

Neuroadaptive technology enables implicit cursor control based on medial prefrontal cortex activity

Thorsten O. Zander^{a,b,1,2}, Laurens R. Krol^{a,b,1}, Niels P. Birbaumer^{c,d}, and Klaus Gramann^{a,e}

^aBiological Psychology and Neuroergonomics, Technische Universität Berlin, 10623 Berlin, Germany; ^bTeam PhyPA (Physiological Parameters for Adaptation) Technische Universität Berlin, 10623 Berlin, Germany; ^cInstitute for Medical Psychology and Behavioural Neurobiology, University Tübingen, 72076 Tübingen, Germany; ^dWyss Center for Bio and Neuroengineering, 1202 Geneva, Switzerland; and ^eCenter for Advanced Neurological Engineering, University of California, San Diego, CA 92093

Edited by Terrence J. Sejnowski, Salk Institute for Biological Studies, La Jolla, CA, and approved November 14, 2016 (received for review March 30, 2016)

The effectiveness of today's human-machine interaction is limited by a communication bottleneck as operators are required to translate high-level concepts into a machine-mandated sequence of instructions. In contrast, we demonstrate effective, goal-oriented control of a computer system without any form of explicit communication from the human operator. Instead, the system generated the necessary input itself, based on real-time analysis of brain activity. Specific brain responses were evoked by violating the operators' expectations to varying degrees. The evoked brain activity demonstrated detectable differences reflecting congruency with or deviations from the operators' expectations. Real-time analysis of this activity was used to build a user model of those expectations, thus representing the optimal (expected) state as perceived by the operator. Based on this model, which was continuously updated, the computer automatically adapted itself to the expectations of its operator. Further analyses showed this evoked activity to originate from the medial prefrontal cortex and to exhibit a linear correspondence to the degree of expectation violation. These findings extend our understanding of human predictive coding and provide evidence that the information used to generate the user model is task-specific and reflects goal congruency. This paper demonstrates a form of interaction without any explicit input by the operator, enabling computer systems to become neuroadaptive, that is, to automatically adapt to specific aspects of their operator's mindset. Neuroadaptive technology significantly widens the communication bottleneck and has the potential to fundamentally change the way we interact with technology.

human-computer interaction | passive brain-computer interfaces | electroencephalogram | predictive coding | neuroadaptive technology

In the European Union, 96% of enterprises rely on computers for their productivity (1). Advances in human-computer interaction (HCI), concerning the effective, efficient, and satisfying use of computer systems, may thus carry great societal benefits, e.g., in terms of productivity. However, although interaction techniques have become increasingly user-friendly—e.g., from punch cards to touch screens—they still depend on the user (operator) to translate their original thought or intention into a sequence of small, explicit commands (2). This translational step, where the human operator must ultimately obey the machine's logic, presents both a communication bottleneck and a source of potential error (3). At the same time, the computer has practically no limitation to the amount of information it can communicate, and is not as adaptable as its user. In these aspects, present-day HCI is asymmetrical (4). Comparing this to human-human interaction, Fischer (5) emphasizes the importance of a shared understanding of the situation and an understanding of the communication partner. In this sense, for a computer system to “understand” its user, it needs a model of that user—a source of information that goes beyond the explicitly given commands. On the basis of such a model, a computer system could adapt its behavior to better suit the current mode of the user (5). This could help alleviate the issue of asymmetry. Relevant information to that end concerns the user's intentions, subjective interpretations, and emotions.

Four decades of developments in brain-computer interfaces (BCIs) (6, 7) have yielded a set of methods that may be used to obtain such information in real time, provided that this information is detectably reflected in brain activity. Specifically, BCIs can detect in real time changes in the electroencephalogram (EEG) and translate these changes into control signals, in line with the principles of physiological computing (8). A subgroup of BCIs, so-called “passive BCIs” (pBCIs) (9), focuses on monitoring otherwise covert aspects of the user state (10) during an ongoing HCI. Neurophysiological correlates of the above-mentioned aspects can be detected and interpreted in the context of the interaction, and can be used to inform the computer about relevant changes in the user's cognition and affect. Using pBCI, thus, a computer can in fact acquire information about its operator other than the explicitly given commands. As such, neurophysiological activity can induce appropriate changes in the machine in real time, essentially serving as an implicit command, without requiring the user to exert any conscious effort in communicating to the computer (9).

Previous and present-day BCI systems use information derived about the user state in only an ad hoc fashion: momentary information derived by the BCI from the EEG is directly interpreted as a specific user intention (11, 12), situational interpretation (13), or a change in the cognitive (14) or affective (15) state. These implementations represent one-to-one mappings of user states to machine behavior. We propose, however, that a machine using

Significance

The human brain continuously and automatically processes information concerning its internal and external context. We demonstrate the elicitation and subsequent detection and decoding of such “automatic interpretations” by means of context-sensitive probes in an ongoing human-computer interaction. Through a sequence of such probe-interpretation cycles, the computer accumulates responses over time to model the operator's cognition, even without that person being aware of it. This brings human cognition directly into the human-computer interaction loop, expanding traditional notions of “interaction.” The concept introduces neuroadaptive technology—technology which automatically adapts to an estimate of its operator's mindset. This technology bears relevance to autoadaptive experimental designs, and opens up paradigm-shifting possibilities for human-machine systems in general.

Author contributions: T.O.Z., L.R.K., and K.G. designed research; L.R.K. performed research; T.O.Z. contributed new reagents/analytic tools; T.O.Z., L.R.K., and K.G. analyzed data; and T.O.Z., L.R.K., and N.P.B. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

Freely available online through the PNAS open access option.

¹T.O.Z. and L.R.K. contributed equally to this work.

²To whom correspondence should be addressed. Email: tzander@gmail.com.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1605155114/-DCSupplemental.

pBCI can detect both general user states and transient, event-related responses, and can use these to continuously and accumulatively learn about its operator. Specifically, we propose that the machine collates the neurophysiological responses of its operator (i.e., implicit inputs) and coregisters them against the events and contexts that evoked them. This allows the machine to build and continuously update a specific and context-sensitive model of that operator (10). The goal is to combine the information gathered from multiple responses to different events to gain insights into higher-level aspects of the operator's cognition.

One aspect of higher cognition that may be inferred in this manner is described by the theory of human predictive coding. Predictive coding holds that there exists a continuous, automatic prediction of future (neuronal) events, as well as a continuous comparison of those predictions with their corresponding final perception (16–18). Discrepancies resulting from these comparisons inform the brain of the correctness of its predictions and actions, providing a fundamental mechanism—prediction error minimization—to shape and optimize behavior. The corresponding predictive signals are assumed to be carried by the dopaminergic system. Changes in the continuous evaluation of events and actions lead to changes in the dopaminergic input to the anterior cingulate cortex, (dis)inhibiting its neurons and eliciting a detectable response (19). Predictions of what is expected to happen, in this sense, relate closely to what is intended to happen. This makes the correlates of predictive coding a fundamental source of information concerning user intention—an aspect of the user's cognition that is highly relevant to HCI.

In this paper, we demonstrate that by collating passive BCI output and context information, it is possible to develop, step by step, a user model that accurately reflects correlates of predictive coding and reveals task-relevant subjective intent.

Specifically, we demonstrate that a user model can be developed and used to guide a computer cursor toward the intended target, without participants being aware of having communicated any such information. Using a passive BCI system, the participant's situational interpretations of cursor movements were classified and interpreted, in the given context, as directional preferences. A user model was generated to represent these context-dependent directional preferences, and this model was then used to guide the cursor toward the intended movement direction.

Results

The experimental paradigm involved a form of cursor control. The cursor moved discretely over the nodes of a (4×4 or in later stages 6×6) grid. For each movement the cursor could travel up

to eight directions, horizontally, vertically, and diagonally, to one of the adjacent nodes. Each movement served both to move the cursor and to elicit a neurophysiological response, reflecting the subjective correctness of that movement. In essence, each movement thus also served as a probe for information. One of the grid's corners was designated the target. For each movement, it could thus be determined at what angle of deviance relative to the target the cursor had moved. This was used for an objective interpretation of the cursor's behavior. We describe the paradigm in detail in *SI Appendix*.

The event-related potential (ERP) following each probe (i.e., each cursor movement) is shown in Fig. 1A. A one-way analysis of variance of the systematic peak differences around 180 ms indicated a significant main effect of angular deviance from the target direction on peak amplitude [$F(7,126) = 47.243$, $P < 0.001$]. Specifically, the peak amplitudes (Fig. 1B, upper curve) differed significantly ($P < 0.001$) between both the lowest and the highest angular deviation from the target direction as used by the classifier. In between, the peak amplitudes scaled linearly with angular deviance, as fitted by a linear regression model using each group's mean angular deviance as a predictor (slope coefficient $b = -0.0035$, $F = 45.28$, $P < 0.001$; $R^2 = 0.33$). Further posthoc comparisons corrected for false discovery rate additionally indicated that significant differences between adjacent groups ($P < 0.05$) were found mostly for groups of lower angular deviance, whereas differences between the three largest-deviance groups (124° and up) were not significant. The results of all posthoc comparisons are listed in *SI Appendix*, Table S3. In summary, the probe elicited systematic variations in event-related amplitudes, depending on the goal congruency of the presented stimulus.

To enable real-time detection of the individual, single-trial neuroelectric responses, we calibrated a discriminative classification system. Calibration was based on two classes of responses representing the extreme ends of the spectrum, with angular deviances of 0° making up the one class, and angular deviances of $\geq 135^\circ$ the other (Fig. 2A). This classification system used a subject-dependent linear combination of all 64 available channels, taking into account full scalp information. It automatically generated appropriate spatial filters for eight 50-ms time windows—starting at 50 ms after stimulus presentation—using supervised machine learning and linear discriminant analysis. This set of filters weighted each electrode in each time window, depending on its relevance to classification. The recorded signal projected through them thus allowed an optimal discrimination between the two classes. This projected recorded signal—from all 64 channels, between 50 and 450 ms after stimulus presentation—defined the

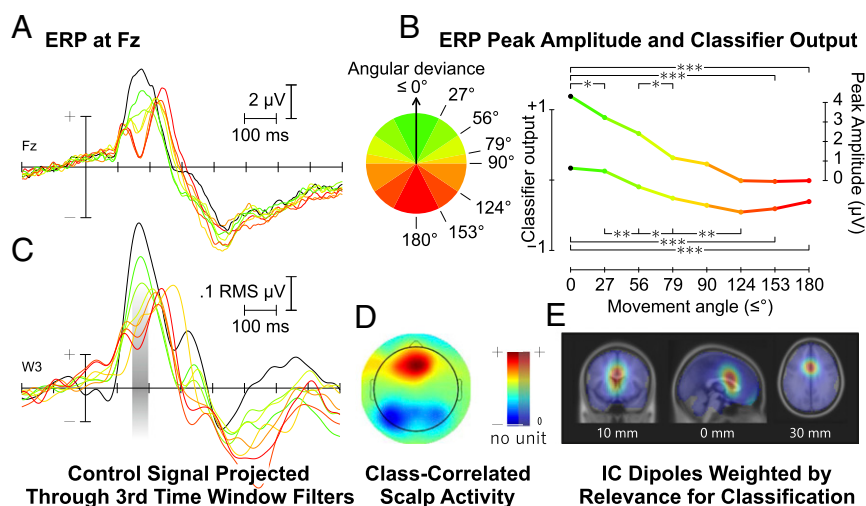
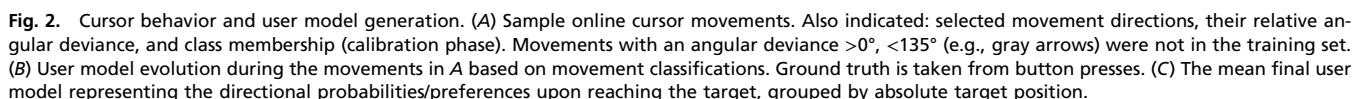


Fig. 1. Neurophysiological analysis. (A) Grand average ERP at Fz ($n = 19$) time-locked to cursor movement, divided into eight groups depending on angular deviance. (B) Peak amplitudes around 180 ms for the ERP in A, and mean classifier output for cursor movements sorted by angular deviance with selected significant differences indicated ($***P < 0.001$, $**P < 0.01$, $*P < 0.05$). (C) Grand average ERP ($n = 19$) projected through the sources focused on in the third time window (150–200 ms; indicated in gray). (D) Scalp map of difference-between-classes activity that contributed to classification in the third time window. (E) Source localization for the third time window.



The signal thus carried task-relevant information. For a true test of this signal's single-trial reflection of individual judgments of cursor movements, and thus its usefulness in creating a user model describing subjective intent, we created a closed-loop, online version of the original offline paradigm. Following each single cursor movement, an individually calibrated classification system classified the evoked response. The extracted information was used for reinforcement learning on the side of the cursor (20), modifying the probabilities of upcoming cursor movements

Online application thus significantly increased the goal congruency, confirming that the signal the classification system focused on was situationally relevant. Although the cursor only made binary interpretations of the classifier's output, this output was continuous: a scale, from -1 to $+1$, correlating to the movements' degree of goal congruency. This is illustrated in Fig. 1B (lower curve). The

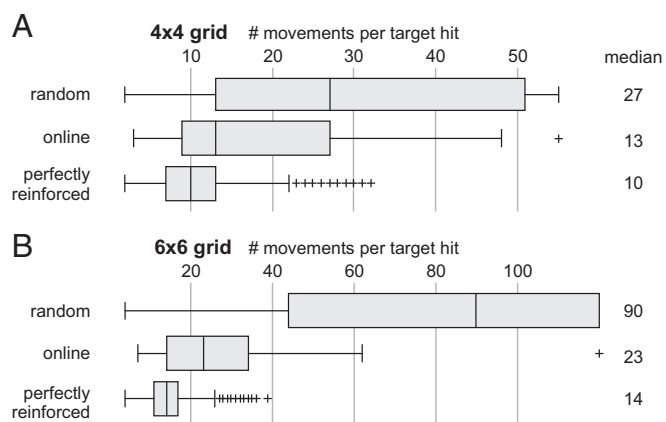


Fig. 3. Performance measure distributions for nonsupported, online, and perfectly reinforced cursor movements on the two online grid sizes. (A) Performance on the 4×4 grids. (B) Performance on the 6×6 grids. All differences between the three conditions are significant ($P < 0.025$). Whiskers cover $\pm 2.7\sigma$.

classifier output differs significantly ($P < 0.001$) between the classes used by the classifier. In between, the classifier output scaled linearly with angular deviance, as fitted by a linear regression model using each group's mean angular deviance as predictor ($b = 0.0035$, $F = 295.42$, $P < 0.001$; $R^2 = 0.76$). See *SI Appendix, Table S4* for further comparisons.

Even though the linearly scaled information was not taken into account, binary classifications still resulted in a graded user model, describing the appropriateness of the different cursor movements depending on the intended target's position. To illustrate this, Fig. 2*A* visualizes the cursor's movements over a grid during one of the online runs with the target in the southwest corner. Fig. 2*B* shows how the individual directional preferences/probabilities in the user model are updated after every cursor movement, showing the progression toward a clearly identified preference for the southwest corner. Fig. 2*C* illustrates the mean final user models for all participants for the four different target positions. It is clear that the user models accurately represent the intended target position. The mean final user model across all participants is illustrated in *SI Appendix, Fig. S4*, with statistics in *SI Appendix, Table S2*. *SI Appendix, Fig. S9* shows one more example of online cursor behavior.

Discussion

We have demonstrated that binary classifications of subjective interpretations of cursor movements can be aggregated into a user model reflecting, in the given context, directional intent. Based on this model, the cursor was effectively guided toward the target. Participants were not aware of their influence on the cursor. Although not used explicitly in this study, analysis shows that more fine-grained information may be available in the elicited responses, encoded in the linear dependency of the response on the angular deviance.

An approach combining independent component analysis (ICA, refs. 21, 22; *Materials and Methods*), supervised machine learning, and higher-order statistics not only gave insight into the individual single-trial responses, but also enabled error-minimized source localization and signal back-projection as well as real-time single-trial analysis. These characteristics could be used to validate the classification system as well as the user model. Firstly, the online application of the classification system increased the paradigm's goal congruency to near the optimum: The gap from no reinforcement to optimal reinforcement could be bridged by over 80% using the presented classification system, both on the 4×4 grid and on the theretofore unseen 6×6 grid. This significant reduction in the number of steps required to

reach a target provides evidence that the classification of brain responses following a cursor movement was based on task-relevant information: Each movement indeed elicited a response enabling the identification of subjective directional preferences.

Secondly, the neurophysiological analysis, based on the classification system's filter set, revealed that the underlying signal predominantly stemmed from the mPFC, and reflected the experimental paradigm's logic. Further interpretation of the neurophysiological response points to its likely generator process. Given the signal's time course, its localization, and the evoking stimuli, the response is in line with the framework of predictive coding. We hypothesize that in the present study, participants consistently predicted—for lack of information that would indicate otherwise—that the cursor would perform the only action that would have been appropriate, i.e., that it would move in the direction of the target. Interestingly, however, our findings imply an extension of the general framework of predictive coding. A focus on “negative” signals is central to current interpretations and findings related to predictive coding: Indications of discrepancies, of prediction errors, are seen as central to learning by reinforcement, in turn explaining the large range of rich human behavior and intelligence (23). The sensitivity of the ERP amplitude to the quality of the cursor movements, however, seems to indicate that neural activity generated within the mPFC provides a range of graded responses to both positive and negative movements. In this context, these reflect the observer's directional preferences, modeling an important, task-relevant factor of their subjective cognition. This points to a continuous response range within the mPFC that not only detects deviations from a predicted event to adapt future behavior, but also confirms correct predictions to reinforce adequate behavior or sharpen perceptive hypotheses. This activity thus reflects complex aspects of the operator's cognition, and can be highly informative for external systems that have access to it.

Taken together, these results demonstrate effective cursor control through implicit interaction: While participants were unaware of having any influence on the cursor, the presented stimuli elicited informative neuronal responses that allowed the system to establish a user model from which the participants' intentions could be derived. The computer system adapted its behavior to fit this model—thus becoming neuroadaptive. The necessary information could also have been provided explicitly and volitionally, but conscious interpretation can involve any number of additional considerations and processes (e.g., resolving competing interpretations from different judgment strategies), and would require an explicit decision as well as its translation into a command to inform the machine. Direct access to such interpretations circumvents these time-consuming and effortful steps, proving advantageous even with simple binary decisions, as implemented here. Communicating more fine-grained information, as seems also to be available, would be even more difficult using traditional input techniques, but equally effortless using the method presented here. As such, neuroadaptive technology based on passive BCI bypasses the communication bottleneck present in traditional HCI, effectively widening it by allowing interaction to take place through implicit channels. This decreases the asymmetry present in current HCI paradigms.

At this point, we would like to speculate about possible implications and future extensions of the findings and the line of thought presented here. Our current method essentially quantified subjective directional preferences, supplying a single value that indicated, in the given context, whether a person interpreted a single cursor movement as being supportive of reaching the target or not. This can be seen as a real-time assessment of subjective satisfaction/dissatisfaction with the presented probe stimulus, thus allowing the generation of a user model representing subjective intent. Interestingly, one can imagine a computer system that intelligently decides what probe to present, to gather information. A system with an incomplete user model, for example, could present a probe to gauge the user's response and thus gather the missing

information. Such an act of active learning (24) would invert the traditional HCI cycle: The probe may be understood as a command—a request for feedback—direct from the machine to the user's brain, inducing its own interpretation, which results in the machine indeed receiving the requested feedback. In the demonstration presented here, each cursor movement served as such a probe, and allowed the gradual development of a user model, but a more intelligent selection of probes may improve the system's efficiency.

With such a fundamental process as for example predictive coding underlying a neuroadaptive system, a large scope of potential applications can be imagined. Any process or path that can be divided into a sequence of one-dimensional (e.g., positive-negative) responses could potentially be covered implicitly (*SI Appendix, Fig. S10*). And, as human predictive coding shows, a great deal can be achieved based on such information using, for example, the relatively simple process of reinforcement. The 2D grid used here could be replaced by any n-dimensional space representing different system parameters. It is tempting to envision how such neuroadaptive systems could transform work and leisure activities in everyday settings. An implementation analogous to the current demonstration (but going beyond the currently presented results), using affective interpretations rather than cognitive ones (25), could be an adaptive, open-ended electronic book. While reading, the reader would interpret the story as it unfolded, thus automatically responding to events with a detectable affective state. Based on what the reader apparently finds enjoyable, a neuroadaptive system could then change the content of subsequent pages. With a sequence of such adaptations, the story is gradually steered in the reader's preferred direction. However, the reader would not actively be directing the story, and would not even need to be aware of the system's existence.

Similarly, the general method demonstrated here is of value to neuroadaptive experimental paradigms. Such paradigms can use the real-time feedback supplied by the classification system to adapt to individual strategies, rather than enforcing a uniform logic over all participants. Probe stimuli can be used to first inspect the subjective relevance of different experimental aspects, for example, and then adaptively go into detail, presenting more fine-grained nuances of these aspects, to model how they influence the brain dynamics of that individual.

A word of caution is in order. Neuroadaptive systems can be said to be systems with an agenda, having a goal of their own (8). By autonomously initiating each interaction cycle using a specifically selected probe stimulus, they would be in a position to "guide" the interaction such that specific information can be gathered, and to change the interactive experience based on that or other information. When designing such systems, care should be taken that this agenda is not adverse to the user's intention. Furthermore, the fact that it can rely on automatic, unconscious responses represents a potential danger to informed consent. Users should always have access to full information concerning the system's goals and actions.

The benefits of closed-loop neuroadaptive technology, however, may be vast. It enables experimental paradigms to model and adapt to relevant individual aspects in real time. For technology in general, this concept could represent a paradigm shift in that it skips translational effort, grants the machine initiative and agency, and may even function outside of conscious awareness. This offers designers the prospect to completely rethink the notion of interaction and the possibilities offered by it. In almost a century of neurophysiological research, a number of correlates of cognitive processes have been identified in the EEG, some of which can already be detected in single trials using passive BCI methodology (as, e.g., refs. 11, 13–15, 26, 27). We are looking forward to investigating which of these could be meaningfully elicited and interpreted to inform personalized user models, as per the concept of neuroadaptive technology presented here. Commercial systems and experimental paradigms specifically

designed for this type of implicit interaction—a cybernetic convergence of human and machine intelligence—could offer new functionality and scientific results we cannot currently foresee.

Materials and Methods

Experimental Procedure and Setup. All participants were informed of the nature of the experiment and the recording and anonymization procedures before signing a consent form. The Ethics Committee of the Department of Psychology and Ergonomics at the Technische Universität Berlin approved the experiment and the procedures.

A gray grid was shown on a black background, with a red target node indicated in one of the grid's corners, and a red circular cursor visible on one of the nodes (Fig. 2A and *SI Appendix, Fig. S1*). The cursor's starting position on each grid was one node away from the corner opposite the target's, in a straight line to the target. In each trial, the cursor moved from its current node to one of the adjacent nodes. A 1-s animation within the cursor served as a countdown. The cursor would then instantaneously jump to the next node, highlighting in white its new position and the grid line between the two nodes. This configuration remained visible for 1 s. Following that, the highlights disappeared and the cursor would remain at its new position for 1 s more before the next trial. *Movie S4* shows animated stimuli as seen by the participant.

Throughout the experiment, participants were instructed to judge each individual cursor movement as either "acceptable" or "not acceptable" with respect to reaching the target, and to indicate their judgment by pressing either "v" or "b," respectively, on a computer keyboard using the index finger of one hand. These button presses were logged by the system but were not used as input.

EEG was recorded using 64 active Ag/AgCl electrodes mounted according to the extended 10–20 system. The signal was sampled at 500 Hz and amplified using BrainAmp DC amplifiers (Brain Products GmbH).

Participants first performed 5 blocks of 120 trials on grids of 4×4 nodes. If the target had not been reached after 55 trials in one grid, a new grid was started. Fifty-five is twice the median number of random movements required to reach a target on a 4×4 grid. The EEG recorded during these five blocks served to calibrate the classifier, as discussed below. In online sessions, this classifier was then applied to one more block of 120 trials on 4×4 grids, and one last block of 120 trials on 6×6 grids. No maximum number of trials other than the block's length was set for the 6×6 online blocks.

During calibration blocks, the cursor moved randomly. During online application of the pBCI, the directional probabilities were altered based on the classification of each movement as either "correct" or "incorrect," biasing the cursor to repeat movements classified as correct.

A total of 19 participants participated in this study, with an average age of 25.4 ± 3.4 . Seven were female. All had normal or corrected to normal vision. The first 3 only performed offline calibration trials, whereas the following 16 additionally performed online trials.

Additional details are provided in *SI Appendix*.

Classifier. A classifier was individually calibrated on data from the initial 600-trial recording of random cursor movements. Movements with an angular deviance of 0° were labeled as class 1, and movements with an absolute deviance of 135° or more were labeled as class 2. A regularized linear discriminant analysis classifier was trained to separate classes (13).

The open-source toolbox BCILAB (28) version 1.01 was used to define and implement the pBCI. Features were extracted through the windowed means approach, using the average amplitudes of eight sequential time windows of 50 ms each, between 50 and 450 ms after each cursor movement (13). For this feature extraction, the data were first resampled at 100 Hz and band-pass filtered using fast Fourier transform from 0.1 to 15 Hz. Ensuring that the features were independent and identically distributed, a 5×5 -times nested cross-validation with margins of 5 was used to select the shrinkage regularization parameter, and to generate estimates of the classification system's online reliability.

Additional details are provided in *SI Appendix*.

Identifying Scalp Projections. Following Haufe et al. (29), linear discriminant analysis (LDA) patterns $A = (a_j)_j$ were generated for each participant from the LDA filters $M = (m_j)_j$ originally used for online classification by conjugation with the features' covariance matrix C : $A = CMC^{-1}$. Spatial interpretation of these patterns for each time window reflects a mixture of scalp activations related to discriminative source activity $\hat{A} = (\hat{a}_j)_j$, and class-invariant noise representation N , with $A = \hat{A} + N$. The latter was filtered out by weighting each pattern entry a_j with the correlation of its associated feature activity vector over trials F_j to the binary vector of true class labels

$L: \hat{a}_j = \text{corr}(F_j, L) \cdot a_j$ The resulting correlated pattern $\hat{A} = (\hat{a}_j)_j$ can be visualized by topographic plots for each time window, as in Fig. 1D.

Additional details are provided in [SI Appendix](#).

Localization. The classification model used was a multivariate approach, an LDA, optimized for the discriminability of the extracted features between classes. Each feature represents data at a single sensor for one of the chosen time windows. Hence, applying the methodology recently introduced by Haufe et al. (29) interprets the classification model at sensor level, and reveals further insight into the relevant underlying processes.

To identify the sources producing the signal, the backward model, i.e., the LDA filter, was combined with an ICA. The ICA unmixing matrix $W = (I_1, I_2, \dots, I_n)$ was determined on manually cleaned data for each participant by using the Adaptive Mixture ICA (AMICA) Toolbox (30), such that $s = Wx$, where s represents the source activation related to a given scalp activation x . For each time window, the relevance for classification R_i of each independent component I_i can then be determined by distributing the LDA filter weights to the independent components via W , weighted by two factors. The first factor compensates for the amplitude alignment of the LDA filter weights to the feature amplitudes. It is determined by calculating the variance over trials of the feature \hat{F}_i extracted from the time series of the independent component: $V_i = \text{var}(\hat{F}_i)$. A second weight is determined for filtering out noise representations by weighting the independent components with the correlation of their feature activity to the true class labels (as described above for electrode activity): $R_i = V_i \cdot \text{corr}(\hat{F}_i, L) \cdot W$.

To then localize these sources, equivalent dipole models that describe the most likely position of the source in a standard head model were identified for selected components by using the EEGLAB plug-in DIPFIT 2.x (31). Components were selected by a threshold criterion for residual variance (RV) of the dipole model ($RV < 0.15$) and visual inspection of the activation spectra, time courses, and scalp topographies. Only dipolar components clearly reflecting cortical, ocular, or muscular activity were included in the analysis. For every time window, each of the 371 resulting dipoles was weighted by the relevance R_i of its associated independent component. The areas of high relevance were then described by a weighted dipole density plot using the EEGLAB plug-in dipoleDensity (32) by plotting the dipole density per cubic millimeter weighted by the relevance R_i of each included dipole with a smoothing kernel of 12 mm.

Movie S2 shows the results of this analysis for the full time course under investigation.

The above-mentioned process of dipole selection did not markedly influence the analysis. Compared with all 1,191 dipoles and averaged over the 8 time windows, the 820 rejected dipoles (68.8%) carried 7.5% of the weights distributed by the classifier. Relative to all 1,191 dipoles, a total of 87 dipoles received a relevance weight larger than an SD of 1 in at least one of the time windows. These 7.3% of dipoles carried 77.8% of the total weight distributed by the classifier. Four of these highly weighted dipoles (4.6%) were rejected in the process explained above and not included in the analysis. These four represent 1.7% of the weight included in the analysis. Three belonged to the same subject.

Additional details are provided in [SI Appendix](#).

Performance Measures and Statistical Methods. Cursor performance was operationalized as the number of movements required to reach the target. Only completed grids are included in the analysis, i.e., either when the target was hit or the maximum number of trials was reached. Online, this latter event occurred a total of seven times to seven participants on the 4×4 grids, and to two participants on the larger grids. Out of 120 cursor movements per grid size per participant, this resulted in 88 pBCI-supported target hits for the smaller grid, and 47 for the larger one.

Random cursor movement data are nonparametrically distributed and vary greatly. Therefore, we used a resampling approach where the available sample of pBCI-supported measures was repeatedly compared with a new random sample of the same size of nonsupported performance measures, using a Wilcoxon rank-sum test. Out of 50,000 such comparisons, 98% of tests were significant at $\alpha = 0.025$ for the smaller grid; for the larger grid, 100% of tests were significant at this level.

The perfect performance was simulated by automatically reinforcing the cursor as in the online sessions, with all movements with an angular deviance of less than 45° reinforced positively and all others negatively. The same procedure as above yielded significant differences to the pBCI-supported measures for both the 4×4 grid (99.9% of tests significant at $\alpha = 0.025$) and the 6×6 grid (100%).

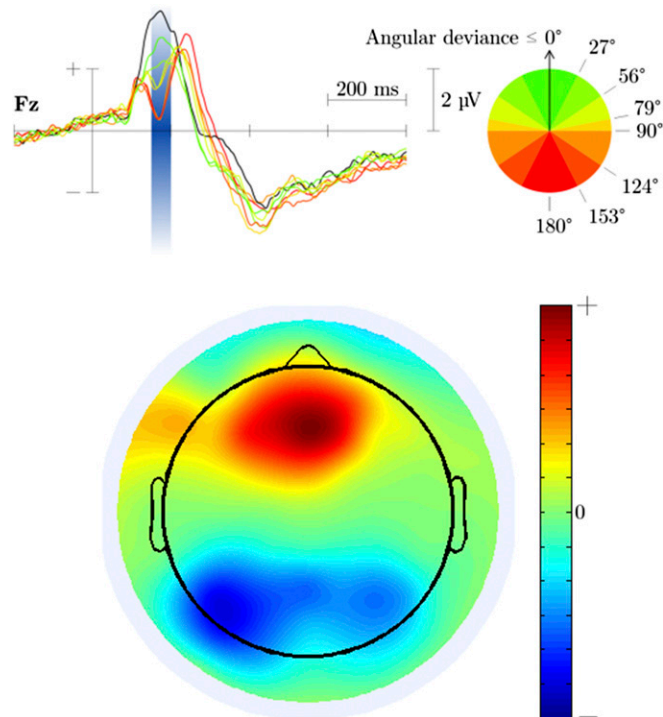
Additional details are provided in [SI Appendix](#).

ACKNOWLEDGMENTS. We thank S. Haufe, S. Makeig, M. Miyakoshi, J. Faller, and C. Wienrich for contributions to the various stages of the research process and C. Kothe, S. Jatzev, M. Rötting, and Brain Products GmbH for support. Part of this work was supported by the Human Factors and Ergonomics Society Europe Chapter VRC Corporation Exchange Grant (Reference EC/2014.11).

- Eurostat (2015). Computers and the Internet: Enterprises - Summary of EU Aggregates (NACE Rev. 2 activity) [isoc_ci_eu_en2]. Available at ec.europa.eu/eurostat/en/web/products-datasets/-/ISOC_CI_EU_EN2. Accessed January 11, 2016.
- Rogers Y, Sharp H, Preece J (2011) *Interaction Design: Beyond Human-Computer Interaction* (Wiley, Chichester), 4th Ed.
- Tufte ER (1990) *Envisioning Information* (Graphics Press, Cheshire, CT).
- Suchman LA (1987) *Plans and Situated Actions: The Problem of Human-Machine Communication* (Cambridge Univ Press, Cambridge, UK).
- Fischer G (2001) User modeling in human-computer interaction. *User Model. User-Adap.* 11(1–2):65–86.
- Vidal JJ (1973) Toward direct brain-computer communication. *Annu Rev Biophys Bioeng* 2(1):157–180.
- Wolpaw J, Wolpaw EW, eds (2012) *Brain-Computer Interfaces: Principles and Practice* (Oxford Univ Press, Oxford).
- Fairclough SH (2009) Fundamentals of physiological computing. *Interact Comput* 21(1):133–145.
- Zander TO, Kothe C (2011) Towards passive brain-computer interfaces: Applying brain-computer interface technology to human-machine systems in general. *J Neural Eng* 8(2):025005.
- Zander TO, Jatzev S (2012) Context-aware brain-computer interfaces: Exploring the information space of user, technical system and environment. *J Neural Eng* 9(1):016003.
- Schultze-Kraft M, et al. (2016) The point of no return in vetoing self-initiated movements. *Proc Natl Acad Sci USA* 113(4):1080–1085.
- Zander TO, Gärtnert M, Kothe C, Vilimek R (2010) Combining eye gaze input with brain-computer interface for touchless human-computer interaction. *Int J Hum Comput Interact* 27(1):38–51.
- Blankertz B, Lemm S, Treder M, Haufe S, Müller K-R (2011) Single-trial analysis and classification of ERP components—a tutorial. *Neuroimage* 56(2):814–825.
- Gerjets P, Walter C, Rosenstiel W, Bogdan M, Zander TO (2014) Cognitive state monitoring and the design of adaptive instruction in digital environments: Lessons learned from cognitive workload assessment using a passive brain-computer interface approach. *Front Neurosci* 8:385.
- Chanel G, Kierkels JJ, Soleymani M, Pun T (2009) Short-term emotion assessment in a recall paradigm. *Int J Hum Comput Stud* 67(8):607–627.
- Clark A (2013) Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behav Brain Sci* 36(3):181–204.
- Friston K (2010) The free-energy principle: A unified brain theory? *Nat Rev Neurosci* 11(2):127–138.
- Brown H, Friston K, Bestmann S (2011) Active inference, attention, and motor preparation. *Front Psychol* 2(218):218.
- Holroyd CB, Coles MGH (2002) The neural basis of human error processing: Reinforcement learning, dopamine, and the error-related negativity. *Psychol Rev* 109(4):679–709.
- Sutton RS, Barto AG (1998) Reinforcement learning: An introduction. *IEEE Trans Neural Netw* 9(5):1054.
- Bell AJ, Sejnowski TJ (1995) An information-maximization approach to blind separation and blind deconvolution. *Neural Comput* 7(6):1129–1159.
- Makeig S, Jung T-P, Bell AJ, Sejnowski TJ (1996) Independent component analysis of electroencephalographic data. *Neural Information Processing Systems*, eds Touretzky DS, Mozer MC, Hasselmo ME (MIT Press, Cambridge, MA), Vol 8, pp 145–151.
- Hawkins J, Blakeslee S (2007) *On Intelligence* (Macmillan, London).
- Settles B (2012) Active learning. *Syn Lect AI Mach Learn* 6(1):1–144.
- Zander TO, Jatzev S (2009) Detecting affective covert user states with passive brain-computer interfaces. *3rd International Conference on Affective Computing and Intelligent Interaction and Workshops, 2009* (IEEE, Amsterdam), pp 1–9.
- Protzak J, Ihme J, Zander TO (2013) A passive brain-computer interface for supporting gaze-based human-machine interaction. *Universal Access in Human-Computer Interaction: Design Methods, Tools, and Interaction Techniques for eInclusion*, eds Stephanidis C, Antonia M (Springer, Berlin), Vol 8009, pp 662–671.
- Zander TO (2011). Utilizing brain-computer interfaces for human-machine systems. Doctoral dissertation (Universitätsbibliothek der Technischen Universität Berlin, Berlin).
- Kothe CA, Makeig S (2013) BCILAB: A platform for brain-computer interface development. *J Neural Eng* 10(5):056014.
- Haufe S, et al. (2014) On the interpretation of weight vectors of linear models in multivariate neuroimaging. *Neuroimage* 87(0):96–110.
- Palmer J (2012) AMICA—Adaptive Mixture ICA. Available at https://scn.ucsd.edu/~jason/amica_web.html. Accessed April 25, 2014.
- Oostenveld R, Delorme A (2003) DIPFIT 2.x. Available at https://scn.ucsd.edu/wiki/EEGLAB_extensions_and_plugin-ins. Accessed April 25, 2014.
- Miyakoshi M (2003) dipoleDensity. Available at https://scn.ucsd.edu/wiki/EEGLAB_extensions_and_plugin-ins. Accessed April 25, 2014.

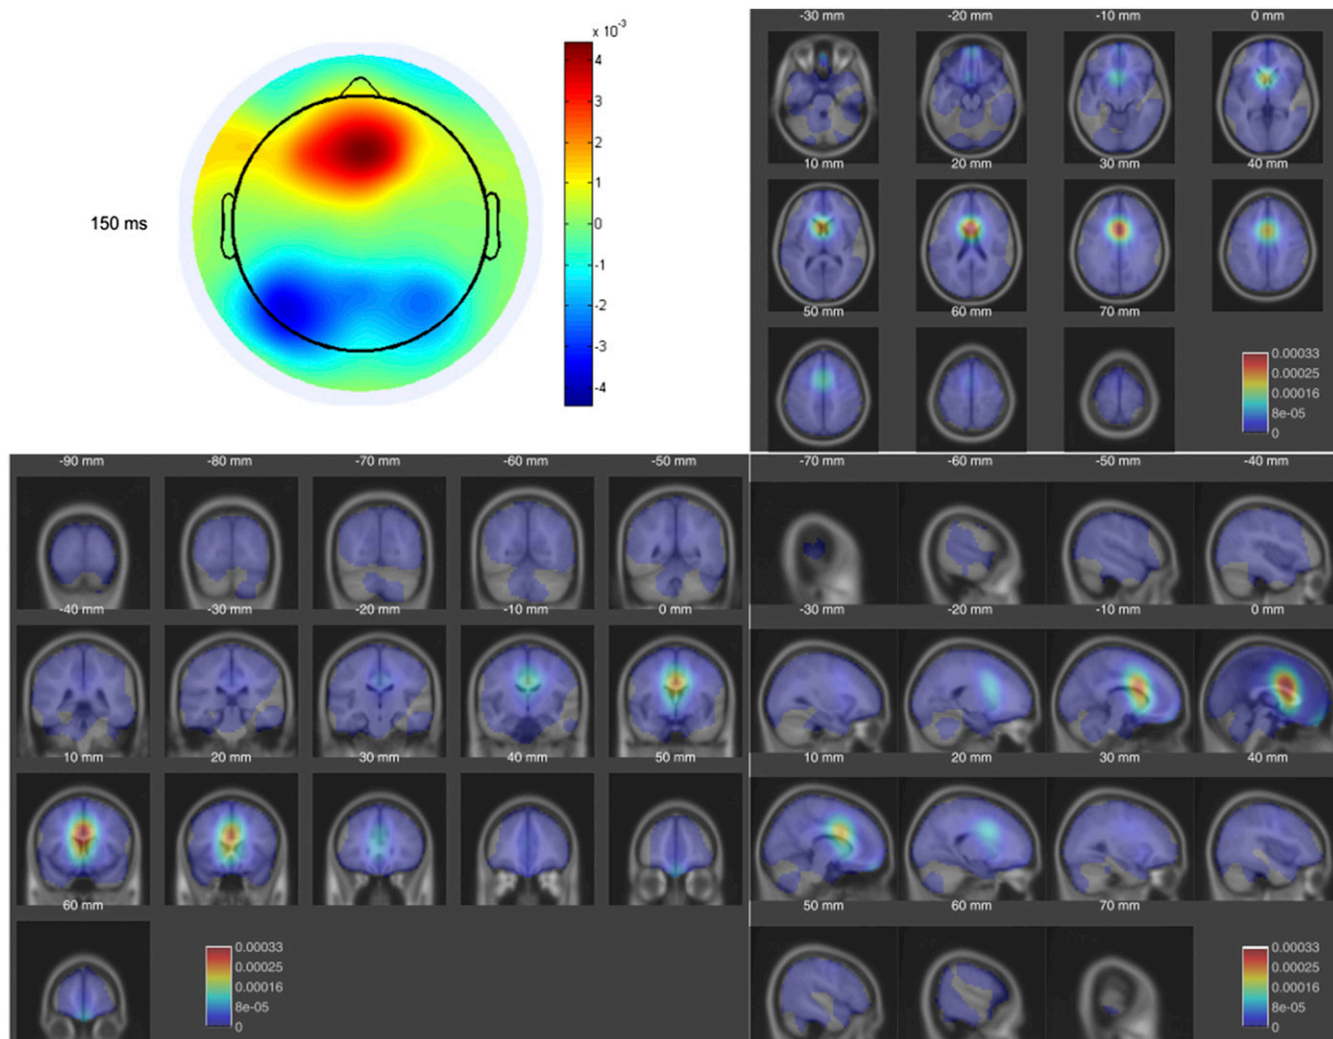
Supporting Information

Zander et al. 10.1073/pnas.1605155114



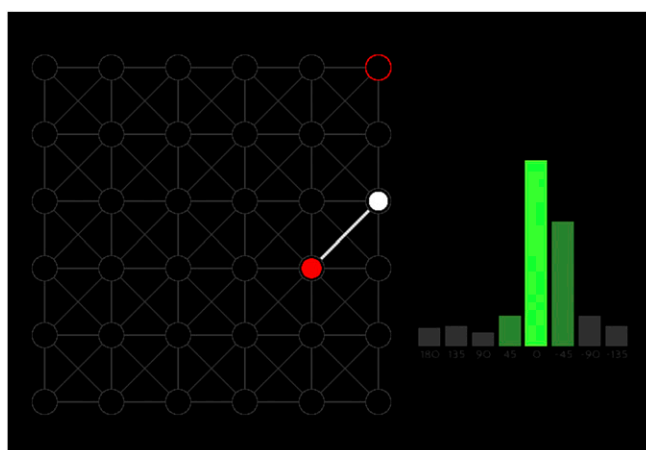
Movie S1. Event-related potential at electrode Fz, and class-correlated scalp activity derived from the LDA filter weights of sequential 50-ms time windows, spanning the 400 ms used for classification.

[Movie S1](#)



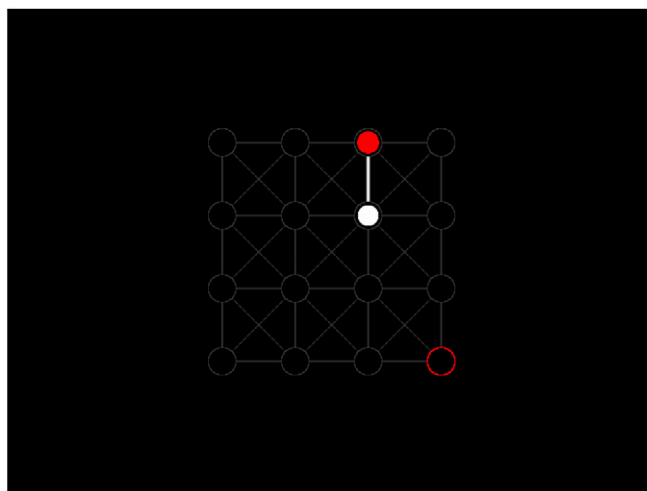
Movie S2. Dipole densities weighted by relevance for classification in sequential 50-ms time windows, spanning the 400 ms used for classification.

[Movie S2](#)



Movie S3. Online experimental stimuli accompanied by classification output, illustrating the online process and outcome.

[Movie S3](#)



Movie S4. Offline experimental stimuli.

[Movie S4](#)

Other Supporting Information Files

[SI Appendix \(PDF\)](#)

Neuroadaptive technology enables implicit cursor control based on medial prefrontal cortex activity

Thorsten O. Zander, Laurens R. Krol, Niels P. Birbaumer, Klaus Gramann

Table of Contents

Materials and Methods	2
1. Experimental Set-Up	2
2. Experimental Paradigm	2
3. Participants	4
4. Experimental Procedure	4
7. Performance Measures and Statistics	8
8. Linear Dependency of Peak Amplitude.....	9
Figures	10
Fig. S1. Experimental Paradigm and Procedure.....	10
Fig. S2. Angular Deviance	11
Fig. S3. Class Definition of the Classifier's Training Set	12
Fig. S4. Directional Probabilities	13
Fig. S5. LDA-projected ERPs	14
Fig. S6. Influence of Eye Movements on the ERP at Fz	17
Fig. S7. Mean Peak Amplitudes at Fz	18
Fig. S8. Scalp Activity Series.....	19
Fig. S9. Example Online Cursor Behavior.....	20
Fig. S10. The neuroadaptive control loop	21
Tables	22
Table S1. Online Classification Accuracy.....	22
Table S2. Post Hoc: Final Direction Probabilities.....	23
Table S3. Post Hoc: Peak Amplitude Differences at Fz.....	24
Table S4. Post Hoc: Classifier Output.....	25
Table S5. Pairwise Comparisons: Projected ERP Peak Amplitude Differences	26
References	26

Materials and Methods

1. Experimental Set-Up

Two computers were used for this experiment: One to display the stimuli and another to record EEG. Stimuli were shown on an Iiyama ProLite B2776HDS 27'' display using a resolution of 1920×1080 pixels, located approximately one meter away from the participants. The keyboard used was a standard US layout computer keyboard.

EEG was recorded continuously using 64 active Ag/AgCl electrodes (actiCAP; Brain Products GmbH, Munich) mounted according to the extended 10-20 system on an elastic cap (EASYCAP GmbH, Munich). The signal was sampled at 500 Hz and amplified using BrainAmp amplifiers (Brain Products GmbH, Munich). All electrodes were referenced to FCz.

EEG data was recorded and combined with the paradigm's markers using the Lab Streaming Layer framework (Swartz Center for Computational Neuroscience [SCCN], University of California, San Diego [UCSD]).

2. Experimental Paradigm

The cursor's movements were restricted by the nodes of a visible grid. The cursor started out on one of these nodes and, for each movement, could only move to one of the adjacent nodes. Depending on the cursor's position in the grid, then, there were up to eight possibilities for one movement. The grid consisted of grey open circles connected by grey lines on a black background. The cursor in this paradigm was a red, filled circle. The cursor's size was 4.5% of the display's total height, the grid nodes' size 6%, and the horizontal and vertical grid lines' length 15%. Line thickness between the grid nodes was set to 2 pixels. 'Red' was pure RGB(255,0,0), 'grey' RGB(51,51,51) and 'black' RGB(0,0,0). The dimensions of the grid are discussed below.

An animation allowed the participants to be able to anticipate the moment of each movement. Over the course of one second, a white 'ghost cursor' would grow inside of the actual cursor. As soon as this ghost reached the same size as the actual cursor, it would instantaneously be repositioned to the chosen adjacent node, while also highlighting the grid line connecting the two nodes in white. 'White' here was pure RGB(255,255,255), and the highlighted grid line's thickness was increased to 3 pixels. The movement remained visible for one second, with the red original cursor still on the initial node, the white ghost cursor on the new node, and a white line connecting them. Following that, all whites disappeared and the (red) cursor would instantaneously move to and remain at its new position, on the new node, for another second, before the animation would start over for the next movement. In all, a single trial was three seconds in length.

Grids of three different dimensions were used in this study: One by three nodes 1×3), four by four nodes (4×4), and six by six nodes (6×6).

On each grid, a single red node in one of the corners indicated the current target. Once the cursor had landed on this target node, a new grid was started with the same dimensions, but a different layout. For the 1×3 grids, each new grid was rotated either 45, 90, or 135 degrees relative to its predecessor, and a new target corner was chosen at random ('corner' here is one of the two ending nodes of the grid row). The larger grids did not rotate, but a new target was selected for each new grid such that no two subsequent grids had a target in the same corner. Each newly started grid was displayed for one second before the first movement was initiated, for the participants to orient themselves.

In all grids, the cursor's starting position was one node away from the corner opposite the target's, in a straight line to the target.

Supplementary Figure S1 illustrates the stimuli; Supplementary Movie 3 contains animated stimuli as seen by the participant.

For each individual cursor movement, a movement direction was chosen randomly from the set of possible movements. By default, all possible directions had an equal probability of being selected, and during the offline (calibration) sessions, these initial probabilities remained unchanged---the cursor was unbiased and moved randomly. However, during online application of the pBCI, the directional probabilities were altered based on the classification of each movement as either 'correct' or 'incorrect', biasing the cursor to repeat movements classified as 'correct'.

Specifically, all movement directions were represented a specific number of times in the set from which each new movement was randomly selected. After classification, the respective directions' numbers were increased or decreased such that the resulting probability of each subsequent direction is given by

$$p = \frac{mx}{n + (m - 1)x}$$

where x is the respective direction's initial number of shares in the set, n is the full set's initial size, and m is a multiplier. For a movement classified as 'correct', m was 2.0 for that movement's direction and 1.5 for the two adjacent directions; for movements classified as 'incorrect', m was 0.5 and 0.75, respectively. All directions started with a share of 100 elements each in the set. All shares below 1 were rounded up at the time of selection.

The adaptation of adjacent directions represents a certain degree of leniency which the participants were assumed to show in their judgements: If the target is a few nodes north of the cursor's current position, a first move to the north-west is imperfect but may still be acceptable. Therefore, the feedback received from a movement to the north-west should not only alter this specific direction's probability, but should also influence subsequent movements to both the north and the west.

All visible on-screen events were marked in the EEG stream and additional information about each cursor movement was logged, including distance to the target and the movement's direction

relative to the target direction (i.e., a straight line unconstrained by the grid represented a 0° movement).

The paradigm was written in Python using the Simulation and Neuroscience Application Platform 1.02-beta (SCCN, UCSD; <https://github.com/sccn/SNAP>).

3. Participants

A total of nineteen participants participated in this study, with an average age of 25.4 years \pm 3.4. Seven were female, and a total of two participants were left-handed. Participants were asked, alternately, to use either their left or their right hand to indicate their judgements. This resulted in ten participants not using their dominant hand. All had normal or corrected to normal vision. Instructions were given in writing, in German. While all knew German, thirteen did not have German as their mother tongue; these were verbally given additional, standardized instructions.

Nineteen participants performed calibration trials, while sixteen of them additionally performed online trials.

4. Experimental Procedure

All participants were informed of the nature of the experiment and the recording and anonymization procedures before signing a consent form. The ethics committee of the Department of Psychology and Ergonomics at the Technische Universität Berlin approved the experiment and the procedures.

After preparation and setting up the EEG cap, which took roughly one hour per participant, participants were seated in a padded chair in a dimly lit room, about one meter away from the stimulus display. In writing, they were instructed to judge each individual cursor movement on the display as either ‘acceptable’ or ‘not acceptable’, with respect to reaching the target, and to indicate their judgement by pressing either ‘v’ or ‘b’, respectively, on a computer keyboard using one and the same finger of one hand. The participants performed this task during all blocks.

Participants were first given four blocks of 50 trials on 1×3 grids. Here, the cursor performed only one movement per grid: Regardless of whether or not that movement reached the target, a new grid was subsequently started. This was because after a movement away from the target, only one movement possibility would remain, i.e. that trial would have no informational value. Breaks between these blocks were self-paced, and participants were given time to practice before the first block was started.

Following these first blocks, participants received additional instructions for the larger grids, emphasizing their task to judge every movement individually and independently of the cursor’s movement history. Participants were again given some time to practice, on a 4×4 grid, and then performed five blocks of 120 trials on grids of that size. Here, if the target had not been reached after 55 trials in one grid, a new grid was started. 55 is twice the median number of random movements required to reach a target on a 4×4 grid.

Taking symmetry and rotation into account, a total of 43 unique cursor movements are possible on the 4×4 grid, i.e., 43 unique pairs of distance and angular deviance. In between the five

calibration blocks, in four sessions, participants were shown, in random order, all 43 unique cursor movements. For each of these, participants were asked to rate it as either ‘acceptable’ or ‘not acceptable’, and on a scale of 1 to 5, with 5 being ‘the best possible movement in that situation’, and 1 being ‘the worst’. Here, unlike during the blocks, there was no time limit imposed on their answers.

EEG recorded during these latter five blocks served to calibrate the classifier, as discussed in Supplementary Method 5. This classifier was then applied to one more block of 120 trials on 4×4 grids, and one last block of 120 trials on 6×6 grids. No maximum number of trials other than the block’s length was set for these online blocks.

5. Feature Extraction and Classification

A BCI based on supervised machine learning needs to be calibrated before it can be applied. This calibration is typically performed on sets of recordings, usually EEG epochs, which are known to contain the signals that need to be detected later. On the basis of these epochs, a classifier is calibrated to optimally distinguish between the different classes of source signals.

The classes the classifier was calibrated on were formed on the basis of the movements’ angular deviance from a straight line towards the target, as illustrated in Supplementary Figure S2. ‘Correct’ movements were those with an absolute angular deviance of 0°; ‘incorrect’ were those with an absolute angular deviance of 135° or more, as in Supplementary Figure S3. These class definitions were determined based on the first three participants’ data.

Note that the participants’ judgements, indicated using button-presses, were ignored: Only a movement’s angle with respect to the target determined its class. Movements between 0 and 135° were not included for calibration.

The open-source toolbox BCILAB (23) version 1.01 was used to define and implement the pBCI. Features were extracted through the windowed means approach. This approach extracted features from the time course of each of the 64 channels by subsampling the ERPs of each epoch. This was done by dividing the time course into a sequence of 8 consecutive 50 ms windows (8 time windows starting at 50 ms after cursor movement) and replacing each window by its average. The resulting features were thus one value for each channel in each time window. Linear discriminant analysis (LDA) was then applied to these features generated from all available calibration trials to distinguish between the features belonging to the two classes. The outcome of the LDA is a linear weighting of all features. In other words, in each time window all channels receive a weight according to their relevance for classification in that specific time window. The linear combination of all features of one single trial with these weights then gives a number between -1 and 1, indicating whether this trial is classified as belonging to class 1 or to class 2.

For the feature extraction, the data was first resampled at 100 Hz, and band-pass filtered from 0.1 to 15 Hz. Linear discriminant analysis was regularized by shrinkage. A [5,5]-times nested cross-validation with margins of 5, ensuring the independence and identity of the feature distributions, was used to select the shrinkage regularization parameter, and to generate estimates of the model’s online reliability.

6. Extended Feature Analysis

The method used by the pBCI to extract discriminative information (described in Supplementary Method 5, above) can be analyzed to reveal further insights into the relevant underlying processes. The classification model used here is a multivariate approach, a linear discriminant analysis (LDA), optimized for the discriminability of the extracted features between classes. Each feature represents data at a single sensor for one of the chosen time windows. Hence, the methodology recently introduced by Haufe *et al.* (1) could be applied to interpret the classification model.

For one of the eight 50 ms time windows, Figure 1 in the main manuscript illustrates three aspects of the information used by the pBCI: its scalp topography, its estimated source within the brain, and the activity of that source as filtered from the EEG, i.e. the selected signal of interest.

Signal of interest. Panel (C) of Figure 1 in the main manuscript shows the EEG activity projected by the LDA filter for the third time window, i.e., the combined activity of all 64 channels as filtered by the classifier. Supplementary Figure S5 illustrates this for all eight time windows.

In Supplementary Figure S5, the ERP is shown for eight groups of cursor movements, time-locked to that movement, as well as the difference wave between the two classes ('correct' and 'incorrect'). The difference wave represents the actual signal (and the optimization target) used by the classifier, in the corresponding time window. The time window of the LDA filter used is indicated in grey.

Each figure indicates the response to different cursor movements of those processes whose activity was filtered from the EEG in that time window, over the course of a full second.

Identifying scalp projections. In panel (A) of Figure 1 in the main manuscript, the scalp map shows, on average for all participants, the (interpolated) activation pattern which illustrates how the signal of interest is expressed in the 64 channels (I), in the third time window. Supplementary Movie 1 shows this activation pattern over the full length of the used 400 ms.

For each participant, LDA patterns $A = (a_j)_j$ were generated from the LDA filters $M = (m_j)_j$ originally used for online classification by conjugation with the features' covariance matrix C : $A = CM C^{-1}$. Spatial interpretation of these patterns for each time window reflects a mixture of scalp activations related to discriminative source activity $\hat{A} = (\hat{a}_j)_j$ and class-invariant noise representation N , with $A = \hat{A} + N$. The latter was filtered out by weighting each pattern entry a_j with the correlation of its associated feature activity vector over trials F_j to the binary vector of true class labels L :

$$\hat{a}_j = \text{corr}(F_j, L) \cdot a_j$$

The resulting correlated pattern $\hat{A} = (\hat{a}_j)_j$ can be visualized by topographic plots for each time window, as in Figure 1 of the main manuscript, Supplementary Figure S5, and Supplementary Movie 1.

Identifying relevant sources. To identify the sources underlying these topographic representations, the backward model, i.e. the LDA filter, was combined with an independent component analysis (ICA). The ICA unmixing matrix $W = (I_1, I_2, \dots, I_n)$ was determined on previously manually cleaned data for each participant by using the Adaptive Mixture ICA (AMICA) Toolbox (2), such that $s = Wx$, where s represents the source activation related to a given scalp activation x . For each time window, the relevance for classification R_i of each independent component I_i can then be determined by distributing the LDA filter weights to the independent components via W , weighted by two factors. The first factor is compensating for the amplitude alignment of the LDA filter weights to the feature amplitudes. It is determined by calculating the variance over trials of the feature \hat{F}_i extracted from the time series of the independent component: $V_i = \text{var}(\hat{F}_i)$. A second weight is determined for filtering out noise representations by weighting the independent components with the correlation of their feature activity to the true class labels (as described above for electrode activity).

$$R_i = V_i * \text{corr}(\hat{F}_i, L) * WM$$

Localizing relevant sources. To localize the identified sources, equivalent dipole models that describe the most likely position of the source in a standard head model were identified for selected components by using the EEGLAB plug-in DIPFIT 2.x (3). Components were selected by a threshold criterion for residual variance of the dipole model ($RV < 0.15$) and visual inspection of activation spectra, time courses, and scalp topographies. Only components reflecting cortical, ocular, or muscular activity were included in the analysis.

For each time window, each of the 371 resulting dipoles was weighted by the relevance R_i of its associated independent component, described above. The areas of high relevance were then described by a weighted dipole density plot using the EEGLAB plug-in dipoleDensity (4).

Supplementary Movie 2 shows the result of this analysis over the full length of the used 400 ms, by plotting the dipole density per cubic millimeter weighted by the relevance R_i of each included dipole with a smoothing kernel of 12 mm.

The above-mentioned process of selection did not markedly influence this analysis. Compared to all 1191 dipoles and averaged over the eight time windows, the 820 rejected dipoles (68.85%) carried 7.5% of the weights distributed by the classifier. Relative to all 1191 dipoles, a total of 87 dipoles received a relevance weight larger than 1 standard deviation in at least one of the time

windows. These 7.3% of dipoles carried 77.82% of all total weight distributed by the classifier. Four of these highly weighted dipoles (4.6%) were rejected in the process explained above and not included in the analysis. These four represent 1.67% of the weight included in the analysis. Three belonged to the same subject.

Behavioral analysis. Participants indicated their evaluation of each movement as ‘acceptable’ or ‘not acceptable’ through button presses. A paired-samples t-test was conducted to compare the subjects’ mean reaction times. There was a significant difference between ‘acceptable’ button presses (mean 369.7 ms±66.9) and ‘not acceptable’ button presses (mean 310.7 ms ±60.7); $t(18) = 8.347$, $p < 0.001$. This calculation included only reaction times between 100 and 1000 ms.

No systematic effect reflecting this delay can be found in the ERPs used for classification however (see Figure 1, and Supplementary Figure S5). This, in combination with an absence of relevance weights in the motor cortex (Supplementary Figure S5) and the ERPs' independence of eye movements (Supplementary Figure S6), leads us to conclude that the behavioral delay has had no effect on the passive BCI.

7. Performance Measures and Statistics

The performance of the cursor was operationalized as the number of movements required to reach the target. Trials at the end of a block that did not contribute to reaching a target or hitting the maximum number of trials for that attempt, were discarded. This measure does include those attempts that hit the maximum number of trials, which happened a total of seven times to seven participants on the 4×4 grid, and to two participants on the 6×6 grid. This resulted in 88 pBCI-supported data points for the 4×4 grid, and 47 for the 6×6 grid.

As a comparison, a sample of around 7000 data points was taken using the same measure from non-supported (random) blocks. This data is non-parametric and varies greatly. Therefore, a resampling approach was taken where the available sample of pBCI-supported measures was repeatedly compared to a new random sample of the same size of non-supported performance measures. This comparison was done through a Wilcoxon rank-sum test. Out of 50 000 such comparisons, 98.02% of tests were significant at the $\alpha=0.025$ level for the 4×4 grid performance measures; for the 6×6 grid, 100.00% of tests were significant at this level. On a 4×4 grid, non-supported cursor movement required on average (median) 27 trials to reach a target, whereas with BCI support, this decreased to an average of 13. On a 6×6 grid, the median values were 90 and 23, respectively.

The pBCI-supported performance was furthermore compared to ‘perfectly supported’ performance, where the cursor was reinforced as during the online sessions, except automatically, with perfect accuracy. Here, all movements with an angular deviance less than 45° were judged to be ‘correct’ and reinforced accordingly, and all others were ‘incorrect’, and their probabilities decreased as described in Supplementary Method 2. The same procedure as above yielded significant differences to the pBCI-supported measures for both the 4×4 grid (99.98% of tests at $\alpha=0.025$) and the 6×6 grid (100%). The median perfectly supported performance measure on the 4×4 grid is 10 movements, and 14 movements on the 6×6 grid.

Additionally, the mean directional probabilities upon reaching the target in both online grids (combined) were calculated. These are illustrated in Supplementary Figure S4. Supplementary Table S2 lists the Bonferroni-adjusted results of pairwise post hoc tests from a one-way ANOVA

($F(7,105)=57.520, p<0.001$) on this data. On average, the classifier has been able to reinforce the ‘correct’ directions significantly more strongly than ‘incorrect’ directions.

8. Linear Dependency of Peak Amplitude

Figures 1 and 2 of the main manuscript show a linear scaling of peak amplitudes with respect to absolute angular deviance away from the target. The eight groups of angular deviations were selected such that for each single participant, at least 50 trials were present in each group, i.e., for an optimal signal to noise ratio while maintaining maximal angular resolution.

For the derivation of the single electrode ERP, BrainVision Analyzer was used (Version 2.0.2.5859, Brain Products GmbH, Munich, 2012). The raw EEG data was first band-pass filtered from 0.5 to 45 Hz, and decomposed into statistically maximally independent source signals through ICA. Individual components that resembled eye movements and eye blinks were manually selected for removal based on their time course and topography (as e.g. described in 26). Figure 1 shows the grand average ERP over all nineteen participants at Fz with these components removed. For a comparison, Supplementary Figure S6 shows the grand average ERP at Fp1, Fp2, and Fz for only the removed components. No systematic response of the sort seen in Figures 1 or 2 is visible, pointing to cortical causes of the differences. At Fz, where the cortical components peak around +3.9 μV for e.g. the 0° condition, the eye components show an amplitude of -0.37 μV .

Statistical analysis focused on the systematic peak differences seen at Fz around 180 ms. A one-way ANOVA indicated a significant influence of angular deviance on peak amplitude ($F(7,126) = 47.243, p < 0.001$). Post-hoc comparisons corrected for false discovery rate are listed in Supplementary Table S3 and illustrated in Supplementary Figure S7. The peak amplitudes differ significantly ($p < 0.001$) between the classes used by the classifier. In between, the peak amplitudes scale linearly with angular deviance, as fitted by a linear regression model using each group’s mean angular deviance as predictor ($b = -0.0035, F = 45.28, p < 0.001; R^2 = 0.33$). Classifier output followed a similar trend, as in Figure 1 of the manuscript.

Figures




Fig. S1. Experimental Paradigm and Procedure

The experimental paradigm and procedure. *Left*: Three different grid sizes were used in the experiment. Data from five blocks in 4×4 grids was used to calibrate the classifier, which was then applied to another 4×4 block, and one block on a 6×6 grid. Data from the 1×3 grids has not been used in this paper. Each cursor movement consisted of the cursor moving from one node to one of the directly adjacent ones. The cursor could move horizontally, vertically, and diagonally over the grid. Since the target was known and indicated, for each movement, it was possible to determine a measure of correctness for each movement by means of calculating the angular deviance, as in Supplementary Figure S2. *Right*: Detail of a 4×4 grid showing the cursor's movement animation.




Fig. S2. Angular Deviance

Illustration of the angular deviance of two possible cursor movements (north, and north-west) from the optimal path straight towards the target. Angular deviance was measured as the absolute deviation, in degrees, of the movement direction from a straight line to the target.




Fig. S3. Class Definition of the Classifier's Training Set

The classifier was calibrated on a subset of trials from a 600-trial calibration session. This set was selected by angular deviance: Movements with an absolute angular deviance to the target of 0° were included as 'correct', while movements with a deviance of 135° or more were included as 'incorrect'. This figure illustrates which of all possible movements on the 4×4 grid were included, and in which category: Green represents 'correct', red 'incorrect'. From the total of 600 trials per participant on the 4×4 grids, this selection left 62.7 ± 7.8 'correct' and 124.4 ± 8.9 'not acceptable' trials per participant for the classifier to be calibrated on.




Fig. S4. Directional Probabilities

Mean directional probabilities, relative to the target from the cursor's starting position, upon hitting the target, in the BCI-supported condition, averaged over participants. Error bars represent the standard deviation. See Supplementary Table S2 for significance tests.




Fig. S5. LDA-projected ERPs

(Figure spans three pages; this is page 1 of 3.)

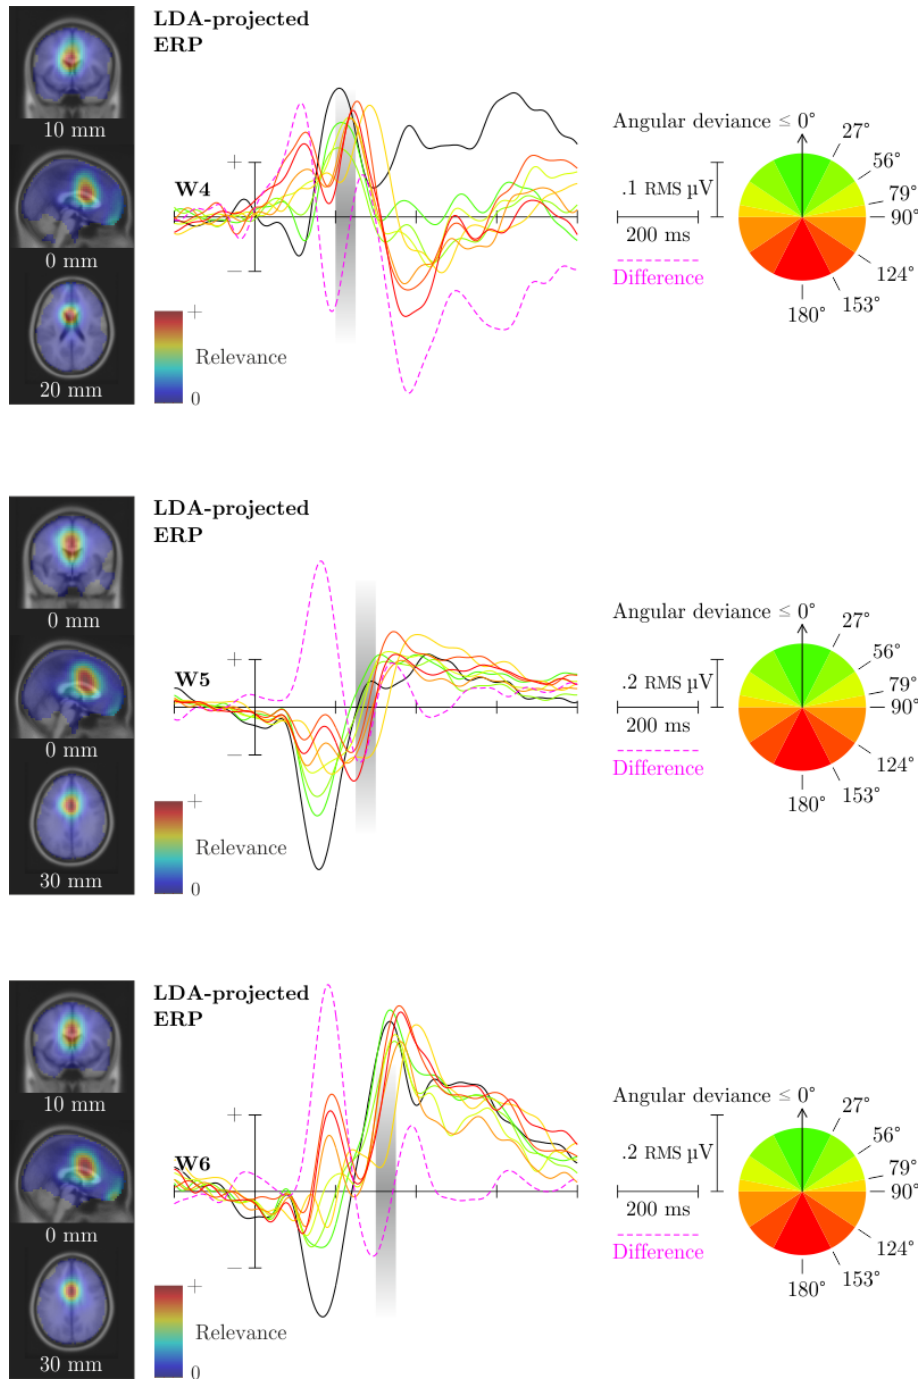


Fig. S5. LDA-projected ERPs

(Figure spans three pages; this is page 2 of 3.)

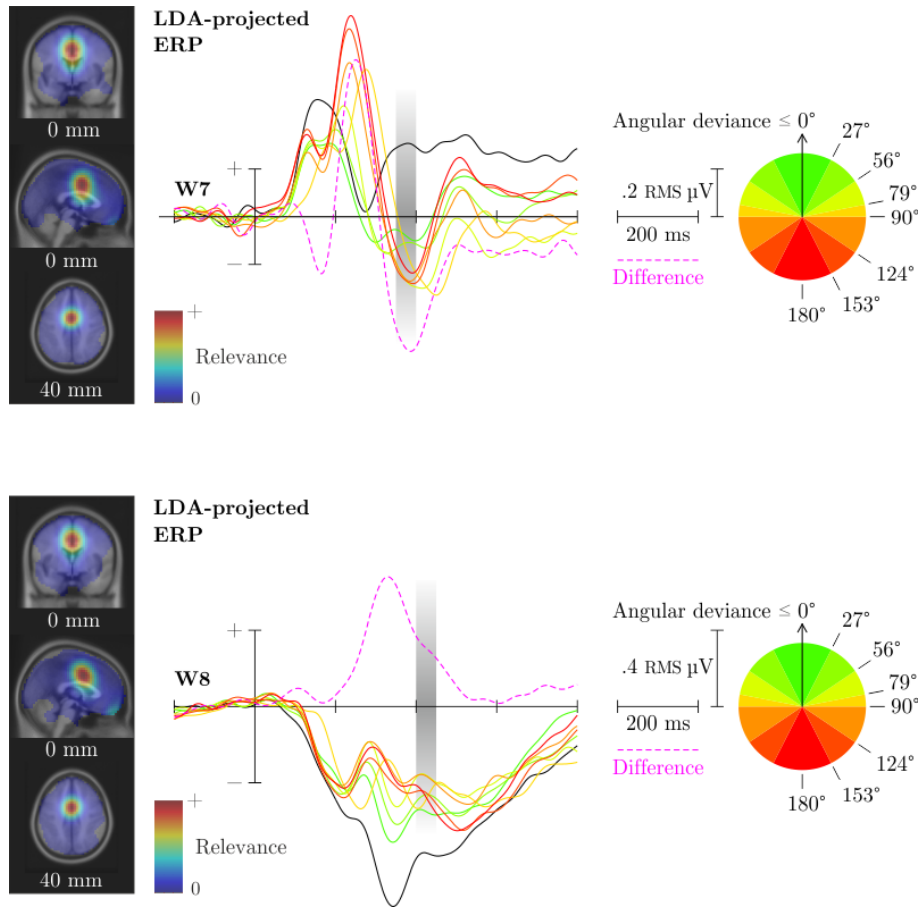


Fig. S5. LDA-projected ERPs

(Figure spans three pages; this is page 3 of 3.)

The classifier was calibrated on features spanning eight time windows of 50 ms each. For each time window, the above figures show on the *left*: Source localisation via weighted dipole density of the process(es) focused on in that time window; *right*: Full-length ERPs, combined as per that time window's LDA filter, of eight groups of cursor movement, as well as the projected difference wave between 'correct' and 'incorrect' classes (i.e. the filter's optimisation target). Each figure's actual time window, which actually contributed to classification, is highlighted in grey.




Fig. S6. Influence of Eye Movements on the ERP at Fz

The grand average ERP ($n = 19$) at Fp1, Fp2, and Fz for only those components that contained strong eye movements and were removed for neurophysiological analysis. No systematic response is visible.




Fig. S7. Mean Peak Amplitudes at Fz

Peak amplitudes around 180 ms for the eight groups of cursor movements. The peak amplitudes differ significantly ($p < 0.001$ ***) between the classes used by the classifier. In between, the peak amplitudes scale linearly with angular deviance, as fitted by a linear regression model using each group's mean angular deviance as predictor ($b = -0.0035$, $F = 45.28$, $p < 0.001$; $R^2 = 0.33$). Statistically significant differences between adjacent groups are indicated as well ($p < 0.05$ *). See Supplementary Table S3 for exact figures of all comparisons. Classifier output followed a similar trend, as in Figure 4 of the manuscript.




Fig. S8. Scalp Activity Series

Class-specific grand average scalp activities ($n = 19$), time-locked to cursor movement, at the middle of each time of the eight time windows used for classification. *Left*: Scalp activities of class 1, containing only cursor movements that went directly towards the target. *Right*: Scalp activities of class 2, containing cursor movements with an angular deviance of 135° or more. EEG data was first band-pass filtered from 0.1 to 5 Hz as per the classification approach and re-referenced to the common average. See Figure 1 in the main manuscript and Supplementary Movie S1 for the class-correlated scalp activity that the classifier focused on.

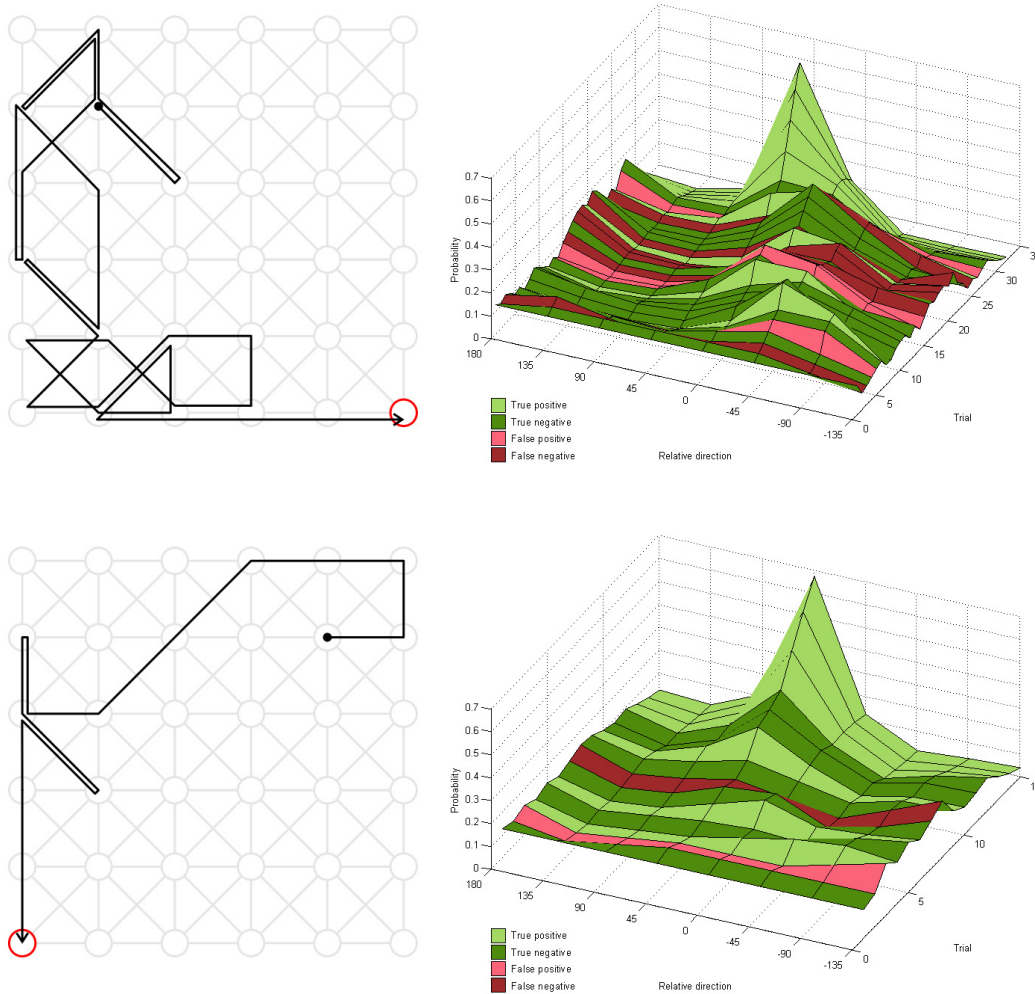


Fig. S9. Example Online Cursor Behavior

Visualization of two online target selections. *Left*: Trace of cursor movements over the grid. Each cursor movement only progressed one node over the grid; extended straight lines thus reflect multiple movements. *Right*: Progression of normalized directional probabilities relative to 0° , with 0° being the direction towards the target at the start of the grid. All directions started out with equal probabilities, indicated at trial 0. Each subsequent trial shows the directional probabilities after that trial's cursor movement was classified and the cursor was reinforced accordingly. The colors reflect the confusion matrix with the participant's button presses being used as ground truth. True positive (light green) indicates the cursor movement was correctly classified as 'correct'. True negative (dark green) indicates a correctly classified 'incorrect' movement; false positive (light red) an incorrectly classified 'correct' movement; false negative (dark red) an incorrectly classified 'incorrect' movement. The top row shows a relatively slow online run of 31 trials on a six-by-six grid. A number of misclassifications, especially near the middle, delay the probabilistic model's convergence towards the desired target direction. However, the erroneous trials do not lead to an adverse bias while the correct classifications do systematically point towards the target, and the correct direction is found. The lower row shows a shorter run of 14 trials with only two misclassifications.

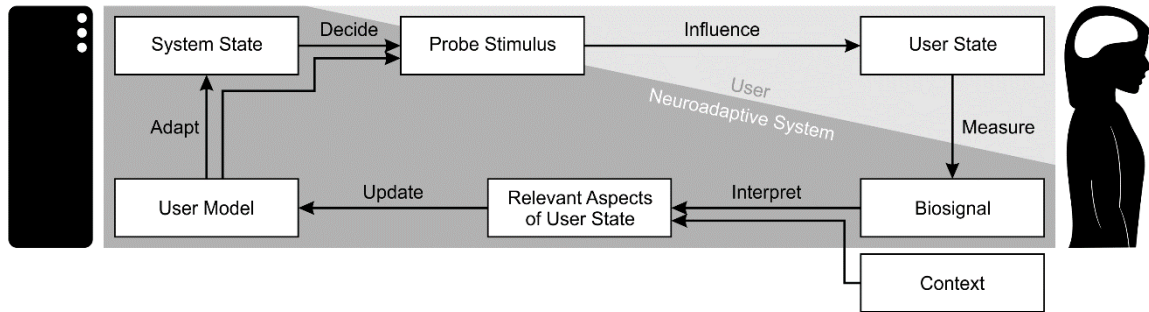


Fig. S10. The neuroadaptive control loop

A stimulus or parameter change provokes an automatic response from the user, whose biosignals are monitored. The user's response can be classified from the gathered data, allowing the system to interpret the response in light of the previously gathered user information and the current context. With this information, the system updates its user model. Based on the user model and the current system status, the system may decide on a new probe stimulus.

Tables

Participant	4 × 4 Grids			6 × 6 Grids		
	TP	TN	CR	TP	TN	CR
07	.91	.86	.89	.	.	.
08	.91	.60	.77	.	.	.
09	.81	.53	.71	.	.	.
10	.83	.79	.81	.76	.75	.76
11	1.00	.83	.93	.88	1.00	.93
12	.83	.63	.76	.82	.64	.73
13	.67	.47	.59	.87	.38	.67
14	.75	.63	.70	.84	.36	.69
15	.82	.69	.75	.71	.82	.75
16	.52	.50	.52	.39	.10	.29
17	.58	.62	.59	.50	.68	.58
18	.67	.62	.65	.44	.53	.48
19	.80	.67	.75	.92	.65	.80
20	.71	.80	.74	.70	.62	.67
21	.88	.56	.76	.74	.45	.63
22	.60	.67	.63	.40	.50	.43
Mean	.77	.65	.72	.69	.58	.65

Table S1. Online Classification Accuracy

Online classification rates for the 4×4 grid, and the 6×6 grid. TP = True Positives, TN = True Negatives, CR = Classification Rate (combined TP and TN).

Supplementary Table S1 lists the classification rates for the online blocks, i.e. what percentage of classifier output agreed with the given definitions of ‘correct’ and ‘incorrect’ movements. These definitions were the same as the ones used for calibration: Movements with an angular deviance to the target of 0° were defined to be ‘correct’, and those with a deviance of 135° or more were taken to be ‘incorrect’. The labels of all other movements were subject to individual interpretation, and have therefore not been included in this analysis.

The table lists these rates for all participants that completed online blocks. From the total of nineteen, this excludes the first three participants, who only performed offline calibration blocks. A further three participants only performed 4×4 grids online.

	Pairwise Comparison						
	180°	135°	90°	45°	0°	-45°	-90°
135°	1.000
90°	0.004	0.003
45°	0.000	0.000	0.000
0°	0.000	0.000	0.000	1.000	.	.	.
-45°	0.000	0.000	0.304	1.000	0.000	.	.
-90°	0.000	0.000	1.000	0.000	0.000	0.006	.
-135°	1.000	1.000	0.005	0.000	0.000	0.000	0.000

Table S2. Post Hoc: Final Direction Probabilities

Bonferroni-adjusted pairwise comparisons of the mean directional probabilities upon reaching the target.

Supplementary Figure S2 shows the mean directional probabilities upon reaching the target in both online grids, sorted by their angular deviance to the target, which was fixed relative to the cursor's starting position. Supplementary Table S2 lists the Bonferroni-adjusted results of pairwise post hoc tests from a one-way ANOVA ($F(7,105) = 57.520, p < 0.001$) on this data. On average, the classifier has been able to reinforce the 'correct' directions significantly more strongly than 'incorrect' directions.

	Pairwise Comparison						
	0°	≤ 27°	≤ 56°	≤ 79°	≤ 90°	≤ 124°	≤ 153°
≤ 27°	0.032
≤ 56°	0.001	0.075
≤ 79°	0.000	0.000	0.020
≤ 90°	0.000	0.000	0.004	0.309	.	.	.
≤ 124°	0.000	0.000	0.000	0.023	0.071	.	.
≤ 153°	0.000	0.000	0.000	0.021	0.066	0.507	.
≤ 180°	0.000	0.000	0.000	0.023	0.071	0.520	0.531

Table S3. Post Hoc: Peak Amplitude Differences at Fz

Pairwise comparisons, adjusted for false discovery rate, of the mean peak amplitudes at Fz for the eight angular categories.

A one-way ANOVA indicated a significant influence of angular deviance on peak amplitude ($F(7,126) = 47.24, p < 0.001$). In Supplementary Table S3 are listed the post-hoc comparisons—one-sided t-tests with pooled standard deviations, corrected for false discovery rate—between the eight individual groups of cursor movement.

	Pairwise Comparison						
	0°	≤ 27°	≤ 56°	≤ 79°	≤ 90°	≤ 124°	≤ 153°
≤ 27°	0.328
≤ 56°	0.000	0.002
≤ 79°	0.000	0.000	0.023
≤ 90°	0.000	0.000	0.000	0.115	.	.	.
≤ 124°	0.000	0.000	0.000	0.006	0.125	.	.
≤ 153°	0.000	0.000	0.000	0.308	0.311	0.829	.
≤ 180°	0.000	0.000	0.000	0.311	0.829	0.979	0.952

Table S4. Post Hoc: Classifier Output

Pairwise comparisons, adjusted for false discovery rate, of the mean peak amplitudes at Fz for the eight angular categories.

A one-way ANOVA indicated a significant influence of angular deviance on classifier output ($F(7,105) = 28.32, p < 0.001$). Below are listed the post-hoc comparisons—one-sided t-tests with pooled standard deviations, corrected for false discovery rate—between the eight individual groups of cursor movement.

	Pairwise Comparison						
	0°	≤ 27°	≤ 56°	≤ 79°	≤ 90°	≤ 124°	≤ 153°
≤ 27°	0.004
≤ 56°	0.000	0.007
≤ 79°	0.000	0.000	0.009
≤ 90°	0.000	0.000	0.000	0.093	.	.	.
≤ 124°	0.000	0.000	0.000	0.168	0.599	.	.
≤ 153°	0.000	0.000	0.000	0.014	0.161	0.134	.
≤ 180°	0.000	0.000	0.000	0.004	0.046	0.404	0.227

Table S5. Pairwise Comparisons: Projected ERP Peak Amplitude Differences

Pairwise comparisons of the mean amplitudes in the third time window (150-200 ms) of the LDA-projected ERP.

Supplementary Figure S5 shows LDA-projected ERPs. At Fz, significant amplitude differences were found at around 180 ms following cursor onset (see Figure 1 of the main manuscript and Supplementary Figure S7). This falls within the third time window used by the classification system (150 to 200 ms following cursor movement). This table shows the results of pairwise comparisons using one-tailed permutation tests of the mean amplitudes of the LDA-projected ERPs in that time window, between the eight individual groups of cursor movement.

References

1. Haufe, S., Meinecke, F., Görgen, K., Dähne, S., Haynes, J.-D., Blankertz, B., & Bießmann, F. (2014). On the interpretation of weight vectors of linear models in multivariate neuroimaging. *NeuroImage*, 87(0), 96–110.
2. Palmer, J. (2012). AMICA—Adaptive Mixture ICA. Retrieved 2014-04-25, from http://sccn.ucsd.edu/~jason/amica_web.html.
3. Oostenveld, R. & Delorme, A. (2003). DIPFIT 2.x. Retrieved 2014-04-25, from http://sccn.ucsd.edu/wiki/EEGLAB_Extensions_and_plugin-ins.
4. Miyakoshi, M. (2003). dipoleDensity. Retrieved 2014-04-25, from http://sccn.ucsd.edu/wiki/EEGLAB_Extensions_and_plugin-ins.