simultaneously recorded neurons in the motor cortex. *Nat Neurosci* 2(7): 664–670.
5. Koralek AC, Jin X, Long JD, 2nd, Costa RM, Carmena JM (2012) Corticostriatal plasticity is necessary for learning intentional neuroprosthetic skills. *Nature* 483(7389):331–335.
6. Leuthardt EC, Schalk G, Wolpaw JR, Ojemann JG, Moran DW (2004) A brain-computer interface using electrocorticographic signals in humans. *J Neural Eng* 1(2):63–71.
7. Wolpaw JR, McFarland DJ, Neat GW, Forneris CA (1991) An EEG-based brain-computer interface for cursor control. *Electroencephalogr Clin Neurophysiol* 78(3):252–259.
8. Fabiani GE, McFarland DJ, Wolpaw JR, Pfurtscheller G (2004) Conversion of EEG activity into cursor movement by a brain-computer interface (BCI). *IEEE Trans Neural Syst Rehab Eng* 12(3):331–338.
9. Miller KJ, et al. (2010) Cortical activity during motor execution, motor imagery, and imagery-based online feedback. *Proc Natl Acad Sci USA* 107(9):4430–4435.
10. Hochberg LR, et al. (2006) Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature* 442(7099):164–171.
11. Carmena JM, et al. (2003) Learning to control a brain-machine interface for reaching and grasping by primates. *PLoS Biol* 1(2):E42.
12. Ganguly K, Carmena JM (2009) Emergence of a stable cortical map for neuroprosthetic control. *PLoS Biol* 7(7):e1000153.
13. Schalk G, et al. (2008) Two-dimensional movement control using electrocorticographic signals in humans. *J Neural Eng* 5(1):75–84.
14. Wolpaw JR, McFarland DJ (2004) Control of a two-dimensional movement signal by a noninvasive brain-computer interface in humans. *Proc Natl Acad Sci USA* 101(51): 17849–17854.
15. Moritz CT, Perlmutter SI, Fetz EE (2008) Direct control of paralysed muscles by cortical neurons. *Nature* 456(7222):639–642.
16. Green AM, Kalaska JF (2011) Learning to move machines with the mind. *Trends Neurosci* 34(2):61–75.
17. Wolpaw JR (2007) Brain-computer interfaces as new brain output pathways. *J Physiol* 579(Pt 3):613–619.
18. Fitts PM (1954) The information capacity of the human motor system in controlling the amplitude of movement. *J Exp Psychol* 47(6):381–391.
19. Doyon J, Ungerleider L (2002) *Neuropsychology of Memory* (Guilford Press, New York), 3rd Ed.
20. van Mier H, et al. (1998) Changes in brain activity during motor learning measured with PET: Effects of hand of performance and practice. *J Neurophysiol* 80(4): 2177–2199.
21. Ray S, Crone NE, Niebur E, Franaszczuk PJ, Hsiao SS (2008) Neural correlates of high-gamma oscillations (60-200 Hz) in macaque local field potentials and their potential implications in electrocorticography. *J Neurosci* 28(45):11526–11536.
22. Crone NE, Miglioretti DL, Gordon B, Lesser RP (1998) Functional mapping of human sensorimotor cortex with electrocorticographic spectral analysis. II. Event-related synchronization in the gamma band. *Brain* 121(12):2301–2315.
23. Wolpaw JR, McFarland DJ, Vaughan TM, Schalk G (2003) The Wadsworth Center brain-computer interface (BCI) research and development program. *IEEE Trans Neural Syst Rehab Eng* 11(2):204–207.
24. Schlaug G, Knorr U, Seitz R (1994) Inter-subject variability of cerebral activations in acquiring a motor skill: A study with positron emission tomography. *Exp Brain Res* 98(3):523–534.
25. Jenkins IH, Brooks DJ, Nixon PD, Frackowiak RS, Passingham RE (1994) Motor sequence learning: A study with positron emission tomography. *J Neurosci* 14(6):3775–3790.
26. Buneo CA, Andersen RA (2006) The posterior parietal cortex: Sensorimotor interface for the planning and online control of visually guided movements. *Neuropsychologia* 44(13):2594–2606.
27. Doyon J, Benali H (2005) Reorganization and plasticity in the adult brain during learning of motor skills. *Curr Opin Neurobiol* 15(2):161–167.
28. Johnson L, et al. (2012) Sleep spindles are locally modulated by training on a brain-computer interface. *Proc Natl Acad Sci USA* 109(45):18583–18588.
29. Mayka MA, Corcos DM, Leurgans SE, Vaillancourt DE (2006) Three-dimensional locations and boundaries of motor and premotor cortices as defined by functional brain imaging: a meta-analysis. *Neuroimage* 31(4):1453–1474.
30. Lancaster JL, et al. (2000) Automated Talairach atlas labels for functional brain mapping. *Hum Brain Mapp* 10(3):120–131.
31. Lancaster JL, et al. (1997) Automated labeling of the human brain: A preliminary report on the development and evaluation of a forward-transform method. *Hum Brain Mapp* 5(4):238–242.
32. Hikosaka O, Nakamura K, Sakai K, Nakahara H (2002) Central mechanisms of motor skill learning. *Curr Opin Neurobiol* 12(2):217–222.
33. McFarland DJ, Miner LA, Vaughan TM, Wolpaw JR (2000) Mu and beta rhythm topographies during motor imagery and actual movements. *Brain Topogr* 12(3):177–186.
34. Vidaurre C, Sannelli C, Müller K-R, Blankertz B (2010) Machine-Learning-Based Co-adaptive Calibration for Brain-Computer Interfaces. *Neural Comput* 23:791–816.
35. Vidaurre C, Sannelli C, Müller K-R, Blankertz B (2011) Co-adaptive calibration to improve BCI efficiency. *J Neural Eng* 8(2):025009.
36. Müller-Gerking J, Pfurtscheller G, Flyvbjerg H (1999) Designing optimal spatial filters for single-trial EEG classification in a movement task. *Clin Neurophysiol* 110(5): 787–798.
37. Moran D (2010) Evolution of brain-computer interface: Action potentials, local field potentials and electrocorticograms. *Curr Opin Neurobiol* 20(6):741–745.
38. Schalk G, McFarland DJ, Hinterberger T, Birbaumer N, Wolpaw JR (2004) BCI2000: A general-purpose brain-computer interface (BCI) system. *IEEE Trans Biomed Eng* 51(6): 1034–1043.