Science and Engineering Ethics (2019) 25:1217–1233
https://doi.org/10.1007/s11948-018-0061-1

ORIGINAL PAPER

Critiquing the Concept of BCI Illiteracy

Margaret C. Thompson[1,2]

Received: 20 April 2018 / Accepted: 10 August 2018 / Published online: 16 August 2018

Abstract

Brain–computer interfaces (BCIs) are a form of technology that read a user's neural signals to perform a task, often with the aim of inferring user intention. They demonstrate potential in a wide range of clinical, commercial, and personal applications. But BCIs are not always simple to operate, and even with training some BCI users do not operate their systems as intended. Many researchers have described this phenomenon as "BCI illiteracy," and a body of research has emerged aiming to characterize, predict, and solve this perceived problem. However, BCI illiteracy is an inadequate concept for explaining difficulty that users face in operating BCI systems. BCI illiteracy is a methodologically weak concept; furthermore, it relies on the flawed assumption that BCI users possess physiological or functional traits that prevent proficient performance during BCI use. Alternative concepts to BCI illiteracy may offer better outcomes for prospective users and may avoid the conceptual pitfalls that BCI illiteracy brings to the BCI research process.

Keywords Brain–computer interface · Brain–machine interface · BCI illiteracy · Research ethics · User-centered design

Introduction

Brain–computer interfaces (BCIs) have garnered attention as a potential communication and control interface for individuals with a variety of health conditions, including spinal cord injury, stroke, and autonomic lateral sclerosis (ALS). They are also attractive for commercial applications such as gaming, monitoring personal behavior, and more. BCIs are a general term for technologies that measure a user's brain activity in order to perform a task. Examples of BCIs include robotic arms

✉ Margaret C. Thompson
mthomps2@uw.edu

1 Department of Electrical Engineering, University of Washington, 185 Stevens Way, Paul Allen Center - Room AE100R, Campus Box 352500, Seattle, WA 98195-2500, USA

2 Center for Sensorimotor Neural Engineering, Box 37, 1414 NE 42nd St., Suite 204, Seattle, WA 98105-6271, USA

controlled by a user's intention to move, "spellers" that allow individuals to select letters to complete words or phrases that they wish to communicate, and monitors that measure brain activity and provide feedback to promote relaxation. Many BCIs rely on the user being able to voluntarily alter their brain activity in order to differentiate between a discrete menu of options (such as 26 letters on a keyboard) or a continuous range of outputs (such as the range of positions that a robotic arm might take, or a user's level of stress).

Under the umbrella of BCI technology, there are different neural mechanisms that may be used as a control signal. For example, in sensorimotor-rhythm (SMR) based BCIs, users may imagine moving a limb in order to generate predictable changes in spatial and temporal oscillations of neural populations in their motor cortex. An algorithm that has been calibrated or trained to recognize changes in these recorded signals then triggers a response based on recognition of the user's brain state. In another example, a P300 (a *positive* amplitude peak that occurs approximately 300 ms after the presentation of the stimuli) BCI may be constructed to utilize a response to an unexpected or "oddball" stimulus that the user is focusing on. By knowing when the stimulus is presented and when the P300 is sensed, the BCI can determine what the user was focusing on. BCIs can be constructed from other predictable neural responses to auditory, visual, tactile, and other stimuli.

Despite advances in many aspects of BCI performance, a subset of potential BCI users are unable to operate some or all types of BCIs despite undergoing the same training as successful users. This phenomenon has been termed *BCI illiteracy* and framed as a trait possessed by BCI users who do not reach a certain level of performance. Many researchers engaged in BCI research have applied the concept to subsets of users in their respective BCI studies (Ahn et al. 2013; Allison et al. 2007; Carabalona 2017; Kaufmann et al. 2013; Kübler and Müller 2007; Nijboer et al. 2008; Pfurtscheller et al. 2010; Shu et al. 2018; Suk et al. 2014; Yao et al. 2017). Some researchers claim that BCI illiteracy affects between 15 and 30% of people who attempt BCI use, regardless of the type of BCI used (Blankertz et al. 2010).

The implications of categorizing users as BCI illiterate may not be initially apparent. It may seem that BCI illiteracy is a benign descriptor of poor performance or even a helpful way to categorize users. However, in using BCI illiteracy as a concept, BCI researchers face a number of scientific and ethical challenges, the outcomes of which stand to directly impact BCI users and the progression of BCI research and development. This paper presents the argument that BCI illiteracy is an inadequate concept for explaining many cases of poor user performance in BCI systems. The concept of BCI illiteracy relies on poor assumptions and fails to consider key aspects about BCI systems and their relationship to their users.

Definitions of BCI Illiteracy and Notes on Terminology

BCI illiteracy is a proposed condition wherein users of BCI technology fail to reach proficiency in using a BCI in a standard training period. Through their experience across a variety of BCI studies, researchers estimate that 15–30% of BCI users could be labeled as BCI illiterate (Blankertz et al. 2010; Viduarre and Blankertz 2010).

Such difficulties are seen across multiple BCI types that rely on different neural signals; in fact, it is possible for an individual to be considered BCI illiterate with one form of BCI while being able to reach proficiency on another (Allison and Neuper 2010). BCI illiteracy as a concept relies on comparing the performance of an individual user to a norm which may vary with each BCI system. These norms have two key features. First, the norm sets a performance threshold representing proficiency. Second, the norm sets a finite time period during which users are expected to reach proficient performance. Users may be classified as illiterate if their performance with the BCI does not meet this norm: they do not reach a certain level of performance, or they do not reach it fast enough within the confines of the study and its training period. In light of the norms used to determine how users are labeled, discussing BCI illiteracy often encourages the use of normative terminology towards prospective users, such as *proficient* or *successful* users in contrast to *deficient*, *unsuccessful*, or *illiterate* users.[1]

The term BCI illiteracy evokes comparisons between learning to use BCIs and written or spoken language acquisition.[2] To avoid conceptual snares in how BCI use may or may not relate to language acquisition, some researchers have chosen to use the terms BCI inefficiency, BCI aphasia, or normative descriptions of performance (Allison and Neuper 2010). Criticisms of BCI illiteracy presented in this manuscript relate to the concept and its function in BCI research rather than the particular terminology used.[3]

Motivations for Categorizing Users as BCI Illiterate

Although it appears that BCI illiteracy has not explicitly been defended as a helpful concept in the field of BCI research, its use has been relatively unchallenged by BCI researchers. Potential benefits to using BCI illiteracy may be outlined by considering how researchers apply the concept in their research processes. Researchers may want to identify which users perform poorly on their BCIs, understand why they perform poorly, and determine how they may alter their BCIs in order to improve performance. Currently, using BCI illiteracy as a paradigm, they may identify which users perform poorly by labeling these users as BCI illiterate relative to a certain proficiency norm; they may examine groups of users labeled as BCI illiterate to find explanations or commonalities of why these users performed poorly; and using these explanations to guide their efforts, they may attempt to make modifications to the

[1] At times this critique will reference similar normative phrasing; the use of these terms in this paper are meant relative to and entangled with existing norms in BCI research rather than as independent, self-sufficient categories.

[2] A more detailed discussion of the similarities or differences between BCI skill and literacy, and the accuracy of such terminology, is provided by Brendan Allison and Christa Neuper (2010).

[3] This paper uses the phrase BCI illiteracy, as it is by far the most popular of these terms, but the arguments presented will apply to any concept similar to BCI illiteracy regardless of the specific terminology used.

BCI system (including hardware, software, and training protocols) to improve performance for those users.[4]

In addition to understanding why some users do not fit well with some BCI technologies, using BCI illiteracy in a predictive sense could identify BCI users who would likely perform poorly and prevent them from receiving devices or systems that they have difficulty operating, especially in the case of implanted systems that may require surgeries. System operability and effectiveness have significant influence in predicting continued use of assistive technologies such as BCIs (Riemer-Reiss and Wacker 2000), thus users who would likely be labeled as illiterate may be at risk of abandoning assistive BCIs. Several studies have already been undertaken to develop predictors of BCI illiteracy (Ahn and Jun 2015; Ahn et al. 2013; Blankertz et al. 2010; Suk et al. 2014), though there is currently no definitive way to screen for any type of BCI illiteracy on a patient-by-patient basis.

While employing BCI illiteracy as a tool in BCI research may sound straightforward, multiple challenges arise from how the concept is framed and used in practice. Prior work by Brendan Allison and Christa Neuper identifies multiple issues with BCI illiteracy, including inconsistent thresholds for proficiency, poor differentiation of the numerous causes for poor BCI performance, and more. They have taken these issues as opportunities to make BCI illiteracy more rigorous as a category, for example by producing standards for proficiency levels and outlining categories of causes of illiteracy (Allison and Neuper 2010). However, there are more fundamental issues at play in the way that BCI illiteracy is used. Specifically, BCI illiteracy relies on the assumption that prospective users ought to be able to use a given BCI system. In cases where they are unable to operate a BCI system, labeling these users as BCI illiterate clearly places the locus of deficiency on the user. BCI illiteracy reflects and affects the way that BCI systems are designed, and it may be misguided to approach it as a well-posed idea that needs only to be standardized. Instead, researchers ought to critically examine the way that the concept of BCI illiteracy affects the BCI systems that are produced and consider alternative concepts that may better serve a wide variety of BCI users.

Labeling Users as BCI Illiterate

In the BCI research process, the first step at which BCI illiteracy is often employed as a concept is in the identification or classification of users who are not compatible with the BCI system. A user might be labeled as BCI illiterate after participating in a BCI study if their performance was not considered proficient on that particular BCI system. As outlined above, proficient performance is a norm with two components: an accuracy level that users are expected to achieve, and a finite time period (generally the length of the experiment) within which users are expected to achieve this accuracy.

[4] Individual BCI researchers may not go through all of these steps—some just stop at labeling, some only propose to design BCIs for illiterate populations, etc.

The way that this norm is determined remains largely unjustified in BCI studies and may be inconsistent between studies (Allison and Neuper 2010). To verify that they are in fact exerting control over the BCI, users are expected to perform above "chance" accuracy (that is, the expected accuracy if the user were exerting no control over the system); beyond this criterion, the choice for proficiency thresholds may in reality be relatively arbitrary. For example, in a two-target task where chance accuracy is 50%, the level for proficiency (and the distinction between proficient and illiterate users) may be cited as 60% accuracy (Guger et al. 2003) or 70% accuracy (McCane et al. 2014). The decisions for these thresholds are made by researchers but are not necessarily justified when results are presented. Lynn McCane and colleagues, who demonstrated a communication BCI for ALS patients, set their proficiency threshold at 70%, citing Eric Sellers and colleagues that this was necessary for "acceptable communication" (McCane et al. 2014); yet Eric Sellers and colleagues only noted that "while accuracy levels less than 60% provide substantial [information conveyed for each user selection], the time needed to produce useful communication might be unacceptable" (Sellers et al. 2006: 251). Other studies do not provide any overt justification for their performance thresholds and may instead rely on trends from other research studies.

Manipulating proficiency thresholds can substantially change the population of users labeled as illiterate. Allison and Neuper examine one such case where, without explicit justification, the researchers chose a proficiency threshold of 60% and observed 6.7% of subjects below the threshold; a threshold of 70% in the same study would have labeled 48.7% of subjects as deficient (Allison and Neuper 2010). While a careful reviewer might question the 60% proficiency threshold, 6.7% failure rate sounds like a mostly-good BCI that only fails a few users; on the other hand, a 48.7% failure rate sounds like a design flaw in the BCI that would be difficult to attribute to users.

Standards and Justifications for BCI Performance Thresholds

Recent work in metrics for evaluating BCI performance has looked for methods of comparing performance more effectively across studies (Thomas et al. 2013). These efforts would seem to answer criticisms such as those raised by Brendan Allison and Christa Neuper that the thresholds used for BCI illiteracy are not standardized across studies (Allison and Neuper 2010). But standardized thresholds, while important for comparisons, are not enough. In addition to being standardized, and among other possible criteria, performance thresholds should be *justifiable*; researchers ought to be able to support, beyond simple comparisons to chance accuracy, why they believe that level of performance is necessary.

Justification for performance thresholds could be derived directly from the users that BCIs are expected to serve. Notably, user input on performance thresholds is almost entirely absent. The accuracy levels chosen for proficiency are not explicitly validated as acceptable to BCI recipients, just as lower accuracy values have not clearly been deemed unacceptable. Depending on the recipients in question—whether they are able-bodied, familiar with technology, etc.—the desired level of accuracy may vary. It may not be possible to find a universal standard for performance, but understanding the differing needs of various user populations could improve efforts to produce specialized BCIs.

Aside from direct user input, proficiency levels could be drawn relative to alternative existing solutions that BCIs are meant to augment or replace. There are few existing comparisons of desired performance levels—in terms of communication rate, accuracy, ease of use, etc.—between BCIs and other systems. For example, BCIs for communication could compare communication rates to other selection systems, such as partner-assisted scanning, for individuals with locked-in syndrome. Until a justified standard for acceptable performance—which may vary depending on the BCI system, the use case, and the user population in question—is determined, the performance thresholds set by BCI researchers remain speculative.

Lastly, it is important to note accuracy metrics used in performance thresholds may be defined relative to different aspects of system performance. For example, *classification accuracy* is the ability of the system to correctly determine the user's choice, whereas a quantity like *typing accuracy* in a spelling system is the ability of the system to output a desired string of letters. These quantities may differ depending on error handling and other system characteristics. Many BCI studies are vague regarding these definitions, making it more difficult to compare performance across studies. Furthermore, users may care more about a quantity like typing accuracy, which manifests in the system output that they observe, rather than a quantity like classification accuracy which may be masked by other system processing. Future research ought to provide justification for not only the accuracy level a threshold uses, but which definition of accuracy is relevant based on user needs and experience.

Consequences of Labeling Users

The impact of BCI illiteracy as a classification mechanism may be considered on two separate scales: that of individual users and that of populations of users. For individuals, there is currently little consequence for being labeled as BCI illiterate outside the confines of the research study. Users may not even be explicitly aware of their label, as these terms may be applied only in summarizing the results of the study. The impact for individual users could change as BCIs become more prevalent, or particularly if screening measures for BCI illiteracy were used in practice in determining who receives a given BCI. If BCIs are used more frequently by the general public in the future, then individuals who cannot receive a usable BCI may be separated from important resources that others could commonly access including assistive devices for communication and control, entertainment opportunities, and more.

Beyond the problems posed by classification for individual users, the process of labeling populations of users as proficient or deficient has implications for future BCI research efforts. Because proficiency thresholds determine which users will be identified as BCI illiterate, researchers using poorly justified thresholds may reach specious conclusions about populations of users labeled as BCI illiterate. This process is in its early stages, but some researchers have undertaken efforts to study how demographics such as age, gender, and other factors correlate with BCI performance (Allison et al. 2010). Although BCI illiteracy predictors may be implemented with the intent "to avoid the frustrating and costly procedure of trying to establish BCI control" (Blankertz et al. 2010: 1304), issues of representation could arise in future BCI studies depending on how such demographic correlations are used. Studying user groups who experience difficulty using BCI systems is a good undertaking, but researchers should think carefully about how they define this group of users so as to avoid missteps in identifying supposed causes for poor performance during BCI use, and how they use this information to influence representation in BCI studies.

Identifying Causes of Poor User Performance in BCI Systems

A subsequent function for BCI illiteracy as a concept may be to attempt to identify common traits—so-called causes of poor performance—in users who have not successfully operated a given BCI system. Once a group of users has been identified who operate below the performance norm for a given system, researchers are interested in understanding what characteristics of these users may have contributed to their performance. Analyzing features of these users collectively may seem to aid in identifying relevant correlations. If there are three users labeled as illiterate in a study, for example, each user may share a single trait that was linked to poor system performance in all three cases, or they may have a diversity of reasons why the BCI system was unsuitable for them. BCI researchers wish to understand causal traits (which they presume the users to possess) that result in user performance below their study's threshold.

The described state of BCI illiteracy is a label for a number of underlying causes. There may be physiological causes, like a fundamental mismatch between the users' capacities and the requirements of the BCI interface. This mismatch may be obvious, such as when an individual is blind and the BCI system uses a visual interface, or it may be less obvious, such as when a BCI system uses P300 signals but a user does not generate such signals. Aside from a decisively physiological cause, a user may not acquire the necessary skills to operate the interface if the control method is challenging or cognitively taxing, or if the training protocol is not effective. It may be hard to disentangle whether the user simply needs a different set of instructions to help them learn the task, or if a more fundamental difference in their neurophysiology is the cause. In either case, it is necessary to reconsider certain intuitions and assumptions underlying BCI illiteracy which have not previously been examined by researchers who employ it as a concept.

Physiological Causes of Poor Performance

Some cases of user difficulty which have been labeled as BCI illiteracy may stem from a mismatch in the user's physiology and the demands of the BCI system. This group of causes could be called neurophysiological. A neurophysiological cause of BCI illiteracy would be one where, for a given patient, the BCI system is incapable of detecting the appropriate neural signals for reasons that may or may not be obvious to researchers. An intuitive distinction is that, regardless of the effort or skill exerted by the user, they are theoretically incapable of operating that BCI (either temporarily or permanently).

Users who would likely be labeled as BCI illiterate—who have been included in a BCI study because they have no obvious excluding impairments—are not immediately differentiable from successful users based on obvious external physical or cognitive symptoms. When BCI illiteracy is employed as a label, it is often the case that researchers do not have a good explanation for why the BCI in question did not work for a subset of users. The differences in brain structure that could lead to BCI illiteracy may not be to the extent that they inhibit other functioning. For some individuals, their brain signals may be difficult to read through their scalp, for example if the relevant neural populations are located in folds of the brain (sulci) from where surface electrodes cannot record (Allison and Neuper 2010). A noninvasive BCI which relies on recording signals from these neural populations may not provide a usable interface for such users. Such physiological differences may only be discovered after a user has enrolled in a BCI study and has failed to achieve proficiency on a BCI. Furthermore, it is possible that such physiological differences would only be discussed theoretically and never verified on a user-by-user basis.

Physiological contributors to poor performance in a BCI system are not necessarily permanent features of the subject. One simple example of how users can move in and out of the category of BCI illiteracy is to consider how transient factors, such as changes in attention level, fatigue, or frustration (Myrden and Chau 2015), or social factors, such as emotional responses, interactions, and social cues experienced during BCI training or use (Sexton 2015), may cause a user's brain signals to change such that they operate their BCI less successfully. Such confounding factors are often referenced as reasons why BCIs that have been tested successfully in a laboratory setting could fail to work as well when tested outside of a controlled environment in the "real world". There is no guarantee that the individual labeled as BCI illiterate would be unable to use the same BCI on a day when they have altered their caffeine intake, had more or less sleep, or altered their personal state in another way. These effects are not necessarily reflective of some permanent, deficient trait that the BCI user possesses. Similarly, there's no guarantee that a proficient user may not later exhibit an illiterate level of performance. An individual with a progressive disease may be prescribed a BCI that they can initially operate, but later lose the ability to use it with proficiency because of changes in their health.

BCI illiteracy has not been constructed to handle these transient factors, and researchers who wish to maintain their use of the concept will need to clarify how they are labeling their users with respect to temporary changes in their performance levels. Are users labeled as BCI illiterate during the hour after they have had a coffee, but not other hours? On days when they have had too little sleep, but not other days? Can previously proficient users suddenly acquire the label of illiteracy? Answers to such questions regarding transient factors that render a user more or less literate could impact the way that researchers conceptualize problems in BCI design.

Non-physiological Causes of Poor Performance

Suppose researchers cannot find a temporary or permanent biological explanation—in brain structure or function—to explain a user's poor BCI performance. In this case, the user is theoretically capable of controlling the BCI, but may lack the skill needed to operate the system or they may not fully understand the instructions given by researchers. Researchers may then turn to factors such as training protocols, user effort, and other non-physiological contributors that could explain deficient performance.

Intuitively, when considering contributors such as training protocols or user effort, it may be helpful to compare BCI skill acquisition with how individuals learn to use other new technologies. Labeling users as BCI illiterate may reflect an erroneous assumption that current BCI systems are optimized to be simple to learn and use. BCI is a relatively young technology that in most cases is unintuitive to learn. Other emerging technologies were also unintuitive to learn when they were first commercialized, such as automobiles and computers. It is presumptuous to believe that current BCI systems are inherently intuitive to use or that the best possible training protocols have been designed.

Research supports the idea that many instances described as BCI illiteracy may be due to the structure of current BCI training protocols. In a recent study, researchers asked individuals to perform a simple, non-BCI motor task (drawing circles and triangles on a tablet screen). The training approach that participants were given to learn the task was designed to mimic those given in most BCI studies; for example, training schemes were not adapted to individual users and users were given limited autonomy in their own training. This study found that, even in the case of a simple motor task which participants were theoretically capable of executing, approximately 15% of participants who were physically capable of performing the drawing task did not successfully learn to draw circles and triangles that the system could differentiate at desired times (Jeunet et al. 2014). This percentage is similar to the ratio of individuals who fail to reach proficiency in BCI tasks with similar training schemes, suggesting that the training protocols (rather than some complexity of the BCI system) may contribute to instances of BCI illiteracy.

In addressing the shortcomings of BCI training schemes, Fabien Lotte and colleagues reference specific training literature to understand how to improve BCI training protocols (2013). For example, current BCI systems may only show a user whether they have successfully hit a target (whether their signals were accurate) with no feedback about how they might improve their control efforts (for example by imagining a different type of movement, imagining a different sensation, etc.). Instead, feedback should explicitly inform the user *how* to correct their output signals, rather than just whether their signals were accurate. Further, goals should be explicitly stated; the user's goal, when clearly stated, may be to generate brain signals with certain qualities or characteristics rather than hit a target with the cursor (Lotte et al. 2013). Improved BCI training schemes should be of particular interest to BCI researchers because they could increase the number of users who are able to reach proficiency without the need for costly improvements to hardware or signal processing. Greater discussion and documentation of existing training approaches could provide insight as to what is working and what is not with respect to training protocols, and inspire researchers on what to try next.

Challenges in Inferring Causes from Users Labeled as Illiterate

In grouping and studying users labeled as illiterate, researchers may suspect these users to share causal traits. But these users may only share in common the ratio between their performance and the proficiency threshold set by researchers. If this ratio does not track some phenomena in the user or some feature of the interface, it is hard to justify its use at all. Information developed around shared traits of these poorly performing users—predictors and screenings, underlying neuroscientific attributes, and more—would then seem suspect.

Current approaches by BCI researchers to develop a body of knowledge regarding causes of BCI illiteracy, methods for illiteracy prediction, and solutions for the condition, could be referred to as the *biologization* of BCI illiteracy as a descriptive category. The category "BCI illiterate" could be understood as a *human kind* described by Ian Hacking; it is a socially constructed category for sorting humans around which people might try to develop knowledge. Hacking describes a common push to biologize these socially constructed human kinds—for example, scanning criminals to research what brain activity is correlated with violent or criminal tendencies (Hacking 1995). Language surrounding BCI illiteracy reflects the process of biologization, such as the paper titled "Towards a cure for BCI illiteracy", implying that BCI illiteracy is a condition to be treated and cured (Viduarre and Blankertz 2010). Research in this vein implicitly assumes that there are underlying characteristics of the user, whether structural or functional, to be discovered that will explain BCI illiteracy. If such features could be identified, researchers might aim to develop fixes for these user characteristics to improve existing BCIs, or create new BCIs that are better suited to these "problematic" variations in physiology.

Developing a body of knowledge around the perceived condition of BCI illiteracy poses a challenge. In the creation of knowledge about BCI illiteracy, researchers study a sample of users that have been labeled as BCI illiterate. This labeling occurs through interaction with an imperfect BCI system, including comparing the users to a proficiency threshold which may not have been determined from a compelling basis. Even if the performance threshold is well justified, the process of identifying user traits based solely on the correlation of a performance metric is ill-defined, especially if those traits are only transient. Understanding the causes of BCI illiteracy is difficult because the inclusion of users under such a label is messy even with rigorous performance standards.

Challenges in Framing the Cause of Poor Performance

In addition to relying on questionable methods for gaining understanding of user populations, BCI illiteracy as a concept relies on the assumption that the user "owns" the problem of poor performance, the implications of which have been noted before. For example, Allison and Neuper note that, "it is also unclear whether ‘BCI illiteracy’ reflects a failure on behalf of the subject or BCI, and whether this distinction is meaningful,” but also acknowledge that from a framing perspective, “‘BCI illiteracy’ implies that failure to use a BCI results from inadequate effort by the user, which is generally not true… some subjects could never learn to use a particular BCI” (Allison and Neuper 2010: 50). Yet even if researchers claim that BCI illiteracy is not meant to blame the user for poor performance, the reality remains that BCI illiteracy is a label applied to the user rather than the BCI system. The locus of deficiency, regardless of intention, is placed on the user.

This distribution of deficiency is not novel; it mirrors a similar conflict between medical and social models of disability. These models pose two alternative frameworks for the locus of disability as well as proper responses to perceived problems. As philosopher Anita Silvers explains,

> …the medical model takes disability to be a problem requiring medical intervention—and as both the prerogative and the responsibility of medical professionals to fix—the social model understands disability as a political problem calling for corrective action by citizen activists who alter other people’s attitudes and reform the practices of the state. (Silvers 2009: 19)

That is, the medical model of disability assumes disability to be a trait possessed by the disabled individual. The social model views disability as a societal factor resulting from a mismatch in an individual’s abilities and the environment with which they must interact. If researchers adopt a medical-type model of poor BCI performance, where the user possesses some temporary or permanent condition that inhibits them from successfully operating their BCI, then the concept of BCI illiteracy may seem appropriate. This inhibition may be perceived to be biological, such as in the case that users cannot generate appropriate signals at the electrode recording site; or it may be perceived to be cognitive or psychological, such as if they cannot properly follow a training protocol to learn to manipulate their generated signals successfully. If, however, researchers adopt a social-type model of these same scenarios, then the problem lies in the lack of fit between user and BCI paradigm, not necessarily in the individual user. These scenarios could be reframed to say instead that a BCI system is flawed if it fails to meet the needs of its intended users. For example, a BCI may have been designed with inappropriate electrode sites for certain users, or a BCI training protocol may be too complicated for a subset of individuals.

Similar intuitions have been raised in other fields of technological development. Describing their approach towards developing effective security frameworks that require the participation of human operators, the People-Centered Security Lead for the National Cyber Security Centre of the UK recently noted that “if security doesn’t work for people, it doesn’t work” (Emma W 2017). In response to problems of framing presented by BCI illiteracy, a similar mentality might be applied to BCI platforms which are meant to serve in assistive or commercial applications. If a BCI doesn’t work for its users, it doesn’t work.

Does the Concept of BCI Illiteracy Allow BCI Researchers to Improve BCIs?

A last step in the BCI research process where BCI illiteracy may be employed as a concept is in the improvement of BCI systems. Having labeled a group of users as illiterate, and having studied this group for potential causes of illiteracy, researchers may wish to close the developmental circle by applying changes to BCI systems that "cure" cases of illiteracy. There is not enough evidence in the field of BCI research to suggest that this step of the process happens effectively. Furthermore, taking this approach to improve BCIs is at the sacrifice of alternative approaches, which may result in better outcomes for BCI users.

Practically speaking, the users labeled as illiterate in one BCI study may not be available for a follow-up study in which improved BCIs are tested. Although this paper has outlined three general steps where BCI illiteracy may be applied as a concept in the research process, often an individual research team is not engaged in all three of those steps. Some focus on design of their own BCI and simply label users who perform poorly within their study (step 1) without further investigating the causes that led to this poor performance; other groups are focused on developing predictors of illiteracy divorced from a particular BCI system (step 2). Thus, assessing whether insights gained from studying illiterate-labeled individuals actually manifest in effective improvements is challenging. Few groups have claimed to accomplish the aim of designing BCIs that reduce the occurrence of BCI illiteracy. As an example, one group applied machine learning techniques to improve performance of SMR BCIs, and noted that these techniques improved performance for illiterate users such that they functioned above a 70% proficiency threshold (Viduarre and Blankertz 2010). Despite these improvements, the users' performance levels for the illiterate-labeled group at the end of the study were still below that of users initially labeled as proficient.

In the example above, it is unclear whether the changes made to the BCI designs were driven by an understanding of the needs of users, and specifically those labeled illiterate. It is relevant to consider the relationship between which aspects of BCI design are the most researched versus the previously discussed possible causes for BCI illiteracy. Surveys of the field have noted that most prior efforts to improve BCI design have focused on changes to signal processing and algorithms (Banville and Falk 2016); more specifically, these efforts have often been implemented through offline studies—that is, those that are run without recruiting users (Pasqualotto et al. 2012). While these efforts may make use of limited face-to-face time with end-users and avoid the challenges of developing novel hardware, they have their limits. Such efforts do not answer highlighted problems with BCI training protocols, and they do not improve the incorporation of other user feedback into BCI design, both of which have been highlighted in a recent review as key areas for improvement in BCI systems (Chavarriaga et al. 2017). This disconnect between identified issues contributing to BCI illiteracy and the improvements pursued by many BCI researchers suggests that knowledge developed from the concept of BCI illiteracy may not be driving real improvements in BCI design.

Ultimately, developing BCIs under the framework of BCI illiteracy not only places the locus of deficiency on the user of the system, it also reflects a research process in which a system is designed by researchers and its effectiveness for users is later assessed. While such an approach does not preclude considering user needs in the initial design of the BCI system, it certainly does not seem to prioritize involving users at all stages of design. Deriving BCI system requirements and metrics directly in partnership with user populations may be preferable for improving how BCI designs are generated by making researchers more knowledgeable about the design problem at hand, in contrast to using the concept of BCI illiteracy to judge between BCI designs that have already been implemented (Whitbeck 1996). BCI illiteracy is one approach to the goal of developing effective BCI systems, but there is reason to question whether it is the most straightforward or effective approach.

Alternatives to BCI Illiteracy

Nowhere in the published BCI literature is the idea of BCI illiteracy defended as a valuable way to label potential users. During early BCI studies, as researchers identified groups of individuals who were unable to effectively use their BCI systems, the concept arose organically without explicit consideration for its effects. BCI researchers were simply attempting to explain a phenomenon that they had witnessed in their experiments. Yet the concept of BCI illiteracy as it is currently structured obscures further scientific investigation into reasons for poor performance in BCI systems. Furthermore, the concept used to explain this phenomenon, whether named "illiteracy," "inefficiency," "aphasia," or something else, places the problem inherently on the user of the BCI system. Even if researchers are interested in developing BCIs with better performance, they do so while framing the problem as a trait which poorly performing users possess, rather than framing it as a failure of the BCI system. Further standardizing BCI illiteracy as a category does not address this fundamental problem of framing.

BCI illiteracy is not the only possible approach to characterizing, understanding, and solving problems of poor performance on BCI systems. Rather than using the conceptual framework of BCI illiteracy, researchers could incorporate BCI users throughout the design process so that they help to define the standard to which BCI systems must adhere. One potential paradigm for this incorporation is user-centered design (UCD). UCD focuses on usability, or "how well a specific technology suits its purpose and meets the needs and requirements of the targeted end-users" (Kübler et al. 2013: 55); in standards for UCD, usability is defined in terms of effectiveness, efficiency, and user satisfaction (Kübler et al. 2013). In the UCD framework, users would suggest a minimum level of usability which the BCI system needs to meet. By leveraging UCD principles to identify usability as a priority or even as a design constraint—that is, a criterion that a final BCI design *must* satisfy for a given population—researchers would not leave room for concepts like BCI illiteracy. BCI systems that do not meet a minimum level of usability for their intended users will be revised or rejected as a failure in design rather than a failure on the part of the user. This is not to say that a single BCI must work for all users, but rather that any usability issues stem from limitations in the BCI rather than the users who interact with it.

Although uncommon, this approach to BCI development is not without precedent. Prior work includes studies where researchers used tailored approaches for different subjects, including flexibility about visual versus auditory modes of feedback (Schreuder et al. 2013) and flexibility about possible control schemes depending on which was most effective on a user-by-user basis (Friedrich et al. 2013). Other studies simply aimed to design a speller system for “layman use” that offered a simpler interface than prior spellers (Kaufmann et al. 2012). In another example, Andrea Kübler and colleagues applied UCD principles to discuss BCI usability with 19 prospective users. Participants readily provided feedback that could inform performance thresholds, such as “five times faster would be acceptable,” or “with my own AT I can write 90 characters per minute” (2014: 15). These statements provide a kind of justification previously lacking in performance thresholds chosen by BCI researchers. Notably, the studies that take this approach do not make references to concepts such as BCI illiteracy.

A UCD approach does not require a BCI system to function perfectly “out-of-the-box” with no training time or effort on the part of the user. Some users could enjoy learning to use their BCI, provided the training was designed in a way that users found meaningful and fulfilling. As in the case of determining other BCI development priorities, potential end users should be the ones to dictate how much of a user training burden is acceptable in a fully designed, commercial BCI system.

A final caveat is that a UCD approach may identify some users that feel other technologies are a better fit for their needs, such as eye-tracking devices for individuals who are not fully locked-in (Pasqualotto et al. 2015). BCI researchers ought to be frank about the limitations of their technologies. BCIs, especially those that are fully implanted, are currently an expensive technology that may be inaccessible to many potential users. Wearable, noninvasive BCIs are generally less effective, or users may find them inconvenient or embarrassing to use in public. Under a UCD approach, it is especially critical that potential recipients of assistive technologies be integrated in the design process from the very start to ensure that BCIs are produced not just because they are an impressive technology that researchers feel potential users should want, but rather because they are serving a need.

Beyond BCI Illiteracy: Guidelines for Improving BCI Research Efforts

If the field of BCI research were to move away from concepts like BCI illiteracy, how would research processes change? Recent efforts to incorporate UCD approaches into BCI research align intuitively with the idea of reframing the BCI system as the locus for improvement. BCI requirements—necessary accuracy and speed of communication, ease of calibration and training, comfort and required effort, etc.—would be driven directly by user needs, and systems that did not meet these needs would be seen as deficient. These approaches have only been applied in a handful of studies (Friedrich et al. 2013; Kaufmann et al. 2012; Kübler et al. 2014; Schreuder et al. 2013), and much remains to be discovered about how users might be effectively incorporated throughout BCI development. Even the process of eliciting such feedback from prospective users with severe impairments such as locked-in syndrome may be an area for research (Peters et al. 2016). Researchers may also face challenges in trying to design BCIs to serve diverse populations with conflicting needs, and methods to balancing such perspectives in the specific case of BCI systems would also be an area for investigation.

In addition to these preliminary efforts, the following guidelines might be useful for navigating situations in the BCI research process where it is tempting to use BCI illiteracy:

1. Define and understand the user population(s) the BCI system aims to serve. Use the needs of these populations to inform BCI design. Aim to understand whether subpopulations have heterogeneous needs that may result in different design considerations for different individuals.
2. When evaluating the performance of BCI systems, clearly define how metrics such as accuracy are calculated. Justify metrics and the thresholds used to define success based on user input, which may be derived directly or in comparison to user feedback on alternative technologies.
3. Define poor performance in terms of shortcomings of the BCI system. For example, a P300 speller BCI is not suitable for users who do not generate a standard P300 response to oddball stimuli. A BCI that doesn't work for its users doesn't work.

Conclusion

BCIs have shown promise for a variety of applications, from assistive technology to personal use or as entertainment platforms. For all of the excitement surrounding BCIs, most researchers in the field would agree that fundamental issues remain to be solved. Many BCIs are not reliable enough for use outside a controlled experimental setting, or they are more cumbersome than they are assistive. There are clearly opportunities for improvement in the design of BCI systems.

It should come as no surprise that, in addition to improvements to BCI design, there may be areas to improve the research processes that generate BCIs. BCI illiteracy is just one example of a conceptual framework that may not be optimal for understanding or developing BCI systems. Though initially employed to understand how users interact with BCI systems, there are clear drawbacks to BCI illiteracy as a concept. These drawbacks include limited conceptual rigor, such as poorly justified performance thresholds or a variety of underlying factors that complicate BCI design which may be conflated. More fundamentally, BCI illiteracy assumes a problematic framing: that the user possesses a state or quality of deficit that may be discovered through their interaction and poor performance on a BCI system. Researchers tasked with developing usable BCI systems should take pause at the notion that they are defining their users in terms of their ability to use these systems.

Users face legitimate difficulties in operating BCI systems. Focusing research efforts on addressing these difficulties is an important pursuit for BCI researchers.

But the conceptual approaches used to frame and investigate these difficulties are not a foregone conclusion, and BCI illiteracy is not the only possible approach. Critiquing concepts such as BCI illiteracy and seeking alternative frameworks will not necessarily result in more usable BCIs or better incorporation of user views in BCI development; but the prospect of exploring and revising the framing of users of BCI systems in order to improve BCI research processes is an opportunity that could provide rich returns and should not be overlooked.

Acknowledgements The author would like to thank the Neuroethics group at the Center for Sensorimotor Neural Engineering, especially Dr. Laura Specker Sullivan, Timothy Brown, Marion Boulicault and Dr. Sara Goering for their help in refining this critique from early to final drafts.

References

Ahn, M., Cho, H., Ahn, S., & Jun, S. C. (2013). High theta and low alpha powers may be indicative of BCI-illiteracy in motor imagery. *PLoS ONE, 8*(11), e80886.

Ahn, M., & Jun, S. C. (2015). Performance variation in motor imagery brain–computer interface: A brief review. *Journal of Neuroscience Methods, 243,* 103–110.

Allison, B., Luth, T., Valbuena, D., Teymourian, A., Volosyak, I., & Graser, A. (2010). BCI demographics: How many (and what kinds of) people can use an SSVEP BCI? *IEEE Transactions on Neural Systems and Rehabilitation Engineering, 18*(2), 107–116.

Allison, B. Z., & Neuper, C. (2010). Could anyone use a BCI? In D. S. Tan & A. Nijholt (Eds.), *Brain–computer interfaces* (pp. 35–54). London, UK: Springer.

Allison, B. Z., Wolpaw, E. W., & Wolpaw, J. R. (2007). Brain–computer interface systems: progress and prospects. *Expert Review of Medical Devices, 4*(4), 463–474.

Banville, H., & Falk, T. H. (2016). Recent advances and open challenges in hybrid brain–computer interfacing: a technological review of non-invasive human research. *Brain–Computer Interfaces, 3*(1), 9–46.

Blankertz, B., Sannelli, C., Halder, S., Hammer, E. M., Kübler, A., Müller, K. R., Curio, G., & Dickhaus, T. (2010). Neurophysiological predictor of SMR-based BCI performance. *Neuroimage*, *51*(4), 1303–1309.

Carabalona, R. (2017). The role of the interplay between stimulus type and timing in explaining BCI-illiteracy for visual P300-based brain–computer interfaces. *Frontiers in Neuroscience, 11,* 363.

Chavarriaga, R., Fried-Oken, M., Kleih, S., Lotte, F., & Scherer, R. (2017). Heading for new shores! Overcoming pitfalls in BCI design. *Brain–Computer Interfaces, 4*(1–2), 60–73.

Emma W. (2017). People: The strongest link. Resource document. National Cyber Security Centre. https://www.ncsc.gov.uk/information/people-strongest-link. Accessed 28 March 2017.

Friedrich, E. V. C., Neuper, C., & Scherer, R. (2013). Whatever works: A systematic user-centered training protocol to optimize brain–computer interfacing individually. *PLoS ONE, 8*(9), e76214.

Guger, C., Edlinger, G., Harkam, W., Niedermayer, I., & Pfurtscheller, G. (2003). How many people are able to operate an EEG-based brain–computer interface (BCI)? *IEEE Transactions on Neural Systems and Rehabilitation Engineering, 11*(2), 145–147.

Hacking, I. (1995). The looping effects of human kinds. In D. Sperber, D. Premack, & A. J. Premack (Eds.), *Symposia of the Fyssen Foundation. Causal Cognition: A Multidisciplinary Debate* (pp. 351–394). New York, NY: Clarendon Press/Oxford University Press.

Jeunet, C., Cellard, A., Subramanian, S., Hachet, M., N'Kaoua, B., & Lotte, F. (2014). How well can we learn with standard BCI training approaches? *A pilot study*, *Presented at 6th International Brain–Computer Interface Conference*, Graz, Austria, September 2014.

Kaufmann, T., Schulz, S. M., Köblitz, A., Renner, G., Wessig, C., & Kübler, A. (2013). Face stimuli effectively prevent brain–computer interface inefficiency in patients with neurodegenerative disease. *Clinical Neurophysiology, 124*(5), 893–900.

Kaufmann, T., Völker, S., Gunesch, L., & Kübler, A. (2012). Spelling is just a click away–a user-centered Brain–computer interface including auto-calibration and predictive text entry. *Frontiers in Neuroscience, 6,* 72.

Kübler, A., Holz, E. M., Riccio, A., Zickler, C., Kaufmann, T., Kleih, S. C., Staiger-Sälzer, P., Desideri, L., Hoogerwerf, E. J., & Mattia, D. (2014). The user-centered design as novel perspective for evaluating the usability of BCI-controlled applications. *PLoS ONE*, *9*(12), e112392.

Kübler, A., Mattia, D., Rupp, R., & Tangermann, M. (2013). Facing the challenge: Bringing brain–computer interfaces to end-users. *Artificial Intelligence in Medicine, 59*(2), 55–60.

Kübler, A., & Müller, K. (2007). An introduction to brain–computer interfacing. In G. Dornhege, J. R. Millán, T. Hinterberger, D. J. McFarland, & K. Müller (Eds.), *Toward brain–computer interfacing* (pp. 1–25). Cambridge, MA: The MIT Press.

Lotte, F., Larrue, F., & Mühl, C. (2013). Flaws in current human training protocols for spontaneous Brain–computer interfaces: Lessons learned from instructional design. *Frontiers in Human Neuroscience, 7,* 568.

McCane, L. M., Sellers, E. W., McFarland, D. J., Mak, J. N., Carmack, C. S., Zeitlin, D., Wolpaw, J. R., & Vaughan, T. M. (2014). Brain–computer interface (BCI) evaluation in people with amyotrophic lateral sclerosis. *Amyotrophic Lateral Sclerosis and Frontotemporal Degeneration*, *15*(3–4), 207–215.

Myrden, A., & Chau, T. (2015). Effects of user mental state on EEG–BCI performance. *Frontiers in Human Neuroscience, 9,* 308.

Nijboer, F., Sellers, E. W., Mellinger, J., Jordan, M. A., Matuz, T., Furdea, A., Halder, S., Mochty, U., Krusienski, D. J., Vaughan, T. M., & Wolpaw, J. R. (2008). A P300-based brain–computer interface for people with amyotrophic lateral sclerosis. *Clinical Neurophysiology*, *119*(8), 1909–1916.

Pasqualotto, E., Federici, S., & Belardinelli, M. O. (2012). Toward functioning and usable brain–computer interfaces (BCIs): A literature review. *Disability and Rehabilitation: Assistive Technology, 7*(2), 89–103.

Pasqualotto, E., Matuz, T., Federici, S., Ruf, C. A., Bartl, M., Olivetti Belardinelli, M., Birbaumer, N., & Halder, S. (2015). Usability and workload of access technology for people with severe motor impairment: A comparison of Brain–computer interfacing and eye tracking. *Neurorehabilitation and Neural Repair*, *29*(10), 950–957.

Peters, B., Mooney, A., Oken, B., & Fried-Oken, M. (2016). Soliciting BCI user experience feedback from people with severe speech and physical impairments. *Brain–Computer Interfaces, 3*(1), 47–58.

Pfurtscheller, G., Allison, B. Z., Bauernfeind, G., Brunner, C., Solis Escalante, T., Scherer, R., Zander, T. O., Mueller-Putz, G., Neuper, C., & Birbaumer, N. (2010). The hybrid BCI. *Frontiers in Neuroscience*, *4*, 3.

Riemer-Reiss, M. L., & Wacker, R. R. (2000). Factors associated with assistive technology discontinuance among individuals with disabilities. *Journal of Rehabilitation*, *66*(3), 44–50.

Schreuder, M., Riccio, A., Risetti, M., Dähne, S., Ramsay, A., Williamson, J., Mattia, D., & Tangermann, M. (2013). User-centered design in brain–computer interfaces—A case study. *Artificial Intelligence in Medicine*, *59*(2), 71–80.

Sellers, E., Krusienski, D., McFarland, D., Vaughan, T., & Wolpaw, J. (2006). A P300 event-related potential brain–computer interface (BCI): The effects of matrix size and inter stimulus interval on performance. *Biological Psychology, 73,* 242–252.

Sexton, C. A. (2015). The overlooked potential for social factors to improve effectiveness of brain–computer interfaces. *Frontiers in Systems Neuroscience, 9*(233), 693–695.

Shu, X., Chen, S., Yao, L., Sheng, X., Zhang, D., Jiang, N., Jia, J., & Zhu, X. (2018). Fast recognition of BCI-inefficient users using physiological features from EEG Signals: A screening study of stroke patients. *Frontiers in Neuroscience*, *12*, 93.

Silvers, A. (2009). An essay on modeling: The social model of disability. *Philosophical reflections on disability* (pp. 19–36). Netherlands: Springer.

Suk, H., Fazli, S., Mehnert, J., Müller, K., & Lee, S. (2014). Predicting BCI subject performance using probabilistic spatio-temporal filters. *PLoS ONE, 9*(2), e87056.

Thomas, E., Dyson, M., & Clerc, M. (2013). An analysis of performance evaluation for motor-imagery based BCI. *Journal of Neural Engineering, 10,* e031001.

Viduarre, C., & Blankertz, B. (2010). Towards a cure for BCI illiteracy. *Brain Topography, 23*(2), 194–198.

Whitbeck, Caroline. (1996). Ethics as design: Doing justice to moral problems. *The Hastings Center Report, 26*(3), 9–16.

Yao, L., Sheng, X., Zhang, D., Jiang, N., Farina, D., & Zhu, X. (2017). A BCI system based on somatosensory attentional orientation. *IEEE Transactions on Neural Systems and Rehabilitation Engineering, 25*(1), 81–90.