Laurens R. Krol, Lena M. Andreessen, Thorsten O. Zander

Passive brain–computer interfaces

A perspective on increased interactivity

WISSEN IM ZENTRUM
UNIVERSITÄTSBIBLIOTHEK

Technische
Universität
Berlin

Passive Brain–Computer Interfaces: A Perspective on Increased Interactivity

Laurens R. Krol (Team PhyPA, Biological Psychology and Neuroergonomics, Technische Universität Berlin, lrkrol@gmail.com),

Lena M. Andreessen (Team PhyPA, Biological Psychology and Neuroergonomics, Technische Universität Berlin, lena.andreessen@gmail.com), and

Thorsten O. Zander (Team PhyPA, Biological Psychology and Neuroergonomics, Technische Universität Berlin, thorsten.zander@tu-berlin.de)

Abstract

Passive brain–computer interfaces (passive BCI; pBCI) have been introduced and formally defined almost a decade ago and have gained considerable attention since then. In this chapter, we clarify some points of confusion and provide a perspective on the past, present, and future of the field of passive BCI. This perspective concerns a key aspect with regard to which various pBCI-based systems differ from each other: interactivity. The more interactive a system is, the more responsive it is, the more autonomous, and the better capable of adaptation. Along these lines, we identify and describe four relevant categories of systems with varying levels of interactivity: mental state assessment, open-loop adaptation, closed-loop adaptation, and automated adaptation. We give examples of past and current research for each of these categories. The latter three are collectively introduced as neuroadaptive systems. This perspective and formal categorisation helps to highlight human–computer interaction aspects that are relevant for the design of pBCI-based systems and points to possibilities for future research and development into passive BCI, implicit interaction, and neuroadaptive technology.

3.1 Passive Brain–Computer Interfaces

In 1977, Jacques J. Vidal demonstrated the real-time detection of "brain events", and offered an example of how these might be used to control a system. Instead of using offline, a posteriori methods to investigate neuroelectric brain activity, Vidal suggested "treating the experiment as a signal detection problem" (Vidal 1977): with continuous access to an electroencephalogram (EEG), a computer classified incoming EEG data as belonging to one of four categories, based on previously learned (and continuously updated) decision strategies. As such, the system was able to recognise specific brain activity in real time.

Although Vidal speculated upon a wide range of potential future applications of this approach, it was in a clinical context that *brain–computer interface* (BCI) methodology gained widespread attention. It was quickly realised that such direct BCIs could potentially provide a means for paralysed or otherwise motor-impaired people to communicate with the outside world without having to use their muscles (Wolpaw et al. 2002). This resolve to help those people who stood to benefit most from this technology led to a strong focus on BCI for direct communication and control, and resulted, inter alia, in different mental speller devices (e.g., Birbaumer et al. 1999; Brouwer & Van Erp 2010; Farwell & Donchin 1988; Furdea et al. 2009; Treder et al. 2011) and brain-activated prostheses (e.g., Müller-Putz & Pfurtscheller 2008; Soekadar et al. 2016; Wessberg et al. 2000; Wolpaw & McFarland 2008). For a long time, in fact, BCI research appeared to be inseparably connected to these applications.

Improvements made over the decades have made the field more interdisciplinary and BCI methodology more reliable. Of particular significance was the introduction of machine learning techniques at the start of the current millennium (e.g., Blankertz et al. 2002a; Lotte et al. 2007; Ramoser et al. 2000). Previously, users of BCI systems were trained to generate specific machine-detectable features in their EEG, such as a specific frequency modulation at a specific electrode site (Birbaumer et al. 1999). These days, training has shifted to the machine, using supervised machine learning based on initial recordings (Blankertz et al. 2002a). As pattern recognition by the machine replaced operant conditioning of the user, the idea was rekindled that this methodology could also be applied to detect and investigate spontaneous, automatic brain activity, and that these detections could then be used as *implicit input* to a system (Rötting et al. 2009; Zander et al. 2014). (We use "spontaneous" in its biological sense, meaning automatic, instinctive, involuntary, and inattentive; see below for a further discussion.)

Early examples of this can be found in suggestions to use error-related potentials to correct classification errors of traditional BCI systems (Blankertz et al. 2002b; Ferrez & Millán 2005, 2008). This was also suggested for response errors (Parra et al. 2003) and machine errors (Zander et al. 2008a) during human–computer interaction (HCI) in general. But it was only in 2008, when BCI enjoyed increased popularity overall, that the approach gained widespread attention. The use of BCI methodology for implicit input to benefit ongoing HCI was explicitly proposed as a worthwhile research endeavour by two research groups at the same conference (CHI 2008, Florence, Italy), independently of each other (Cutrell & Tan 2008; Zander et al. 2008a). Zander et al. (2008b) subsequently presented a framework formally separating this approach from traditional BCI research and proposed a name: *passive BCI*. The concept found acceptance also among clinical BCI researchers in 2012 (Wolpaw & Wolpaw 2012).

A passive BCI (Zander et al 2008b; Zander & Kothe 2011) system derives its output from automatic, spontaneous brain activity, interpreted in the given context (Zander & Jatzev 2012). This interpreted activity is then used as implicit input to support an ongoing task (Rötting et al. 2009). Use of the word "passive" here results from a user-centred perspective on HCI. It refers to the role of the end user of a system with respect to the BCI: the underlying signals being automatic, spontaneous brain activity, it is an inherent and defining aspect that the user exerts no effort to actively, explicitly, or voluntarily elicit or modulate this activity. Instead, the user focuses on the task at hand while a passive BCI system, in the background, monitors their brain activity for informative correlates of relevant cognitive or affective states.

Zander and Kothe (2011) contrast passive BCI with *active* BCI systems, where the brain activity in question is consciously and purposefully modulated in order to control an application (e.g., motor imagery BCIs; Pfurtscheller & Neuper 2001), and with *reactive* BCI systems, which rely on brain activity that is evoked by external stimulation but indirectly modulated through voluntary attention (e.g., P300 amplitudes modulated by attention shifts; Farwell & Donchin 1988).

Care should be taken when presuming that mental states can be readily categorized as "spontaneous" versus "voluntary," as required by these definitions of BCI systems. Similarly, the activity ultimately used by any BCI system may not precisely fit one such category. A user who is aware of a passive BCI system might be influenced by the expectations they have of that system and voluntarily commit attentional resources to make sure that the "spontaneous" activity takes place. A user might also attempt to consciously modulate this activity if results are not as expected. The other way around, an active BCI might rely on, or inadvertently make use of, brain activity that is not fully voluntarily controlled by the user. Meditation, as an example, seems to present an ambiguous mixture of both: it is the

voluntary attempt to induce a state that is usually only achieved when contextual factors align.

The active/reactive/passive distinction used here is a user-oriented one and focuses on conscious intentions. In meditation, the intended purpose of the activity is relaxation itself, not the communication of a state of relaxation for further processing. As a thought experiment, we could ask whether the same user behaviour and brain activity would be observed if the user was not aware of their influence over a system. If the answer is "yes," then we may say that this behaviour and brain activity represent "natural" human activity. The implicit input gathered from it is then dissociated from its function as input. This is taken as a core property of the technology discussed in this chapter: the underlying brain activity arises from natural human activity that is not purposefully modulated to cater to the BCI system. Note, however, that by this measure, the categorisation of a BCI system as (re)active or passive ultimately depends on the individual user and not on the system itself.

Indeed, the example of meditation can in fact also be used for an active BCI—when a state of relaxation is induced with the conscious intention of influencing a BCI system—or for reactive BCI, when, in addition, external stimuli are used to cue or achieve that state.

In the last decade, the passive BCI approach (i.e., using natural human brain activity as implicit input) has led to research into a range of different directions. Because of the general applicability of this approach to a wide variety of applications, a truly comprehensive overview is beyond the scope, and not the intention, of this chapter. Instead, we would like to discuss a trend that can be observed in recent work and thought related to passive BCI: a trend toward increased interactivity between the human user and the machine.

"Interactivity" denotes "the ability of a computer to respond to a user's input" (*Oxford English Dictionary*, Oxford University Press 2016). In the following sections, we will describe four categories of systems that all take a neurophysiological signal as their input, but respond to it with varying degrees of complexity. Each presented category represents an increase in interactivity compared to the previous one. Selected examples from past and current research are given to illustrate the different categories. A (partially) hypothetical example accompanies us throughout the chapter, suggesting how a surgeon in the operating theatre might be helped by a system from each category. In sequence, the given examples thus illustrate the above-mentioned trend toward increased interactivity. The chapter ends with a brief speculation of the future and a summarising conclusion.

3.2 Mental State Assessment

3.2.1 Introduction

BCI systems in general, by one definition, consist of three components: input, output, and a translation algorithm that produces the latter based on the former (Wolpaw et al. 2000). In this definition, input focuses on the measurement of particular aspects of brain activity, which reflect or correlate to a specific mental state. The detection of such a mental state (inferred from the measured activity) is then translated into a specific computer action (output), for example, the movement of a prosthetic arm or the selection of a letter on a screen.

Passive BCIs are distinct with respect to the exact cause of the mental state that is measured, requiring it to arise automatically as the result of a natural perception, activity, or thought, without the conscious intention of influencing the BCI system (see the discussion in Section 3.1). Emotions, for example, are not usually induced voluntarily but experienced as the result of something we perceive, do, or think. Other examples include stress, workload, vigilance, arousal, and surprise.

The measurement of such mental states can be helpful and informative by itself,

however, without any computer actions being initiated. Translation, in this case, merely quantifies the measurements but causes no further system changes that are reported back to the human who is being measured. A subset of traditional BCI methodology can thus be used to obtain information concerning a variety of mental states, i.e., for *mental state assessment* (Müller et al. 2008; van Erp et al. 2012; Zander & Kothe 2011). The resulting state information can then be analysed and studied further. The user is assumed to be behaving naturally, as discussed in Section 3.1.

To perform such an analysis, i.e. to calibrate such a "state detector" or *classifier*, individual calibration data is usually gathered in experimental environments where an operator's mental state can be carefully controlled and induced. Signal processing, feature extraction, and classifier calibration are then performed as per the general BCI approach, in order to later be able to assess, from a new recording in a different environment, to what extent the operator's state reflected the experimentally induced one(s). If, for example, a state of engagement can reliably be induced experimentally and neurophysiological correlates of that state can robustly be detected in those recordings, then an index can be generated that quantifies, for other recordings, to what extent these engagement-related features were present. This quantification can be done for complete recordings at once, but also moment by moment, identifying, for example, at what times during the recording the features were most or least pronounced.

In brief, mental state assessment takes a recording of neurophysiological activity and interprets it by producing a quantitative measure of a given mental state.

To give a hypothetical example of how this could be used, imagine a surgeon is presented with three different designs for a new teleoperated surgical robot. The goal is to evaluate which of these designs will be most efficient under working conditions. The surgeon performs the same operation using all three designs, while her brain activity is registered. Features of this brain activity are later compared to different recordings of her brain activity, collected during carefully controlled sessions of high and low workload. Based on the gathered data, it can be determined which of the three designs evoked the least workload-correlated brain activity. Furthermore, it may even be seen what phases of the operation saw local increases or decreases in such activity.

3.2.2 Examples from the Literature

Mental state monitoring can be particularly useful where more traditional methods have clear disadvantages. In *neuroergonomics*, "the study of brain and behaviour at work" (Parasuraman 2003), the use of BCI-based methods in order to assess an operator's workload levels may replace traditional ergonomics methods where the operator would need to be interrupted for data collection, for example, to fill out a workload-measuring questionnaire. Workload has been widely recognised as a fundamental issue in ergonomics (Wickens et al. 2014), but a clash of definitions and constructs and intersubjective differences prevent a uniform and objective measurement (Young et al. 2015). BCI methodology may here be able to deliver a data-driven approach that circumvents (or complements) conceptual definitions and focuses directly on the neurophysiological correlates of those conditions that are to be measured. In case of workload, a consistent finding reflects a frontal–parietal theta–alpha asymmetry in EEG activity, representing the interaction of the dorsolateral prefrontal cortex and the intraparietal sulcus, which are also described as anterior and posterior attentional systems in controlled attention tasks (Gerjets et al. 2014).

For example, Gevins and Smith (2003) used the Multi-Attribute Task Battery (Comstock & Arnegard 1992) to simulate controlled working conditions of three different load levels and found reliable differences in frontal theta and parietal alpha band power in the

continuous EEG recordings. On this basis, they constructed a cognitive workload index that could then be used to analyse later recordings. It performed in line with their hypothesis on other traditional experimental workload tasks as well as during more natural HCI tasks.

Using such pre-calibrated indices, mental state assessment can be used to, for example, A/B test and compare alternative user interface designs or task conditions with respect to the cognitive states they induce. Frey et al. (2014) review how a number of other constructs (attention/vigilance/fatigue, error recognition, emotions, engagement/flow/immersion) can be used to evaluate user experience during HCI tasks.

Care should be taken that the recorded reference data are ecologically valid and exhibit the neurophysiological features that specifically correlate with the mental state of interest and not with other aspects of the recording's context. For example, Mühl et al. (2014) recorded both high and low workload reference data in different affective contexts and showed that cross-context training improved the classifier's robustness. See also Section 3.4.2 for another approach to make the reference data robust to unrelated influences through a careful selection of reference tasks. Gerjets et al. (2014) and Brouwer et al. (2015) present an extended discussion of pitfalls and lessons learned in mental state assessment research.

Rather than recording reference data indicating extreme ends of the spectrum (e.g., underload vs. overload conditions), reference data can also be used as a baseline indicating optimal conditions (e.g., balanced workload). Zander and Jatzev (2012) used a measure of *deviation from the baseline* as undirected measure of suboptimal performance. Such a detected divergence can then be analysed more closely to find out the underlying factors. In their paradigm, a goal-oriented, gamified HCI task, the Kullback–Leibler divergence of recorded features, increased compared to a baseline measurement in phases where participants felt they had lost control over the interaction. This *loss of control* could perhaps constitute a mixture of workload (compensatory actions) and emotional states (frustration), contributing to the measured neurophysiological differences.

Aside from continuous, frequency-based measurements, features can also be extracted in the time domain of EEG recordings, looking at event-related potentials (ERPs).

Evoked responses to presented stimuli are said to reflect neuronal processing of those stimuli, and the ERPs' morphology may be modulated by affective and cognitive processes (Luck 2014). For example, an alternative workload measurement method is to pose an oddball paradigm as a secondary task (e.g., to count the number of relatively rare high-pitch tones among a monotonous sequence of low-pitch tones), but instruct participants to focus attention primarily on another, primary task. The more cognitive load is demanded by the primary task, the less attention can be invested in the secondary task. This decreased attention is then reflected in an amplitude decrease of the P300 component in the ERP following target stimuli in the oddball task (Kok 2001).

Frey et al. (2016) used both a continuous EEG-based measure of workload, calibrated using an experimentally controlled variant of the *n*-back task (Kirchner 1958), and a secondary oddball task to measure attention. They found that both measures differ significantly across difficulty levels in a 3D wayfinding game, albeit with different sensitivities.

ERP components may also reflect differences in higher-level cognitive processing. This assumption is the basis of ERP-based *guilty knowledge tests*, where the goal of the experiment is to find out whether or not the participant possesses any information that they wish to conceal. As such, guilty knowledge tests are a type of lie detector. Participants are shown a series of stimuli, some related to the concealed information (e.g., the used weapon in a crime), others neutral (other weapons, random objects). Should the participant exhibit deviating responses to the related stimuli compared to the neutral ones, then this could be taken as a clue that the participant does possess some information concerning the case in

question. Based on differences in their ERPs, Farwell and Donchin (1991) devised a system that could accurately determine for 87.5% of participants whether or not they had information concerning mock espionage scenarios that the participants were exposed to as part of the experiment. Ongoing developments in BCI and machine learning algorithms continue to be applied to improve performance in this field (e.g., Abootalebi et al. 2009).

Similar to workload, measures of a product or service's quality, attractiveness, or potential value to the user are traditionally gauged using questionnaires or interviews, and there is difficulty with respect to conceptual definitions and objective values. In order to elucidate consumer decisions, marketing researchers have turned to neuroscientific and BCI-based methodology to investigate human consumption from a neurophysiological perspective—*neuromarketing* (Lee et al. 2007). For example, Knutson et al. (2007) investigated neural correlates of two often opposing forces in purchase decisions: a product's attractiveness and its price. In a functional magnetic resonance imaging study, they found discriminative activity in separate cortical areas for these two product attributes. Based on features extracted from the nucleus accumbens and the mesial prefrontal cortex, they were subsequently able to develop an index predicting purchase behaviour as exhibited by participants during the experiment.

3.2.3 Reflection

The examples of mental state assessment given here all use an initial recording of brain activity as their basis and then apply signal processing and classification or other predictive techniques in order to obtain a continuous quantification reflecting mental states or changes in mental state. By using brain activity, these approaches exhibit a number of relevant differences compared to other methods that attempt to do the same thing (e.g., questionnaires, interviews, introspection, thinking aloud). Brain activity provides a continuous, potentially unobtrusive source of data, and its recording does not interfere with the state to be measured. It enables a functional, data-driven approach to individual, subject-specific state assessment. This can provide a possible alternative to sometimes competing and contradictory theory-driven approaches. It does, however, require special care to interpret the resulting data and validate the measures taken: in data-driven approaches, correlations may be found that are not neurophysiologically plausible—for example, co-varying neural responses to confounding variables may show up in cortical areas unrelated to the state of interest, or models may be overfitted on the available data (e.g., Babyak 2004; Haufe et al. 2014). Another reason to use neurophysiological measurements may be when the information of interest cannot be obtained otherwise, for example, when a continuous recording and quantification is required, or when the human's own indications may be biased.

Mental state monitoring is an essential element in (passive) BCI systems but does not constitute such a system in and of itself. As per the abovementioned definition of BCI systems, the third component—output—is missing: the measurement does not trigger any actions, nor does it influence any system, beyond the measurement itself. Rather, the measurement itself is the goal of the use of the technology, as an additional instrument, to be used for later post hoc analysis.

Section 3.3 will present examples of passive BCI systems that translate the measurements into system actions in real time. These actions simultaneously provide feedback to the participants, adding a first element of interactivity to the system. We will also see that this feedback has inherent advantages for the validation of the measured user state.

3.3 Open-Loop Adaptation

3.3.1 Introduction

HCI describes the interaction between one or more human users and a technical system (Preece et al. 2015). In a typical interaction cycle, an operator (user) gives a command to a machine, which processes the command and gives corresponding feedback to the operator. "Feedback" here denotes information about an action returned to the initial causing source of the action—i.e., a response from the computer back to the user. This can be either the effect of the action itself or a separate informative signal.

Even modern-day interaction techniques, e.g. touch screens and virtual reality controls, in essence still rely on the same principles as traditional techniques: in order to provide input to the computer, users explicitly activate one element after another, in accordance with the computer's logic (keyboard presses, mouse clicks, virtual buttons, menus, sliders, etc.). "Explicit" here reflects that the action was voluntarily executed with the conscious and sole intention of providing input to the computer.

Passive BCI can provide a fundamentally different, *implicit* input channel. "Implicit" here is the antonym to explicit and indicates that the input in question was generated without the source's voluntary intention of doing (merely) that (Rötting et al. 2009; Schmidt 2000; Zander et al. 2016). Passive BCI can be embedded into the HCI cycle by taking mental state assessments as implicit input, processing this input accordingly, and providing corresponding feedback to the user. This feedback creates interactivity with respect to the BCI: the computer now has the ability to respond to the user's implicit input.

In the simplest case, one interaction cycle (input–processing–feedback) can be viewed in isolation, following a simple one-time stimulus–response logic. Whenever a specific input is given to the machine, it performs one and the same action. The action serves only to comply with the input as it was given. This is referred to as open-loop adaptation.

Systems from this first category of adaptive applications apply mental state assessment to obtain a measure of a mental state online and respond to certain states with specific pre-programmed actions in an open-loop fashion.

For our surgeon, the online detection of high phases of workload could, for example, automatically trigger the system to switch on an indicator light. This communicates to the other members of the surgical team that the surgeon is not to be disturbed with lower-priority interruptions (Zander et al. 2017b).

3.3.2 Examples from the Literature

One mental state that has received particular attention in the context of general HCI is the *perception of errors*. Partially due to the artificial and limited nature of traditional interaction techniques (Suchman 1987; Tufte 1990), mistakes are made during ongoing HCI, resulting in frustration, loss of productivity, or, in safety-critical environments, potentially worse (Reason 1990). Upon perceiving the feedback indicating the error, or even already upon executing the erroneous action itself, the human user recognises that a mistake has been made. Such conscious error perception elicits a much-researched neuroelectric response known as the error-related negativity (ERN; Gehring et al. 2012). If a system could detect such a negativity in its human user, and link it to a specific action or feedback signal, it could automatically undo the apparently perceived to be erroneous action or even learn to prevent it in the future.

For example, Parra et al. (2003) report a 21.4% average reduction (albeit ranging from −6 to +49%) in errors made by the combined human–machine system on a speeded reaction task, which is the traditional experimental task to elicit response errors. Here, a passive BCI system calibrated to detect an ERN after each human response informed a decision to

override the human input. Note that, as the −6% figure for one participant indicates, this can lead to an increase in errors when the passive BCI itself makes mistakes.

Testing to what extent these findings would apply to more realistic HCI control scenarios, Iturrate et al. (2012) found similar detectable responses to errors in a series of three experiments that represent essentially a similar task but with different degrees of realism: abstract cursor control, simulated robot control, and actual robot control. Error-related responses have also been found in different HCI contexts and modalities, for example, based on auditory (Zander et al. 2011a) and tactile (Lehne et al. 2009) feedback.

Being reliable in various, realistic contexts, such responses can be used to undo or otherwise correct for errors during HCI, as, for example, Zander et al. (2010) demonstrated using passive BCI error correction support in a gaming context.

ERN detection has also been used by Kreilinger et al. (2012) to inform the continuous motion of a prosthetic arm. Blinking light-emitting diodes (LEDs) would inform the user whether the arm would continue or stop moving within the upcoming second. When a LED indicating continued movement elicited an error response in the user, this could be interpreted as a passive input command to stop moving, and vice versa.

The use of another type of cognitive state was proposed by Protzak et al. (2013) and further investigated by Shishkin et al. (2016) to support gaze-based HCI. In gaze-based HCI, an eye tracker is used to follow the user's gaze, which can thus be used to control, for example, a pointer on a screen. To select or activate an item, often a *dwell time* is used: if the cursor remains for a certain amount of time on one and the same item (i.e., if the user looks at it for that long), it is activated. This can lead to a high number of false activations, as not all gazes are intended for interaction. Both Protzak et al. (2013) and Shishkin et al. (2016) compared brain activity following gaze fixations on an item under two conditions: one where interaction was intended (gaze-based control was enabled), and one where it was not (gaze-based control was disabled). They found significant amplitude differences in the ERP following control versus no-control fixations. Shishkin et al. applied a classifier trained on these differences online, in a gaze-controlled selection task. Although their classification system achieved insufficient sensitivity values to fully replace the "click" of a mouse, this may one day be possible. A computer that can reliably detect its user's intentions in such a way could greatly reduce the amount of explicit commands required from the user.

BCI has also been used in gaming environments to effect changes in the ongoing interaction. For example, van de Laar et al. (2013) incorporated BCI into an online role-playing game. Depending on parietal alpha band power, assumed to reflect the player's state of relaxation versus tension, the character they controlled would shape-shift into either an elf (relaxation) or a bear (tension). These character forms had different functional abilities, thus affecting the user's strategy in the game. Note that, in the cited case, this was a form of *active* BCI, because the experimenters explicitly instructed the participants to control their state of mind in order to consciously shift between shapes. However, since the chosen features (parietal alpha) do naturally correlate to relaxation, this could also be an exemplary implementation of open-loop adaptation based on passive BCI. This highlights the issue discussed in Section 3.1: whether or not a system is active or passive does not depend on the implementation itself but ultimately on the user's behaviour. For similar examples, see Ilstedt Hjelm and Browall (2000), Ilstedt Hjelm (2003), and Zander and Krol (2017).

3.3.3 Reflection

Open-loop adaptive systems use a measurement of the user's state as implicit input in parallel to an ongoing (inter)action. These can be transient states such as error perception or interaction intent, which should then be carefully tied to the context that evoked them, or

more constant states such as moods or other psychophysical conditions, which will then also have a more constant effect on the interaction. As such, we see that the latter is used to implicitly control an application's mode, whereas the former is used to provide more timely implicit executive commands. The use of passive BCI methodology in these interactive systems either supports the user by replacing commands that otherwise needed to be provided explicitly or provides an extra dimension of experience to the interaction that would not have been possible using traditional input modalities.

An analysis of the effects of the BCI on the interaction can also serve as a validation of the system. For example, a quantification of *errors compensated* or *correct selections* provides a clear measure of the efficacy of the system. As mentioned in Section 3.2.3, though, care should still be taken that the implicit input underlying any performance improvement is neurophysiologically plausible.

The examples mentioned here use passive BCI for implicit human-to-computer communication. The implicit input is processed immediately, corresponding actions are performed, and the feedback given to the users completes the interaction cycle. This, however, is also the full extent of the interactivity at this point: the effects do not reach beyond the current interaction cycle, as there is no closed control loop influencing the next cycle. Both the executive and mode switching actions represent open-loop control, where single, independent state detections result in single, fixed actions from the computer. More interactive applications exhibit closed-loop control, where the given feedback also influences the user, as we shall see next.

3.4 Closed-Loop Adaptation

3.4.1 Introduction

In closed-loop systems, output is fed back to the system as new input. In other words, the generated output influences the next cycle. In case of an elementary human–machine system based on passive BCI, the ultimate source of input is the human user's mental state. One method to implement a closed loop is thus to have the produced feedback cause a change in the user state. By influencing the mental state that is measured, the system influences its own next input.

One motivation for such an implementation may be to promote and sustain a desirable psychological state; for example, an optimally supporting balance of factors, or a state of *flow* (Csikszentmihalyi 2008). The adaptive loop may, for example, be programmed to promote a state of positive engagement in order to improve performance. This type of system could modify the demands of the task to maximise the user's motivation and enjoyment. When there is a clearly defined and unambiguously operationalised psychological state that is desirable, then the system can use its adaptive logic and feedback to manipulate the state of the user until the desired target state is reached.

This increases the interactivity of the system in the sense that it is now able to respond purposefully, as well as able to monitor the effect of each response. The given response attempts to qualitatively influence the condition of the ongoing interaction.

Thus, closed-loop adaptive systems apply mental state assessment to obtain a measure of a mental state online and respond to certain states—or changes in states—with actions that influence that same mental state.

We can again turn to our earlier example of workload in the operating theatre to illustrate this category of applications. When a system detects that the surgeon is currently under high load, it can attempt to directly compensate for this by assuming certain tasks on behalf of the surgeon. By relieving the surgeon from some of her tasks, the system

purposefully attempts to lower her load. We will see a similar example from the literature in more detail, described next.

3.4.2 Examples from the Literature

As mentioned in Section 3.2.2, a measure of workload is important in ergonomics research in order to inform design decisions. An adaptive system, however, could contain a range of options that would otherwise be decided upon by the designers and switch between them during online operation depending on current measures of workload. One such option range is the level of automation: a machine can be operated entirely by explicit commands from the operator, or support the human operator to various degrees, up to assuming full autonomy and leaving the user in a mere supervisory role (see Parasuraman et al. 2000 for an overview and framework describing these options). The purpose of automation is to reduce operator mental workload and alleviate fatigue during sustained performance; however, use of automation is associated with negative consequences such as the out-of-the-loop problems and decay of skilled performance (Endsley & Kiris 1995). Striking the right balance is important to the well-being and performance of the operator, but the appropriate level of automation varies with fluctuating operator capacities and varying task demands. Therefore, *adaptive automation* (Byrne & Parasuraman 1996) attempts to match the level of allocation to the current capacity of the operator, measured in real time.

Kohlmorgen et al. (2007), for example, performed EEG measurements during highway driving. Aside from controlling the vehicle, two additional tasks were given to the participants that mimic human–car interaction (an auditory response task) as well contextual distractors during driving (a mental calculation task and a listening comprehension task). All these together defined high-workload phases, which were compared to low-workload phases, where no contextual distractors were present. A classifier calibrated to distinguish these phases was later applied in an online driving condition such that high workload detection resulted in the automatic temporary suspension of the auditory response task. The mean performance on this response task significantly improved due to this intervention.

Similar adaptive logic can be applied to intelligent tutoring systems designed to deliver educational material in a way that sustains the engagement of the student. For example, the pacing of the learning process can be adjusted dynamically in order to sustain motivation without inducing fatigue, frustration, or information overload. In particular, working memory load is said to need careful management for an optimal learning experience (Cowan 2014).

Gerjets et al. (2014) review theory and practice with respect to online working memory assessment and present two studies in the context of adaptive learning environments. Special care was taken to select experimentally controllable calibration tasks. A combination of two tasks was used, as these tasks both targeted the same cognitive resources that were to be measured during realistic learning exercises, while at the same time these tasks differed with respect to their demand on executive resources. As such, Gerjets et al. attempted to develop a classifier that could detect working memory load independent of other task properties. Using this classifier, a classification accuracy of 73% was achieved on single trials of word algebra tasks. Their aim is to apply this classifier in an online adaptive environment, catering to current cognitive abilities of the student.

Such an online application has been presented by Yuksel et al. (2016): they used functional near-infrared spectroscopy (fNIRS) first to differentiate between participants' brain activity playing easy and hard piano pieces and then to detect their cognitive workload online during the learning of a new piece. The adaptive tutor first presented only one line of musical notes and added the next line only when workload levels indicated that the previous

step had been learned to a sufficient degree. Yuksel et al. report significant increases in performance and learning speed using the fNIRS-based adaptive system as compared to self-paced learning.

Ewing et al. (2016) applied the concept of closed-loop adaptation to a game, Tetris, adjusting the game's speed in order to maintain a level of optimal engagement. They implemented different adaptive behaviours (i.e., more or less responsive), but found little difference between them. They discuss the issue of finding an appropriate benchmark: taking a manual condition as control, as they did, may provide an inherent difference in user experience that is not related to the exact behaviour of the closed-loop adaptive system. A better control may be to compare this adaptive system to alternative methods of adaptation, for example, pre-programmed, random, or "yoked" (responding to the user state of another individual). Similar to the earlier distinction between mode switches and executive commands, the examples given above aim to achieve a constant, optimal state in the user, but adaptive systems can also effectuate short-lived closed feedback loops in order to achieve a one-time cognitive state.

Kirchner et al. (2013) suggest closed-loop implicit control of an alarm system, making use of BCI methodology to assess whether or not the alarm has been perceived and/or the user intends to act on it. Based on this information, the system could then decide to re-sound the alarm with increased saliency until it has been perceived. A state of nonperception would thus trigger new feedback until a state of perception would be observed. Kirchner et al. report accurate classification rates of perceived versus missed alarms during a demanding robotic teleoperation scenario. In another study, they used these as measures of current load, reversing this logic: nonperception of new instructions indicated that the user was occupied with previous instructions. The new instructions were then delayed to avoid distractions and promote adequate levels of engagement and performance (Kirchner et al. 2016).

3.4.3 Reflection

Direct access to a quantification of a certain mental state enables systems to monitor and, with appropriate feedback, influence this state also in a closed loop. This approach can be used to promote and sustain desirable mental states or indeed to mitigate an undesired state. The goal of this closed-loop control can be to achieve a certain equilibrium with respect to the state of interest or to push for a given threshold. In the former case, the system's potential influence may need to be bidirectional, allowing for both changes that encourage and changes that alleviate the target state.

When considering what is a closed-loop system and what is not, special attention must be given to the exact state that is targeted, which may be related to but different from the state that is measured. For example, error perception, as a transient state in the form of an error-related potential that lasts for less than a second, cannot necessarily be used as such to define a closed-loop system. This would be different, however, if this transient mental state can be interpreted in the given context as, for example, being indicative of a more persistent state of general dissatisfaction.

A closed feedback loop represents increased interactivity because the feedback in this case does not merely function as an informative response to the user, but in fact *purposefully acts upon* the user. The feedback intends to influence the user, as the user's commands intend to influence the system. Where in our earlier examples we saw implicit human-to-computer communication, now we could speak of an *implicit dialogue* between both agents, cooperating more closely in order to achieve common goals.

This cooperation as exemplified here is still limited by the amount of information available to the system and the simplistic logic of the single loop—for example, the one-

dimensional measure of workload can only have a one-dimensional adaptation of automation levels. A next step up would need to tackle the question how systems can achieve a more complex adaptive, cooperative behaviour based on the limited information passive BCI methodology can obtain at any single time. One such method will be discussed next.

3.5 Automated Adaptation

3.5.1 Introduction

Analysis of ongoing brain activity can give insight into momentary mental states, but by itself, not into their likely causes or the appropriate responses. In the above-mentioned adaptive systems, the context is controlled and specific enough to reasonably assume that, for example, workload is caused by task demands, and increased automation will lead to a reduction in measured load. The interactive functions, however, are limited to these specific conditions.

Human–human interaction, the benchmark for natural communication and cooperation, relies heavily on a shared understanding of the world, the situation, and the communication partner—a shared model of appropriate behaviour and how things are (Fischer 2001). In this, however, HCI is heavily asymmetrical: the human user can potentially be appraised of all details concerning the state of the computer, but hardly any such information is available vice versa (Suchman 1987). For a computer to *understand* its user, it needs information and input that goes beyond the bare necessities for its operation. The applications discussed in this chapter so far provide such information, but use it for pre-programmed one-dimensional response logic—they do not "understand" the user, that is, they build no model of higher, more abstract aspects of the user's cognitive or affective behaviour. A better-informed system has increased ability to respond appropriately to the user; even to act on the given, known information before any new input is received. This thus constitutes an additional increase in interactivity: now also the system, and not merely the human user, can autonomously decide to act.

In this final category of automated adaptive applications, the systems apply mental state assessment alongside other methods of information gathering and build a model to represent aspects of the user's cognitive or affective responses. It is this automatically generated model, finally, that serves as a basis for the system's own autonomous behaviour and adaptations—not (merely) the system's current input.

Our surgeon's load may have been consistently detected to increase beyond sustainable levels after 1 hour when using a specific teleoperated robot. Having learned this, a well-informed system could call for an additional assistant ahead of time.

3.5.2 Examples from the Literature

Zander et al. (2014) and Zander et al. (2016) used single-trial detections of error-related brain activity not to immediately correct perceived mistakes, but to learn, over time, what strategy the user followed to subjectively make the distinction between errors and non-errors. To that end, the system exhibited different behaviours and learned which behaviours evoked positive and which evoked negative responses. These responses were generated in the medial prefrontal cortex, associated with a fundamental source of human intelligence: predictive coding (Hawkins & Blakeslee 2007). It thus built a model of the user's higher-level preferences with respect to the presented behaviours. Already during this learning process, the system adapted to exhibit behaviour most likely to be perceived positively.

Specifically, this approach was applied to two-dimensional cursor control. A cursor

on a computer screen initially moved randomly. Movements in some directions led to positive responses, while movements in other directions evoked negative responses. Over time, the system learned the pattern behind these responses and, as such, learned where the user wanted the cursor to go. The system could thus steer the cursor toward the intended target.

Extending their earlier-mentioned experiments, Iturrate et al. (2015) demonstrated a similar principle using a robotic arm. Participants observed its movements while their error responses were tracked in order to inform the arm of appropriate and inappropriate behaviour. Based on the participant's automatic responses, the system learned using a process of reinforcement and elimination which path or control policy was currently preferred by the participant. To that end, all possible control policies were known to the system upon initialisation, and it then attempted to find a best match between these policies and the participant's responses. The use of policies as opposed to directions allows the system to account for the arm's changing relative position to the target, and would potentially be able to deal with obstacles or other more complex arm behaviour.

This implementation thus differs from the one used by Zander et al. (2014) and Zander et al. (2016) with respect to the form of, and the kind of information contained within, the model, but uses the same general principles of automated adaptation.

3.5.3 Reflection

In the two examples given in this section, rather than either having direct consequences or not, mental state measurements were stored along with a description of the contexts (here: computer actions) during which they were observed. As such, a model was created that correlated user states with different contexts. This model then served as a basis for the system to decide future actions.

These examples were again limited to their respective environments (cursor/arm movements in a limited number of directions), but the principle can be envisioned to apply in a broader sense. With a reliable measure of e.g. satisfaction versus dissatisfaction (as, e.g., in Zander et al. 2016), as well as an accurate recording of relevant context variables (Zander & Jatzev 2012), a hypothetical ubiquitous passive BCI system could generate such a model for a large number of aspects in a user's daily and working life. By collating contextual factors with measures of the user's mental state, the computer learns to understand this user in much the same way as we humans learn from our own experiences. Subsequently, given a known context, the ubiquitous computer could execute that action that has been observed to be most likely to satisfy the user in that context.

Although the ultimate goal of this hypothetical system would be the same as of the earlier-mentioned adaptive automation, i.e. to promote and sustain a desirable state, the methods by which this can be achieved are not set: the parameters of adaptation are learned automatically by the system itself, on the basis of a continuously updated model of its user.

For a large part, it is the understanding (i.e., the models) that we have of our social and physical context that allows us humans to behave appropriately in various contexts. We have started to build these models as soon as we were born and continue to update them as long as we are capable. For this, we use all information available to us, not just the information that others have explicitly deemed relevant to us. By giving computers access to more information and giving them the freedom to build their own models of their users and the context, a form of artificial intelligence may be created that may finally be seen as an autonomous, interactive agent, just as we see ourselves.

In a far-reaching, hypothetical case of a learning, closed-loop, automatically adapting system, we could imagine a system providing what we want before we can act upon those

desires ourselves. With this autonomously operating system originally having learned from our own subjective responses to various actions and contexts, it may at some point become unclear who was the ultimate initiator of the action, by whose authority the action was executed, or who, in fact, "acted." As the lines separating the two interactive agents, human and machine, become blurred, merely the interaction itself remains.

Given also ethical deliberations with respect to, for example, data ownership, informed consent, the potentially harmful influence of machine over man (the same technique could be used to promote an undesired state, resulting in the computer actively working *against* the user), and outcome responsibility, it is prudent to aim to strike a balance between the two agents' influence on each other. Users should always have access to full information concerning the system's models, processes, goals, and actions. Users should also have ultimate control over, and ownership of, data that are derived from their physiology (Fairclough 2014). From a cooperative perspective, the best balance may be one where both agents get to optimally use their strengths while their weaknesses are mitigated by the strengths of the other, while keeping these domains formally separate.

3.6 Summary and Conclusion

In this chapter we have discussed potential applications that build up on the ability to perform real-time analyses of neurophysiological data, first demonstrated as brain–computer interfacing by Jacques J. Vidal. In 1973, envisioning future progress, Vidal wrote that "mental decisions and reactions can be probed, in a dimension that both transcends and complements overt behaviour" (Vidal 1973). This statement appears to accurately capture the motivation of passive BCI research today, in its various forms of applications. In this chapter, these were divided into four distinct categories, based on their level of interactivity.

One category of applications, mental state assessment, uses the quantification of state-specific brain activity as a measurement, for example, to replace or support other measurements (e.g., questionnaires, behavioural data). This approach can provide data unobtrusively, can be individually calibrated, and may be able to access information that would otherwise remain hidden. An example from this category is the generation of a workload index based on brain activity, in order to assess the amount of resources that different instruments demand from the surgeon operating them.

A second category of applications, open-loop adaptation, uses the obtained measurements to inform an open-loop adaptation of the software. Using online state assessment, these systems can respond in real time to the detection of certain states. Continuing the above example, a warning light could switch on in an operating theatre when it is detected that the surgeon is currently experiencing high load.

In a further category of applications, closed-loop adaptation, the software adaptation is designed with the explicit purpose to also effect a user adaptation, creating a closed control loop. This provides the system with the ability to not merely detect but also act upon the detected mental states. To keep the surgeon's workload at a sustainable level, for example, certain systems and operations in the theatre may be adaptively automated.

Finally, a category of applications was described where the control logic of the system exceeds predetermined single-loop control: automated adaptation. The system is given the autonomy to learn and adapt based on a larger amount of information, coupled with a measure of the user's mental state. This also grants the system the ability to act autonomously for, or on behalf of, the user. For the surgeon, an additional assistant may automatically be called as the system predicts workload will exceed otherwise manageable levels.

The first category is neither BCI nor interactive but serves as a basis for the three

following categories. Their interactivity shows in their different abilities to respond to the given input. Since effective passive BCI systems are unobtrusive and inconspicuous, these systems' responses may be similarly hidden. It is for this reason that we refer to *adaptation*, using this as a more generic term that includes traditional feedback (e.g., the movement of a prosthetic limb) but also other ways to process the given input (e.g., updating a user model) or act on the basis of previously learned information. When adaptation is ultimately based on neurophysiological measures, we refer to these systems collectively as *neuroadaptive* systems (Zander et al. 2016).

The different adaptive behaviours are not mutually exclusive. An automated adaptation system might, for example, predict what a user intends to do and execute that action in advance, but then, in an open-loop fashion, undo the action when it is detected that the prediction was in error.

The here-presented categorisation emphasises the observable path of passive BCI applications toward increased interactivity.

The investigation of truly interactive systems, in particular based on natural human brain activity, can only be properly researched and developed in natural, real-world interactive settings. The well-controlled experimental conditions will at some point no longer be sufficient to make significant progress. Luckily, easily applicable, comfortable electrodes with sufficient signal quality and stability are being developed and continue to be improved (e.g., Goverdovsky et al. 2016; Mullen et al. 2015; Zander et al. 2011b, 2017a). Increased interest from high-profile commercial ventures is currently providing an effective development infrastructure and budget, at least for specific business interests—but the field as a whole may benefit, since the underlying issues are the same. Finally, when classifiers can be constructed that are independent of the human users and applicable across contexts, their universal usability will provide easy access to mental state assessment, boosting its usage (e.g., Krol et al. 2016; Wei et al. 2015; Zander 2012). When these development reach maturity, real-world implementations of the adaptive systems discussed here may seriously and dramatically alter the many human–system interactions that dominate our everyday lives (McDowell et al. 2013).

This chapter described a trend toward increased interactivity in past and current passive BCI research and development. The history of HCI as a whole, too, can be seen in light of increased interactivity. In early computer systems, input commands were written beforehand and given to the system as complete programs. This either resulted in the successful processing of the program, or, quite bluntly, it did not. It was only later that humans were given real-time control over the system, through, for example, immediate processing and interrupt options (Suchman 1987). This provided the first interactive human–computer experience, albeit with ultimate agency solely with the human user. A next step toward interactivity thus appeared to be to give agency also to the computer—allowing it to adapt itself and execute commands that the user did not explicitly ask for. Passive BCI systems, discussed here in their various forms, can provide a source of information to make this computer agency in line with the user's intentions. It can provide this next step in human–computer interactivity, and, as the categorization used here intends to show, may provide the ones after that as well, in the form of advanced neuroadaptive technology.

Acknowledgements

References

Abootalebi, V., Moradi, M. H., & Khalilzadeh, M. A. 2009. A new approach for EEG feature extraction in P300-based lie detection. *Computer Methods and Programs in Biomedicine*, *94*(1), 48–57.

Babyak, M. A. 2004. What you see may not be what you get: A brief, nontechnical introduction to overfitting in regression-type models. *Psychosomatic Medicine*, *66*(3), 411–421.

Birbaumer, N., Ghanayim, N., Hinterberger, T., Iversen, I., Kotchoubey, B., Kübler, A., ... Flor, H. 1999. A spelling device for the paralysed. *Nature*, *398*(6725), 297–298.

Blankertz, B., Curio, G., & Müller, K.-R. 2002a. Classifying single trial EEG: Towards brain computer interfacing. In T. G. Dietterich, S. Becker, & Z. Ghahramani (Eds.), *Advances in Neural Information Processing Systems 14* (pp. 157–164). Cambridge, MA: MIT Press.

Blankertz, B., Schäfer, C., Dornhege, G., & Curio, G. 2002b. Single trial detection of EEG error potentials: A tool for increasing BCI transmission rates. In J. R. Dorronsoro (Ed.), *Artificial Neural Networks – ICANN 2002* (Vol. 2415, pp. 1137–1143). Berlin, Germany: Springer.

Brouwer, A.-M., & van Erp, J. 2010. A tactile P300 brain–computer interface. *Frontiers in Neuroscience, 4*, 19.

Brouwer, A.-M., Zander, T. O., van Erp, J. B. F., Korteling, J. E., & Bronkhorst, A. W. 2015. Using neurophysiological signals that reflect cognitive or affective state: Six recommendations to avoid common pitfalls. *Frontiers in Neuroscience, 9*, 136.

Byrne, E. A., & Parasuraman, R. 1996. Psychophysiology and adaptive automation. *Biological Psychology*, *42*(3), 249–268.

Comstock, J. R. J., & Arnegard, R. J. 1992. *The Multi-Attribute Task Battery for Human Operator Workload and Strategic Behavior Research* (Tech. Rep.). Hampton, VA, United States: NASA Langley Research Center.

Cowan, N. 2014. Working memory underpins cognitive development, learning, and education. *Educational Psychology Review*, *26*(2), 197–223.

Csikszentmihalyi, M. 2008. *Flow: The Psychology of Optimal Experience* (Modern Classics ed.). New York, NY, USA: HarperCollins.

Cutrell, E., & Tan, D. 2008. BCI for passive input in HCI. In *Proceedings of the Brain–Computer Interfaces for HCI and Games Workshop at the SIGCHI Conference on Human Factors in Computing Systems (CHI).*

Endsley, M. R., & Kiris, E. O. 1995. The out-of-the-loop performance problem and level of control in automation. *Human Factors, 37*(2), 381–394.

Ewing, K. C., Fairclough, S. H., & Gilleade, K. 2016. Evaluation of an adaptive game that uses EEG measures validated during the design process as inputs to a biocybernetic loop. *Frontiers in Human Neuroscience, 10*, 223.

Fairclough, S. H. 2014. Physiological data must remain confidential. *Nature*, *505*, 263.

Farwell, L. A., & Donchin, E. 1988. Talking off the top of your head: Toward a mental prosthesis utilizing event-related brain potentials. *Electroencephalography and Clinical Neurophysiology*, *70*(6), 510–523.

Farwell, L. A., & Donchin, E. 1991. The truth will out: Interrogative polygraphy ("lie detection") with event-related brain potentials. *Psychophysiology*, *28*(5), 531–547.

Ferrez, P. W., & Millán, J. d. R. 2005. You are wrong!—Automatic detection of interaction errors from brain waves. In *Proceedings of the 19th International Joint Conference on Artificial Intelligence* (pp. 1413–1418). San Francisco, CA: Morgan Kaufmann.

Ferrez, P. W., & Millán, J. d. R. 2008. Error-related EEG potentials generated during simulated brain–computer interaction. *IEEE Transactions on Biomedical Engineering*, *55*(3), 923–929.

Fischer, G. 2001. User modeling in human–computer interaction. *User Modeling and User-Adapted Interaction*, *11*(1–2), 65–86.

Frey, J., Daniel, M., Castet, J., Hachet, M., & Lotte, F. 2016. Framework for electroencephalography-based evaluation of user experience. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (pp. 2283–2294). New York, NY, USA: ACM.

Frey, J., Mühl, C., Lotte, F., & Hachet, M. 2014. Review of the use of electroencephalography as an evaluation method for human–computer interaction. In *Proceedings of the International Conference on Physiological Computing Systems* (pp. 214–223).

Furdea, A., Halder, S., Krusienski, D., Bross, D., Nijboer, F., Birbaumer, N., & Kübler, A. 2009. An auditory oddball (P300) spelling system for brain–computer interfaces. *Psychophysiology*, *46*(3), 617–625.

Gehring, W. J., Liu, Y., Orr, J. M., & Carp, J. 2012. The error-related negativity (ERN/Ne). In S. J. Luck & E. S. Kappenman (Eds.), *Oxford Handbook of Event-Related Potential Components* (pp. 231–291). New York, NY: Oxford University Press.

Gerjets, P., Walter, C., Rosenstiel, W., Bogdan, M., & Zander, T. O. 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. *Frontiers in Neuroscience*, *8*, 385.

Gevins, A., & Smith, M. E. 2003. Neurophysiological measures of cognitive workload during human–computer interaction. *Theoretical Issues in Ergonomics Science*, *4*(1–2), 113–131.

Goverdovsky, V., Looney, D., Kidmose, P., & Mandic, D. P. 2016. In-ear EEG from viscoelastic generic earpieces: Robust and unobtrusive 24/7 monitoring. *IEEE Sensors Journal, 16*(1), 271–277.

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.

Hawkins, J., & Blakeslee, S. 2007. *On Intelligence*. London, UK: Macmillan.

Ilstedt Hjelm, S. 2003. Research + design: The making of Brainball. *Interactions, 10*(1), 26–34.

Ilstedt Hjelm, S., & Browall, C. 2000. Brainball—Using brain activity for cool competition. In *Proceedings of the 1st Nordic Conference of Human–Computer Interaction (NordiCHI)*.

Iturrate, I., Chavarriaga, R., Montesano, L., Minguez, J., & Millán, J. d. R. 2012. Latency correction of error potentials between different experiments reduces calibration time for single-trial classification. In *2012 Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)* (pp. 3288–3291).

Iturrate, I., Chavarriaga, R., Montesano, L., Minguez, J., & Millán, J. d. R. 2015. Teaching brain–machine interfaces as an alternative paradigm to neuroprosthetics control. *Scientific Reports*, *5*, 13893.

Kirchner, E. A., Kim, S. K., Straube, S., Seeland, A., Whrle, H., Krell, M. M., ... Fahle, M. 2013. On the applicability of brain reading for predictive human–machine interfaces in robotics. *PLOS ONE*, *8*(12), 1–19.

Kirchner, E. A., Kim, S. K., Tabie, M., Wöhrle, H., Maurus, M., & Kirchner, F. 2016. An

intelligent man-machine interface—Multi-robot control adapted for task engagement based on single-trial detectability of P300. *Frontiers in Human Neuroscience*, *10*, 291.

Kirchner, W. L. 1958. Age differences in short-term retention of rapidly changing information. *Journal of Experimental Psychology, 55*(4), 352–358.

Knutson, B., Rick, S., Wimmer, G. E., Prelec, D., & Loewenstein, G. 2007. Neural predictors of purchases. *Neuron*, *53*(1), 147–156.

Kohlmorgen, J., Dornhege, G., Braun, M., Blankertz, B., Curio, G., Hagemann, K., ... Kincses, W. 2007. Improving human performance in a real operating environment through real-time mental workload detection. In G. Dornhege (Ed.), *Toward Brain–Computer Interfacing* (pp. 409–422). Cambridge, MA: MIT Press.

Kok, A. 2001. On the utility of p3 amplitude as a measure of processing capacity. *Psychophysiology*, *38*(3), 557–577.

Kreilinger, A., Neuper, C., & Müller-Putz, G. R. 2012. Error potential detection during continuous movement of an artificial arm controlled by brain–computer interface. *Medical & Biological Engineering & Computing*, *50*(3), 223–230.

Krol, L. R., Freytag, S.-C., Fleck, M., Gramann, K., & Zander, T. O. 2016. A task-independent workload classifier for neuroadaptive technology: Preliminary data. In *2016 IEEE International Conference on Systems, Man, and Cybernetics (SMC)* (pp. 003171–003174).

Lee, N., Broderick, A. J., & Chamberlain, L. 2007. What is 'neuromarketing'? A discussion and agenda for future research. *International Journal of Psychophysiology*, *63*(2), 199–204.

Lehne, M., Ihme, K., Brouwer, A. M., van Erp, J. B. F., & Zander, T. O. 2009. Error-related EEG patterns during tactile human–machine interaction. In *2009 3rd International Conference on Affective Computing and Intelligent Interaction and Workshops* (pp. 1–9).

Lotte, F., Congedo, M., Lécuyer, A., Lamarche, F., & Arnaldi, B. 2007. A review of classification algorithms for EEG-based brain–computer interfaces. *Journal of Neural Engineering*, *4*(2), 24.

Luck, S. J. 2014. *An Introduction to the Event-Related Potential Technique* (2nd ed.). Cambridge, MA: MIT Press.

McDowell, K., Lin, C.-T., Oie, K. S., Jung, T.-P., Gordon, S., Whitaker, K. W., ... Hairston, W. D. 2013. Real-world neuroimaging technologies. *IEEE Access, 1*, 131–149.

Mühl, C., Jeunet, C., & Lotte, F. 2014. EEG-based workload estimation across affective contexts. In A.-M. Brouwer, T. O. Zander, & J. B. F. van Erp (Eds.), *Using Neurophysiological Signals That Reflect Cognitive or Affective State* (pp. 227–241). Frontiers Media SA.

Mullen, T. R., Kothe, C. A. E., Chi, Y. M., Ojeda, A., Kerth, T., Makeig, S., ... Cauwenberghs, G. 2015. Real-time neuroimaging and cognitive monitoring using wearable dry EEG. *IEEE Transactions on Biomedical Engineering, 62*(11), 2553–2567.

Müller, K.-R., Tangermann, M., Dornhege, G., Krauledat, M., Curio, G., & Blankertz, B. 2008. Machine learning for real-time single-trial EEG-analysis: From brain–computer interfacing to mental state monitoring. *Journal of Neuroscience Methods*, *167*(1), 82–90.

Müller-Putz, G. R., & Pfurtscheller, G. 2008. Control of an electrical prosthesis with an SSVEP-based BCI. *IEEE Transactions on Biomedical Engineering*, *55*(1), 361–364.

Parasuraman, R. 2003. Neuroergonomics: Research and practice. *Theoretical Issues in Ergonomics Science*, *4*(1–2), 5–20.

Parasuraman, R., Sheridan, T. B., & Wickens, C. D. 2000. A model for types and levels of

human interaction with automation. *IEEE Transactions on Systems, Man, and Cybernetics—Part A: Systems and Humans*, 30(3), 286–297.

Parra, L. C., Spence, C. D., Gerson, A. D., & Sajda, P. 2003. Response error correction—A demonstration of improved human–machine performance using real-time EEG monitoring. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 11(2), 173–177.

Pfurtscheller, G., & Neuper, C. 2001. Motor imagery and direct brain–computer communication. *Proceedings of the IEEE*, 89(7), 1123–1134.

Preece, J., Sharp, H., & Rogers, Y. 2015. *Interaction Design: Beyond Human–Computer Interaction* (4th ed.). Chichester, UK: Wiley.

Protzak, J., Ihme, K., & Zander, T. O. 2013. A passive brain–computer interface for supporting gaze-based human–machine interaction. In C. Stephanidis & M. Antona (Eds.), *Universal Access in Human–Computer Interaction. Design Methods, Tools, and Interaction Techniques for eInclusion* (Vol. 8009, pp. 662–671). Berlin Heidelberg, Germany: Springer.

Ramoser, H., Müller-Gerking, J., & Pfurtscheller, G. 2000. Optimal spatial filtering of single trial EEG during imagined hand movement. *IEEE Transactions on Rehabilitation Engineering*, 8(4), 441–446.

Reason, J. 1990. *Human Error*. Cambridge, UK: Cambridge University Press.

Rötting, M., Zander, T. O., Trösterer, S., & Dzaack, J. 2009. Implicit interaction in multimodal human–machine systems. In C. M. Schlick (Ed.), *Industrial Engineering and Ergonomics* (pp. 523–536). Berlin Heidelberg, Germany: Springer.

Schmidt, A. 2000. Implicit human computer interaction through context. *Personal and Ubiquitous Computing*, 4(2-3), 191–199.

Shishkin, S. L., Nuzhdin, Y. O., Svirin, E. P., Trofimov, A. G., Fedorova, A. A., Kozyrskiy, B. L., & Velichkovsky, B. M. 2016. EEG negativity in fixations used for gaze-based control: Toward converting intentions into actions with an eye-brain–computer interface. *Frontiers in Neuroscience*, 10, 528.

Soekadar, S. R., Witkowski, M., Gómez, C., Opisso, E., Medina, J., Cortese, M., ... Vitiello, N. 2016. Hybrid EEG/EOG-based brain/neural hand exoskeleton restores fully independent daily living activities after quadriplegia. *Science Robotics*, 1(1).

Suchman, L. A. 1987. *Plans and Situated Actions: The Problem of Human–Machine Communication*. Cambridge, UK: Cambridge University Press.

Treder, M. S., Schmidt, N. M., & Blankertz, B. 2011. Gaze-independent brain–computer interfaces based on covert attention and feature attention. *Journal of Neural Engineering*, 8(6), 066003.

Tufte, E. R. 1990. *Envisioning Information*. Cheshire, CT: Graphics Press.

van de Laar, B., Gürkök, H., Bos, D. P.-O., Poel, M., & Nijholt, A. 2013. Experiencing BCI control in a popular computer game. *IEEE Transactions on Computational Intelligence and AI in Games*, 5(2), 176–184.

van Erp, J. B., Lotte, F., & Tangermann, M. 2012. Brain–computer interfaces: Beyond medical applications. *Computer*, 45(4), 26-34.

Vidal, J. J. 1973. Toward direct brain–computer communication. *Annual Review of Biophysics and Bioengineering*, 2(1), 157–180.

Vidal, J. J. 1977. Real-time detection of brain events in EEG. *Proceedings of the IEEE*, 65(5), 633–641.

Wei, C.-S., Lin, Y.-P., Wang, Y.-T., Jung, T.-P., Bigdely-Shamlo, N., & Lin, C.-T. 2015. Selective transfer learning for EEG-based drowsiness detection. In *2015 IEEE International Conference on Systems, Man, and Cybernetics* (pp. 3229–3232).

Wessberg, J., Stambaugh, C. R., Kralik, J. D., Beck, P. D., Laubach, M., Chapin, J. K., ... Nicolelis, M. A. L. 2000. Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. *Nature*, *408*, 361–365.

Wickens, C. D., Hollands, J. G., Banbury, S., & Parasuraman, R. 2014. *Engineering Psychology and Human Performance* (4th ed.). New York, NY: Routledge.

Wolpaw, J. R., Birbaumer, N., Heetderks, W. J., McFarland, D. J., Peckham, P. H., Schalk, G., ... Vaughan, T. M. 2000. Brain–computer interface technology: A review of the First International Meeting. *IEEE Transactions on Rehabilitation Engineering*, *8*(2), 164–173.

Wolpaw, J. R., Birbaumer, N., McFarland, D. J., Pfurtscheller, G., & Vaughan, T. M. 2002. Brain–computer interfaces for communication and control. *Clinical Neurophysiology*, *113*(6), 767–791.

Wolpaw, J. R., & McFarland, D. J. 2008. Brain–computer interface operation of robotic and prosthetic devices. *Computer*, *41*, 52–56.

Wolpaw, J. R., & Wolpaw, E. W. 2012. Brain–computer interfaces: Something new under the sun. In *Brain–Computer Interfaces: Principles and Practice* (pp. 3–12). Oxford, UK: Oxford University Press.

Young, M. S., Brookhuis, K. A., Wickens, C. D., & Hancock, P. A. 2015. State of science: Mental workload in ergonomics. *Ergonomics*, *58*(1), 1–17.

Yuksel, B. F., Oleson, K. B., Harrison, L., Peck, E. M., Afergan, D., Chang, R., & Jacob, R. J. K. 2016. Learn piano with BAch: An adaptive learning interface that adjusts task difficulty based on brain state. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (pp. 5372–5384). New York, NY, USA: ACM.

Zander, T. O. 2012. *Utilizing Brain–Computer Interfaces for Human–Machine Systems* (Doctoral dissertation, Technische Universität Berlin, Berlin, Germany).

Zander, T. O., Andreessen, L. M., Berg, A., Bleuel, M., Pawlitzki, J., Zawallich, L., Krol, L. R., & Gramann, K. 2017a. Evaluation of a dry EEG system for application of passive brain–computer interfaces in autonomous driving. *Frontiers in Human Neuroscience, 11*, 78.

Zander, T. O., Brönstrup, J., Lorenz, R., & Krol, L. R. 2014. Towards BCI-based implicit control in human–computer interaction. In S. H. Fairclough & K. Gilleade (Eds.), *Advances in Physiological Computing* (pp. 67–90). Berlin, Germany: Springer.

Zander, T. O., & Jatzev, S. 2012. Context-aware brain–computer interfaces: Exploring the information space of user, technical system and environment. *Journal of Neural Engineering*, *9*(1), 016003.

Zander, T. O., Klippel, M. D., & Scherer, R. 2011a. Towards multimodal error responses: A passive BCI for the detection of auditory errors. In *Proceedings of the 13th International Conference on Multimodal Interfaces* (pp. 53–56). New York, NY, USA: ACM.

Zander, T. O., & Kothe, C. A. 2011. Towards passive brain–computer interfaces: Applying brain–computer interface technology to human–machine systems in general. *Journal of Neural Engineering*, *8*(2), 025005.

Zander, T. O., Kothe, C. A., Jatzev, S., Dashuber, R., Welke, S., De Filippis, M., & Rötting, M. 2008a. Team PhyPA: Developing applications for brain–computer interaction. In *Proceedings of the Brain–Computer Interfaces for HCI and Games Workshop at the SIGCHI Conference on Human Factors in Computing Systems (CHI).*

Zander, T. O., Kothe, C. A., Jatzev, S., & Gärtner, M. 2010. Enhancing human–computer interaction with input from active and passive brain–computer interfaces. In D. S. Tan & A. Nijholt (Eds.), *Brain–Computer Interfaces* (pp. 181–199). London, UK: Springer.

Zander, T. O., Kothe, C. A., Welke, S., & Rötting, M. 2008b. Enhancing human–machine systems with secondary input from passive brain–computer interfaces. In *Proceedings of the 4th International Brain–Computer Interface Workshop & Training Course* (pp. 144–149). Graz, Austria: Verlag der Technischen Universität Graz.

Zander, T. O., & Krol, L. R. 2017. Team PhyPA: Brain–computer interfacing for everyday human–computer interaction. *Periodica Polytechnica Electrical Engineering and Computer Science, 61*(2), 209–216.

Zander, T. O., Krol, L. R., Birbaumer, N. P., & Gramann, K. 2016. Neuroadaptive technology enables implicit cursor control based on medial prefrontal cortex activity. *Proceedings of the National Academy of Sciences*, *113*(52), 14898–14903.

Zander, T. O., Lehne, M., Ihme, K., Jatzev, S., Correia, J., Kothe, C. A., Nijboer, F. (2011b). A dry EEG-system for scientific research and brain–computer interfaces. *Frontiers in Neuroscience, 5*, 53.

Zander, T. O., Shetty, K., Lorenz, R., Leff, D. R., Krol, L. R., Darzi, A. W., ... Yang, G.-Z. 2017b. Automated task load detection with electroencephalography: Towards passive brain–computer interfacing in robotic surgery. *Journal of Medical Robotics Research*, *2*(1), 1750003.