

Brain-Computer Interfacing for Intelligent Systems

Anton Nijholt, *University of Twente*
Desney Tan, *Microsoft Research*

Advances in cognitive neuroscience and brain-imaging technologies give us the unprecedented ability to interface directly with brain activity. These technologies let us monitor the physical processes in the brain that correspond with certain forms of thought. Driven by society's growing recognition of the needs of people with physical disabilities, researchers have begun using these technologies to build brain-computer interfaces (BCIs)—communication systems that don't depend on the brain's normal output pathways of peripheral nerves and muscles. In BCIs, users explicitly manipulate their brain activity instead of motor movements to produce signals that control computers or communication devices. This research has extremely high impact, especially for disabled individuals who can't otherwise physically communicate.

Although removing the need for motor movements in computer interfaces is challenging and rewarding, we believe the full potential of brain imaging as an input mechanism lies in the rich information it provides about the user's state. Having access to this state is important to researchers because it might let us derive more direct measures of traditionally elusive phenomena such as task engagement, cognitive workload, surprise, satisfaction, or frustration. These measures could open new avenues for evaluating systems and

interfaces. Additionally, knowing the user's state as well as the tasks they're performing might provide key information that would help us design context-sensitive systems that adapt themselves for optimal user support. This could prove useful to healthy users who might be situationally disabled—that is, they might lack full access to traditional, physically based communication modalities. It also opens a whole new domain of niche applications, carefully designed to exploit this novel modality's specific affordances, perhaps in conjunction with more traditional input devices. We believe that games might be an area of early adoption—first, because games have traditionally pushed us to consider completely new usage paradigms, and second, because gamers tend to be fairly tolerant of new technologies. Education could be another such domain.

The four short articles in this issue's Trends & Controversies provide a quick overview of the past, present, and future of BCIs. They are written primarily by European researchers working with noninvasive techniques, which represent a focused subset of the broader research and viewpoints in the field.

Gert Pfurtscheller and Clemens Brunner begin with a state-of-the-art survey. They discuss brain signals that can be measured with various devices, ways to control these signals, and how to train users to do this.

José del R. Millán describes real-time, robust control of brain-actuated robots and neuroprostheses. He focuses on how to optimally blend a human user's mental capabilities with a robot's intelligence to operate complex devices through a low-bit-rate BCI based on electroencephalography.

Brendan Allison and Bernhard Graimann present specific situations in which BCI research aimed at the physically disabled can apply to healthy users.

Finally, Florin Popescu, Benjamin Blankertz, and Klaus-R. Müller ground the opportunities in the hardware, computational, and social challenges we face as we work to create BCIs that work effectively in real-world environments.

Anton Nijholt is full professor of computer science at the University of Twente and chair of its Human Media Interaction subdepartment. Contact him at a.nijholt@ewi.utwente.nl.

Desney Tan is a researcher at Microsoft Research, where he manages the Computational User Experiences group. He also holds an affiliate faculty appointment in the Computer Science and Engineering Department at the University of Washington. Contact him at desney@microsoft.com.

The State-of-the-Art in BCIs

Gert Pfurtscheller and Clemens Brunner,
Graz University of Technology

A brain-computer interface (BCI) is a novel communication system that translates human thoughts or intentions into a control signal. In this way, a BCI provides a new, nonmuscular communication channel that system developers can use in a variety of applications, such as assisting people with severe motor disabilities; supporting biofeedback training in people suffering from epilepsy, stroke,

or attentional deficit hyperactivity disorders (ADHD); or controlling computer games.^{1,2}

Every mental activity—for example, decision making, intending to move, and mental arithmetic—is accompanied by excitation and inhibition of distributed neural structures or networks. With adequate sensors, we can record changes in electrical potentials, magnetic fields, and (with a delay of some seconds) metabolic supply when the activated neuron population exceeds some critical mass. Consequently, we can base a BCI on electrical potentials, magnetic fields, or metabolic/hemodynamic recordings.

Figure 1 presents a schematic of the principal BCI components. The components involve signal acquisition, preprocessing, feature extraction, classification, and an application interface together with the application. When we talk about a BCI, we must consider several component options. Signal recordings can be either invasive or noninvasive. Signal features require analysis and classification methods. Control functions require selecting a suitable mental strategy as well as operational and feedback mechanisms.

Suitable brain signals

Invasive BCI methods place electrodes directly on or inside the cortex. One method records electrical potentials for subsequent analysis of the electrocorticogram (ECoG). Another method places a multiunit electrode array in the cortex to record the neural firing of a small population of neurons. Both signal types have a superior signal-to-noise ratio, need little user training, and are suitable for replacing or restoring lost motor functions in patients with damaged parts of the neuronal system.

Noninvasive BCIs, on the other hand, can use a variety of brain signals as input, such as electroencephalograms (EEG), magnetoencephalograms (MEG), blood-oxygen-level-dependent (BOLD) signals, and (de)oxyhemoglobin concentrations. The EEG, which is basically the sum of many postsynaptic potentials in the cortex, is the most widely used brain signal for operating a BCI system. We can extract two types of changes from the ongoing EEG signals: one is time- and phase-locked (evoked) to an externally or internally paced event, while the other is also time-locked but not phase-locked (induced). To the former class belong the event-related potentials (ERPs), including the P300, steady-state visual evoked potentials (SSVEPs), and slow cortical negative shifts; to the latter class belong the event-related desynchronizations (ERDs) and event-related synchronizations (ERSs).

The MEG can measure brain activity by detecting weak magnetic fields caused by current flows in the cortex. These small magnetic fields in the picotesla to femto-tesla range are measured with multichannel SQUID (superconducting quantum interference device) gradiometers in a shielded environment. This technique combines excellent time resolution with good spatial




Figure 1. The brain-computer interface: (a) schematic of principal BCI components and (b) three applications: playing table tennis (top), using a spelling system (middle), and restoring grasp functions (bottom).

resolution, which can be as fine as 2 to 3 millimeters. Researchers have studied BCIs using MEG data, but they haven't been able to demonstrate significant advantages over EEG-based systems.

Unlike EEG and MEG systems, which detect the electromagnetic activity of cortical neurons, near-infrared spectroscopy (NIRS) measures the metabolic activity of specific cortical regions. NIRS uses light in the near-IR spectrum (typically between wavelengths of 630 to 1,350 nm) to determine the oxygenation of the tissue, and researchers have recently applied it to BCI research. The potential advantages of realizing a BCI with this technique include its insensitivity to typical EEG artifacts such as the electrooculogram (EOG), electromyogram (EMG), and electrode failures. However, the technique also requires several seconds to pass before it can measure the metabolic response, which is a long time compared to EEG and MEG. The spatial resolution also lies in the centimeter range.³

Like NIRS, functional magnetic resonance imaging (fMRI) measures the metabolic changes in the brain. Based on traditional MRI principles, the fMRI neuroimaging technique can also be used to control a BCI. To measure the hemodynamic response, fMRI studies usually use the BOLD signal. The stimulus response time is in the range of some seconds.⁴

After the brain signals have been re-

corded (and possibly preprocessed in suitable ways), the next step is to extract prominent features that describe important discriminative signal properties. This processing stage aims simply to reduce data and adequately transform it such that the subsequent classification process is optimal. Example features used in EEG processing are the power in a specific frequency band (band power), autoregressive parameters, and synchronization measures.

Choosing the mental strategy

Operant conditioning is a learning process with the goal of self-regulating brain potentials (such as slow cortical potential shifts) or brain waves (such as sensorimotor rhythms) with the help of suitable feedback. This process doesn't require continuous feedback, but it does require a reward for achieving the desired brain potential (wave). Researchers have used operant conditioning to realize a communication system for completely paralyzed ("locked-in") patients.

Another frequently used mental strategy is motor imagery. Research results from this strategy provide strong evidence that motor imagery activates cortical areas similar to those activated by executing the same movement. Consequently, we place the EEG electrodes over the primary sensorimotor areas. When a user learns such a

motor imagery task in a number of training sessions, characteristic ERD/ERS patterns are associated with different types of motor imagery and detectable in single trials in an online system.

Other mental tasks besides motor imagery are suitable to modulate the brain signals—for example, mental arithmetic and imaging the rotation of geometric objects. Focused attention or gaze control on visual stimuli, such as flickering lights or flashed letters, is especially suitable to realize spelling devices with a P300-based BCI or to control neuroprostheses with a SSVEP-based BCI.

Self-based and cue-based BCI systems

The mode of operation determines the type of data processing, either in a predefined time window of some seconds following a cue stimulus (synchronous BCI) or continuously sample-by-sample (asynchronous BCI). The cue might contain information for users (for example, it might let them know whether they should imagine moving the left or right hand during training), or it might be neutral. In the latter case, the users are free to choose one of the predefined mental tasks after the cue.

A synchronous BCI system is not available for control outside the cue-based processing window. In the asynchronous mode, no cue is necessary, so the system is continuously available to the users. They can decide freely when they wish to generate a control signal. Such a system is more complex and demanding, and the great challenge is to maximize the intentional control (true positives) while minimizing the nonintentional control (false positives) at the output. We used such an asynchronous BCI successfully to operate a spelling device and to navigate in a virtual environment.

Organizing training and feedback

To employ a BCI successfully, users must first go through several training sessions to obtain control over their brain potentials (waves) and maximize the classification accuracy of different brain states. In general, the training starts with one or two predefined mental tasks repeated periodically in a cue-based mode. In predefined time windows after the cue, we record the brain signals and use them for offline analyses. In

this way, the computer learns to recognize the users' mental-task-related brain patterns. This learning process is highly subject-specific, so each user must undergo the training individually. The learning phase produces a classifier that we can use to classify the brain patterns online and provide suitable feedback to the users. Visual feedback has an especially high impact on the dynamics of brain oscillations that can facilitate or deteriorate the learning process. The training phase is relatively short with P300 or SSVEPs, but can last weeks or even months with mental tasks.²

References

1. J.R. Wolpaw et al., "Brain-Computer Interfaces for Communication and Control," *Clinical Neurophysiology*, vol. 113, no. 6, 2002, pp. 767–791.
2. G. Pfurtscheller, C. Neuper, and N. Birbaumer, "Human Brain-Computer Interface," *Motor Cortex in Voluntary Movements*, A. Riehle and E. Vaadia, eds., CRC Press, 2005, pp. 367–401.
3. S. Coyle et al., "On the Suitability of Near-Infrared (NIR) Systems for Next-Generation Brain-Computer Interfaces," *Physiological Measurement*, vol. 25, no. 4, 2004, pp. 815–822.
4. N. Weiskopf et al., "Principles of a Brain-Computer Interface (BCI) Based on Real-Time Functional Magnetic Resonance Imaging (fMRI)," *IEEE Trans. Biomedical Eng.*, vol. 51, no. 6, 2004, pp. 966–970.

Gert Pfurtscheller was professor of medical informatics and is head of the Brain-Computer Interface Lab at the Graz University of Technology's Institute for Knowledge Discovery. Contact him at pfurtscheller@tugraz.at.

Clemens Brunner is a postdoctoral researcher at the Brain-Computer Interface Lab at the Graz University of Technology's Institute for Knowledge Discovery. Contact him at clemens.brunner@tugraz.at.

Brain-Controlled Robots

José del R. Millán, *IDIAP Research Institute and École Polytechnique Fédérale de Lausanne*

The idea of moving robotic or prosthetic devices not by manual control but by mere "thinking"—that is, by human brain activity—has fascinated researchers for the past 30 years. But only now have experiments

shown the possibility of doing so.

How can brainwaves directly control external devices? The current focus is mainly on invasive approaches that provide detailed, single-neuron activity recorded from microelectrodes implanted in the brain.¹ The motivation for invasive approaches is broad evidence that ensembles of neurons in the brain's motor system—motor, premotor, and posterior parietal cortex—encode the parameters related to hand and arm movements in a distributed, redundant way.

For humans, however, noninvasive approaches avoid health risks and associated ethical concerns. Most noninvasive brain-computer interfaces (BCIs) use electroencephalogram (EEG) signals—electrical brain activity recorded from electrodes on the scalp. The EEG's main source is the synchronous activity of thousands of cortical neurons. Thus, EEG signals suffer from a reduced spatial resolution and increased noise when measurements are taken on the scalp. Consequently, current EEG-based brain-actuated devices are limited by low channel capacity and are considered too slow for controlling rapid and complex sequences of robot movements.

Recently, however, my coworkers and I at the IDIAP Research Institute and the École Polytechnique Fédérale de Lausanne have shown for the first time that online EEG signal analysis, if used in combination with advanced robotics and machine learning techniques, is sufficient for humans to continuously control a mobile robot² and a wheelchair.³

Spontaneous EEG and asynchronous operation

We can classify noninvasive EEG-based BCIs as *evoked* or *spontaneous*. An evoked BCI exploits a strong characteristic of the EEG, the *evoked potential*, which reflects the immediate automatic responses of the brain to some external stimuli. Examples of evoked potentials include P300 and SSVEP (steady-state visual evoked potentials). In principle, evoked potentials are easy to detect with scalp electrodes. However, evoking them requires external stimulation, so they apply to only a limited task range.

In my view, a more natural and suitable alternative for interaction begins with analyzing components associated with spontaneous, *intentional* mental activity. This is particularly the case for controlling robotics devices. As in driving a car, the

subject's attention must focus on driving and not on external stimuli.

Spontaneous BCIs are based on the analysis of EEG phenomena associated with various aspects of brain function related to mental tasks that the subject carries out at will. For example, the subject might imagine limb movements, such as the right or left hand, or cognitive operations, such as arithmetic or language.

But volunteer mental control isn't enough for steering a wheelchair or a prosthesis. These tasks require subjects to also make self-paced decisions. In such asynchronous protocols, the subject can deliver a mental command at any moment without waiting for external cues.^{2,4} This contrasts with synchronous interaction, where the EEG is time-locked to externally paced cues. Only asynchronous controls can send the appropriate mental command at the right time to make the wheelchair turn and cross the desired doorway while it's moving continuously.

The statistical machine learning way

Training is a critical BCI development issue—that is, how do users learn to operate the BCI? Like other groups,^{5,6} we follow a mutual-learning approach to facilitate and accelerate the user's training period. The user and the BCI are coupled together and adapt to each other. In other words, we use machine learning approaches to discover the individual EEG patterns characterizing the mental tasks users execute while learning to modulate their brainwaves in a way that will improve system recognition of their intentions.

We use statistical machine learning techniques at two levels: selecting the features and training the classifier embedded in the BCI. In particular, the statistical classifier achieves error rates below 5 percent for three mental tasks, but correct recognition is 70 percent. In the remaining cases, the classifier doesn't respond because it considers the EEG samples to be uncertain.

Incorporating rejection criteria to avoid making risky decisions is an important BCI concern. From a practical viewpoint, a low classification error is a critical BCI performance criterion. Otherwise, users can become frustrated and stop using it. Furthermore, not executing probable wrong commands increases the BCI's theoretical bit rate and improves the robot's trajectories. The subject won't need to correct



Figure 2. A brain-actuated wheelchair. The subject guides the wheelchair through a maze, using a BCI that recognizes the subject's intent from analysis of noninvasive EEG signals. (photo courtesy of the MAIA project)

wrong turns or bring back the wheelchair to the desired doorway.

A blending of intelligences

How is it possible to control a robot that must make accurate turns at precise moments using signals that arrive at a rate of about one bit per second?

The key aspect of our brain-actuated robots is combining the subject's mental capabilities with the robot's intelligence. That is, the subject delivers a few high-level mental commands (for example, "Turn right at the next occasion"), and the robot executes these commands autonomously using the readings of its onboard sensors. In other words, the EEG conveys the subject's intent, and the robot performs it to generate smooth, safe trajectories.

This approach makes it possible to continuously control a mobile robot—emulating a motorized wheelchair—along nontrivial trajectories requiring fast and frequent switches between mental tasks.² In a few days, two human subjects learned to mentally drive a robot between rooms in a house-like environment and visit three or four rooms in a prescribed order. Furthermore, when the subjects later controlled the robot manually along the same trajectories,

the performance was only marginally better than the mental performance.

More recently, we extended this work to the mental control of both a simulated and a real wheelchair (see figure 2).³ We performed this work in the framework of the European project MAIA (Augmentation through Determination of Intended Action, www.maia-project.org) and in cooperation with Katholieke Universiteit Leuven. In this case, we incorporated shared-control principles to blend the two intelligences.⁷ Although our first brain-actuated robot had a form of cooperative control, shared control is a more principled, flexible framework and gives users a finer degree of control.

Challenges and future research directions

For brain-actuated robots, in contrast to augmented communication through BCI, fast decision making is critical. In this sense, real-time control of brain-actuated devices, especially robots and neuroprostheses, is the most challenging BCI application. While researchers have demonstrated brain-actuated robots in the laboratory, the technology isn't yet ready for use in real-world situations. We still need to improve the BCI's robustness to make it a more practical and reliable technology.

A first line of research is online adaptation of the interface to the user to keep the BCI constantly tuned to its owner.⁸ This would account for the new capabilities—and corresponding new brain signals—that subjects gain with experience. In addition, brain signals change naturally over time. In particular, they can change from one session that supplies the data to train the classifier to the next session that applies the classifier. Online learning can help adapt the classifier throughout its use and keep it tuned to drifts in the signals it receives in each session.

The second line is the analysis of neural correlates of high-level cognitive and affective states such as errors, alarms, attention, frustration, and confusion. The EEG has information about these states embedded in it, together with the mental commands intentionally generated by the user. The ability to detect and adapt to these states would enable the BCI to interact with the user in a much more meaningful way. One of these high-level states is the awareness of erroneous responses. The neural correlate for this awareness arises in the millisecond range, so user commands are executed only if no

error is detected in this short time frame. Recent results have shown satisfactory single-trial error recognition that significantly improves BCI performance.⁹ In addition, this new type of error potential—which is generated in response to errors made by the BCI rather than by the user—can provide performance feedback that, in combination with online adaptation, improves the BCI while it's being used.

Acknowledgments

The Swiss National Science Foundation supported this work through the National Centre of Competence in Research on Interactive Multimodal Information Management and also by the European Information Society Technologies Programme, Future and Emerging Technologies Project FP6-003758. The article reflects only the author's views, and funding agencies aren't liable for any use that might be made of the information it contains.

References

1. J.M. Carmena et al., "Learning to Control a Brain-Machine Interface for Reaching and Grasping by Primates," *PLoS Biology*, vol. 1, no. 2, 2003, pp. 193–208.
2. J.d.R. Millán et al., "Noninvasive Brain-Actuated Control of a Mobile Robot by Human EEG," *IEEE Trans. Biomedical Eng.*, vol. 51, no. 6, 2004, pp. 1026–1033.
3. F. Galán et al., "An Asynchronous and Non-Invasive Brain-Actuated Wheelchair," *Proc. 13th Int'l Symp. Robotics Research*, 2007, pp. 45–54.
4. J.d.R. Millán, "Adaptive Brain Interfaces," *Comm. ACM*, vol. 46, no. 3, 2003, pp. 74–80.
5. B. Blankertz et al., "The Berlin Brain-Computer Interface: Machine Learning Based Detection of User Specific Brain States," *J. Universal Computer Science*, vol. 12, no. 6, 2006, pp. 581–607.
6. G. Pfurtscheller and C. Neuper, "Motor Imagery and Direct Brain-Computer Communication," *Proc. IEEE*, vol. 89, no. 7, 2001, pp. 1123–1134.
7. D. Vanhooydonck et al., "Shared Control for Intelligent Wheelchairs: An Implicit Estimation of the User Intention," *Proc. 1st Int'l Workshop Advances in Service Robotics*, Fraunhofer IRB Verlag, 2003, pp. 176–182.
8. J.d.R. Millán et al., "Adaptation in Brain-Computer Interfaces," *Towards Brain-Computer Interfacing*, G. Dornhege et al., eds., MIT Press, 2007, pp. 303–325.
9. P.W. Ferrez and J.d.R. Millán, "Error-Related EEG Potentials Generated during Simulated Brain-Computer Interaction," *IEEE Trans. Biomedical Eng.*, vol. 55, no. 3, 2008, pp. 923–929.

José del R. Millán is an adjunct professor at the Swiss Federal Institute of Technology in Lausanne (EPFL) and a senior researcher at the IDIAP Research Institute. Contact him at jose.millan@idiap.ch.

Why Use a BCI If You're Healthy?

Brendan Allison and Bernhard Graimann, *University of Bremen*

Most brain-computer interface (BCI) research focuses on restoring communication for severely disabled users.^{1,2} However, BCIs could also treat disabilities such as stroke, autism, epilepsy, or emotional disorders, and they might even become useful to healthy users.^{3,4} At present, BCIs have several serious drawbacks relative to conventional interfaces such as keyboards or mice. They're much slower, less accurate, and operational only at very low bandwidths. They require cables and unfamiliar, expensive hardware, including an electrode cap. The cap requires hair gel and several minutes of preparation and cleanup. Some BCIs require training, are difficult to use, and fail with some subjects or in noisy environments. BCIs often seem intimidating, exotic, Orwellian, or even nerdy. They rarely show up in mainstream markets, and this won't change soon.

Hence, the prevailing view is that BCIs, at best, enable people to send the same information available much more quickly and easily via other interfaces. This perspective is wrong. Here, we'll discuss why healthy people might eventually use BCIs in specific situations. We'll consider BCIs with scalp-mounted electrodes because other neuroimaging approaches are typically impractical.^{1,3,5}

BCIs for healthy users

A few BCI R&D projects envisioned healthy subjects as end users. Modern BCI simulations or games usually allow one or two degrees of freedom or 1D to 2D graded control. Turning, moving, or leaning are often possible, sometimes in a virtual environment. For example, researchers have demonstrated BCIs intended to let healthy users navigate maps while their hands are busy.^{6,7} Game companies such as NeuroSky and Emotiv advertise games that allow people to move a character with

conventional handheld controls and control special features through a BCI.

New BCI subjects sometimes perform effectively within about 10 minutes despite background distraction and electrical noise, but researchers haven't yet studied the effects of intensive usage as might occur in gamers.^{1–3} Nor have they fully studied the precision and timing of translating user intent into control signals through BCIs.

Typical research BCIs allow communication only via electrodes and so exhibit very low bandwidth. Hybrid interfaces could combine BCIs with other interfaces to provide an additional independent signal or modify other commands,¹ which might allow moving while crouching, dodging, firing, communicating, spellcasting, and/or mentally levitating an object.

The BCI "distraction quotient" is unknown in these scenarios. How can BCIs best be integrated with other interfaces? Which BCIs work best with other interfaces, environments, and games? How do these issues vary across users with different personalities, backgrounds, motivations, abilities, experience, training, and other characteristics? These questions will become increasingly important as pressure to build a practical BCI mounts from commercial sources.

Induced disability

Healthy users might communicate via BCIs when conventional interfaces are inadequate, unavailable, or too demanding. Surgeons, mechanics, soldiers, cell phone users, drivers, and pilots can experience induced disability when hand or voice communication is infeasible. BCIs might help them request tools, navigate maps or schematics, access data, or perform otherwise difficult, distracting, dangerous, or impossible tasks.

Hybrid interfaces could also help when conventional interfaces provide insufficient bandwidth. Expert gamers often use many keys at once. Console games require using several fingers on both hands. A major benchmark will be the first BCI that reliably provides supplemental information without impairing mainstream interface performance.

Ease of use in hardware

The keyboard and mouse seem like natural, intuitive, convenient interfaces—when expert users just happen have them handy. Users who wear electroencephalography (EEG) sensors might find BCIs easier to

use. EEG sensor technology is becoming more practical.¹ New electrodes require little or no gel, scalp contact, or preparation and cleanup time. As electronics and signal processing improve, smaller, better, cheaper sensors and amplifiers could operate with devices or clothing on or near the head. Bluetooth, the ubiquitous wireless Internet, and related technologies facilitate wireless BCIs. BCIs might eventually become more convenient and accessible than cell phones, watches, remote controls, or car dashboard interfaces.

Laziness is the wayward child of invention. Laziness can induce disability, and it can be very motivating. Although television has viable interfaces, people typically prefer more portable alternatives that provide no advantage except remote control.

BCIs could also help people who retype words or sentences (rather than cut and paste via mice) by letting them instead select, drag, or click via the BCI, thus avoiding temporarily disengaging from the keyboard. BCIs could allow sending messages without the hassle of a keyboard, microphone, or cellphone numberpad. Humanity might finally escape the various inconveniences of finding handheld interfaces or pressing buttons.

Ease of use in software

The activities that control most BCIs and conventional interfaces differ fundamentally from desired outputs. Noticing flashes or moving fingers across a keyboard isn't like natural communication. However, some BCIs allow walking or turning by imagining foot or hand movements,^{2,7} and these might offer new frontiers of usability for all users. As with other interfaces, research should address which mental activities seem most natural, easy, and pleasant for different users in different situations.

Otherwise unavailable information

Available interfaces have heavily influenced all software. Operating systems would look very different if eye trackers and voice commands were the dominant interfaces. Just as keyboards and mice are inherently suited to typing and dragging, BCIs are inherently better suited to certain tasks. The error-related negativity and P300 that often develop after a subject recognizes a mistake could allow real-time error recognition.¹ The P300, steady-state visual evoked potential (SSVEP), and other signals reflect

regional attention. Software might magnify, link, remember, or jump to interesting areas of the screen or auditory space. EEG-based assessment of global attention, frustration, alertness, comprehension, exhaustion, or engagement could enable software that adapts much more easily to the user. The challenge of developing new opportunities for integrating BCI-based signals into conventional and emerging operating systems might be as fun as Douglas Engelbart's daunting task of integrating the mouse into a world then dominated by keyboards.

Improved training or performance

Some BCIs train subjects to produce specific activity over sensorimotor areas, so BCI training might improve movement training or performance. Subjects' athletic and motor background and skills might influence BCI parameters. These avenues might be useful for motor rehabilitation or finding the right BCI for each user.^{3,4}

Confidentiality

BCIs might be the most private communication channel possible. With other interfaces, eavesdropping simply requires observing the necessary movements. This important security problem also shows up in competitive gaming environments. For example, many console gamers have chosen an offensive football play, then noticed an adjacent opponent select a corresponding defensive play after overt peeking.

Speed

Relevant EEGs are typically apparent one second before a movement begins and might precede the decision to move.¹ Future BCIs might be faster than natural pathways. Further research should provide earlier movement prediction with greater precision and accuracy, integrate predicted with actual movements smoothly, and evaluate training and side effects.

Novelty

Some people might use a BCI simply because it seems novel, futuristic, or exciting. This consideration, unlike most others, loses steam over time. BCIs will become more flexible, usable, or better hybridized as research continues. However, as BCIs improve, public perception will follow a pattern reminiscent of microwaves and cell phones. BCIs will first be exotic, then novel, widespread, unexceptional, and finally boring.

Healthy target markets

Most healthy BCI users today are research scientists, friends, research subjects, and visitors at expositions. A few people order commercial BCIs, forming a crucial fifth category in which no BCI expert prepared the software or hardware for individual users.

Gamers are likely early adopters. They often wear headgear, enjoy novelty and technical challenges, have money and time available for peripherals and training, and are competitive and increasingly numerous. Specific military or government personnel follow technology validated elsewhere. Highly specialized users such as surgeons, welders, or mechanics are also likely second-generation adopters. Electrooculograms, electromyograms, electrocardiograms, and other signals might supplement EEG control in many BCI and related applications.

More mainstream applications, such as error correction hybridized with word processors, are more distant. These approaches require new software development, much better EEG sensors, and encouraging validation. BCIs might instead seem unreliable, useless, unfashionable, dangerous, intrusive, or oppressive, spurred by inaccurate reporting. Websites such as bci-info.org, proper dissemination of results, and positive appearances at conferences, expositions, interviews, or other events can educate people and reduce miscommunication.

BCIs won't soon replace conventional interfaces, but they might be useful to healthy users in specific situations. Integrating them with other interfaces raises many questions best addressed with parametric research involving different users, interfaces, mental activities, goals, output devices, and training parameters.

References

1. B.Z. Allison, E.W. Wolpaw, and J.R. Wolpaw, "Brain-Computer Interface Systems: Progress and Prospects," *Expert Rev. of Medical Devices*, vol. 4, no. 4, 2007, pp. 463–474.
2. G. Pfurtscheller et al., "15 years of BCI Research at Graz University of Technology: Current Projects," *IEEE Trans. Neural Systems and Rehabilitation Eng.*, vol. 14, no. 2, 2006, pp. 205–210.
3. N. Birbaumer and L.G. Cohen, "Brain Computer Interfaces: Communication and Restoration of Movement in Paralysis," *J. Physiology*, vol. 579, pt. 3, 2007, pp. 621–636.
4. B. Graimann, B.Z. Allison, and A. Gräser, "New Applications for Non-invasive Brain-Computer Interfaces and the Need for Engaging Training Environments," *Proc. Int'l*

Conf. *Advances in Computer Entertainment Technology*, ACM, 2007, pp. 25–28.

5. J.R. Wolpaw et al., “BCI Meeting 2005: Workshop on Signals and Recording Methods,” *IEEE Trans. Neural Systems and Rehabilitation Eng.*, vol. 14, no. 2, 2006, pp. 138–141.
6. L.J. Trejo, R. Rosipal, and B. Matthews, “Brain-Computer Interfaces for 1D and 2D Cursor Control: Designs Using Volitional Control of the EEG Spectrum or Steady-State Visual Evoked Potentials,” *IEEE Trans. Neural Systems and Rehabilitation Eng.*, vol. 14, no. 2, 2006, pp. 225–229.
7. R. Scherer, G.R. Müller-Putz, and G. Pfurtscheller, “Self-Initiation of EEG-Based Brain-Computer Communication Using the Heart Rate Response,” *J. Neural Eng.*, vol. 4, 2007, pp. L23–L29.

Brendan Allison is a researcher at the University of Bremen’s Institute of Automation. Contact him at allison@iat.uni-bremen.de.

Bernhard Graimann is a researcher at the University of Bremen’s Institute of Automation. Contact him at graimann@iat.uni-bremen.de.

Computational Challenges for Noninvasive Brain Computer Interfaces

Florin Popescu, *Fraunhofer Institut für Rechnerarchitektur und Softwaretechnik (FIRST) Berlin*

Benjamin Blankertz and Klaus-R. Müller, *Berlin Institute of Technology*

Electroencephalography (EEG) is unique among functional brain-imaging methods in that it promises a means of providing a cost-efficient, safe, portable, and easy-to-use brain-computer interface (BCI) for both healthy users and the disabled. An already-extensive corpus of experimental work has demonstrated that, to a degree, EEG-based BCI can detect a person’s mental state in single trials of mental imagination using sophisticated mathematical tools; but this work has also outlined clear challenges. The first challenge is the rather limited information transfer rate (ITR) achievable through EEG, which is—in the most optimistic of cases—about an order of magnitude lower than invasive BCI methods currently provide. That said, the potential benefits of brain implant-based BCI haven’t yet proved worth the associated cost and risk in the most disabled patients, let alone healthy users.

EEG seems for now the only practical brain-machine interaction choice (cost and ITR limitations hamper other noninvasive methods). As such, we ask here not how further signal-processing and machine-learning improvements might increase the ITR.^{1,2} BCI researchers already know that many complex technical problems remain: such problems have been the field’s main concern up to now. Nor will we discuss EEG-BCI applications. Instead, we concentrate on outlining the challenges that remain in adapting EEG-BCI from the laboratory to real-world use by healthy subjects.

Dry electrodes

The most elementary EEG-BCI challenge for healthy users isn’t—at first glance—a computational one. Standard EEG practice involves the tedious application of conductive gel on EEG electrodes to provide accurate measurements of the microvolt-level scalp potentials that constitute EEG signals. Without “dry-cap” technology, the proper set-up of BCI sessions in, say, a home environment, is too tedious and messy to be practical. Some dry electrode designs that use a combination of EEG and electromyogram (EMG) have been announced for home entertainment use. The EMG originates from body and face muscles; in BCI studies, it’s considered an artifact. Although EMG is stronger and easier to read than EEG, it doesn’t truly constitute a mental interface. Our research group has developed an EEG-BCI dry-cap design and tested its performance (and the absence of muscle artifacts) in a controlled study.³

For ease-of-use and cost reasons, all foreseeable systems will use fewer electrodes than found on standard EEG caps today. The computational challenges we’ve addressed include optimal placement of the reduced number of electrodes and robustness of BCI algorithms to the smaller set of recording sites. With only six unipolar electrodes, we can achieve about 70 percent of full-gel-cap BCI performance at sites above the motor cortex, while being able to discount any potential influence of muscle and eye movement artifacts.

Most other remaining dry-cap challenges are of an engineering design nature, excluding perhaps the computational reduction of artifacts produced not by unrelated electrophysiological activity but by measured low-frequency voltage variations caused by the head’s physical movement.

BCI illiteracy

A long-standing problem of BCI designs that detect EEG patterns related to a voluntarily produced brain state is that such paradigms work with varying success among different subjects or patients. We distinguish mental-task-based BCI, such as “movement imagination” BCI, from paradigms based on involuntary stimulus-related potentials such as P300. These stimulus-related potentials are limited to very specific applications, such as typing for locked-in patients, and they require constant focus on stimuli extraneous to the task at hand.

In a recent study, with 10 untrained users,² our research group took a close look at how fast the users achieved their best performance (by skill acquisition) during a small number of BCI sessions and how much this performance varied among subjects. We confirmed the results in a follow-up study with 13 novice subjects.⁴ Although machine learning techniques allow use of minimal calibration data recording (< 20 minutes) before the BCI system is ready to use, the subjects’ peak-performance plateaus, even after multiple sessions, varied greatly. Using this and other unreported data by many research groups, we estimate that

- about 20 percent of subjects don’t show strong enough motor-related mu-rhythm variations for effective asynchronous motor-imagery BCI,
- another 30 percent exhibit slow performance (< 20 bits per minute), and
- up to 50 percent exhibit moderate to high performance (20–35 bits/min.).

It’s still a matter of debate as to why BCI systems exhibit “illiteracy” in a significant minority of subjects and what can be done about it in terms of signal processing and machine learning algorithms. From internal investigations (as well as the results of BCI Competition II, data set Ib⁵), BCI illiteracy in a subject appears to depend not so much on the algorithm used but on a property inherent in the subject.

EEG is sensitive to sources in cortical folds, so it might not be able to read motor-imagery activity in some subjects because the particular cortical region involved is tangential to the scalp. An observation consistent with this explanation is that in certain subjects some *classes*—that is, types of imagined movements—are detectable and

others not. Calibration sessions should therefore select subject-specific classes along with frequency bands necessary for feature generation to minimize the illiteracy problem.

Idle class

Most commonly, BCI controllers involve two classes, which can move a monitor-displayed cursor toward, say, left and right. Although these controllers can perform asynchronously—that is, at their own independent pace—useful cursor control is difficult. The user must either continuously imagine one of the two classes or lose control of the cursor.

Besides self-pacing, BCI would greatly benefit from integrating an “idle” or “rest” class with the BCI’s active classes—that is, those corresponding to mentally imagining a particular task and implying the desire to transmit the activation of a corresponding command. This would keep the cursor from responding when no active class (from a set of two or more) is activated.

The idle state might take one of two forms: a relax state, where the subject stays still and tries to “think of nothing,” or a state where the subject can do almost any mental task other than those belonging to the active classes. In the case of deliberate relaxation, usability is obviously limited, although signal processing is easier, given that relaxation tends to increase EEG power in the alpha band. For example, researchers have shown that alpha band modulation played a strong role in detecting relaxation when subjects closed their eyes during an idle state.⁶

Relying on alpha power modulation is complicated by the involuntary variation of background alpha in physiological as opposed to experimental conditions—for example, due to fatigue. Furthermore, relaxing itself induces drowsiness.

A neurofeedback-style, low-frequency modulation approach has shown promise as an idle-state paradigm, but it requires intensive subject training, exhibits limited ITR, and has only one active class.⁷ The Graz group has begun work toward idle-state control with a relax cue,⁸ but so far there is little hard data on idle-state duration and accuracy. This is important, because two-class classifier output (a noisy signal) is usually integrated until it hits a threshold (for example, left or right cursor movement). The level of this threshold

offers a clear trade-off between high idle-class accuracy (that is, the thresholds are high) and fast speed of response or high ITR (the thresholds are low). Remaining challenges are to find a classifier that can induce a rest state without a relax cue and to optimize the relationship between classifier output and BCI command. Because of physiological variations in background EEG activity, where fatigue is a main factor, we believe an adaptive classifier and controller are necessary for maximal performance. Our group has undertaken some efforts toward optimizing a true idle-state BCI paradigm by balancing idle-class accuracy and ITR.⁹

Future challenges and implementations

While these three computational challenges are, we believe, the most urgent, other improvements might also be beneficial. Although 20 minutes of calibration for a novel subject isn’t excessive, usability would benefit from knowing the minimal number of calibration trials needed to achieve moderate performance and rule out BCI illiteracy, such that a classifier can then adapt to the user during normal use. For applications such as gaming, or voluntary self-paced interaction with an unstructured environment, this adaptation should work even in cases where class labels aren’t available (unsupervised adaptation).¹⁰

We envisage an EEG BCI scenario in which users purchase an affordable computer peripheral that is simply placed on the head and requires no gel. New users will undergo a one-time calibration procedure that takes maximally 10 minutes, ideally even less. They then proceed to use the BCI system in a game environment to, for example, control a robot or wheelchair. The system’s performance slowly adapts to the user’s brain patterns, reacting only when he or she intends to control it. At each repeated use, the system recalls parameters from previous sessions, so recalibration is rarely, if ever, necessary.

We strongly believe such a system, capable of an average performance of about 15 to 20 bits/min, is achievable within the next few years. Challenges such as BCI illiteracy are likely to be only partially met. Still, if this percentage decreases further, it shouldn’t prevent noninvasive BCI systems from reaching a large user population, healthy or disabled. ■

References

1. B. Blankertz et al., “The Noninvasive Berlin Brain-Computer Interface: Fast Acquisition of Effective Performance in Untrained Subjects,” *NeuroImage*, vol. 37, no. 2, 2007, pp. 539–550.
2. G. Dornhege et al., eds., *Toward Brain-Computer Interfacing*, MIT Press, 2007, p. 83.
3. F. Popescu et al., “Single Trial Classification of Motor Imagination Using Six Dry EEG Electrodes,” *PLoS ONE*, vol. 2, 2007, p. e637.
4. B. Blankertz et al. “The Berlin Brain-Computer Interface: Accurate Performance from First-Session in BCI-Naive Subjects,” to be published in *IEEE Trans. Biomedical Eng.*, 2008.
5. B. Blankertz et al., “The BCI Competition 2003: Progress and Perspectives in Detection and Discrimination of EEG Single Trials,” *IEEE Trans. Biomedical Eng.*, vol. 51, no. 6, 2004, pp. 1044–1051.
6. J.d.R. Millán and J. Mouriño, “Asynchronous BCI and Local Neural Classifiers: An Overview of the Adaptive Brain Interface Project,” *IEEE Trans. Neural System Rehabilitation Eng.*, vol. 11, no. 2, 2003, pp. 159–161.
7. J.F. Borisoff et al., “Brain-Computer Interface Design for Asynchronous Control Applications: Improvements to the LF-ASD Asynchronous Brain Switch,” *IEEE Trans. Biomedical Eng.*, vol. 51, no. 6, 2004, pp. 985–992.
8. G.R. Muller-Putz et al., “Brain-Computer Interfaces for Control of Neuroprostheses: From Synchronous to Asynchronous Mode of Operation,” *Biomedizinische Technik (Berl)*, vol. 51, 2006, pp. 57–63.
9. S. Fazli et al., “Asynchronous, Adaptive BCI Using Movement Imagination Training and Rest-State Inference,” *Proc. Artificial Intelligence and Applications (AIA 08)*, ACTA Press, 2008, pp. 85–90.
10. M. Krauledat et al., “Reducing Calibration Time for Brain-Computer Interfaces: A Clustering Approach,” *Proc. Advances in Neural Information Processing Systems (NIPS 06)*, vol. 19, MIT Press, 2007, pp. 753–760.

Florin Popescu is a research scientist at Fraunhofer FIRST’s Intelligent Data Analysis Group. Contact him at florin.popescu@first.fraunhofer.de.

Benjamin Blankertz is a researcher in the Machine Learning Laboratory at the Berlin Institute of Technology and in the Intelligent Data Analysis group at Fraunhofer FIRST. Contact him at blanker@cs.tu-berlin.de.

Klaus-R. Müller is a professor and chair of the Machine Learning Department at the Berlin Institute of Technology and head of the Intelligent Data Analysis group at Fraunhofer FIRST. Contact him at krm@cs.tu-berlin.de.