Physiological Computing and Intelligent Adaptation

Stephen H Fairclough

Liverpool John Moores University, UK.

Abstract

Physiological Computing describes a category of human-computer interaction where physiological data from the brain and body are transformed into input control to inform software adaptation. These physiological data are used to provide a dynamic representation of the user with respect to extending the body schema (sense of personal agency) and the body image (perception of cognitions, emotion, motivation etc.). A dynamic user representation can be used as a form of input control and to guide a process of intelligent adaptation. This chapter will provide a historical perspective on the concept and describe the cybernetic logic of the closed-loop technology. Important aspects of the adaptive loop will be described such as wearable sensors, the process of real-time classification and how a process of monitoring, analysis and adaptation may enhance human-computer interaction.

Keywords: Physiological Computing, Psychophysiology, Neuroergonomics, Intelligent Adaptation

1. INTRODUCTION

Physiological computing is characterised by a live connection between technology and the human nervous system. This act of monitoring renders the machine privy to a variety of data: electrochemical activity from the epidermis, fluctuations in muscular tension, the haemodynamics of the cardiovascular system and the electrocortical fluctuations of the brain. The connection between person and technology corresponds to an act of *digital embodiment*. By connecting to a computer, the human extends the boundaries of the central nervous system, communicating directly with technology via those physiological processes that underpin thoughts, emotions and actions.

Digital embodiment has a number of implications for the way in which we interact with technology. Conventional human-computer interaction is asymmetrical with respect to the flow of information. The user can interrogate a huge range of data concerning the internal processes within the computer (e.g. RAM use, disk space etc.) whilst the computer remains essentially blind to the psychological intentions and experience of the user. Continuous monitoring of the central nervous system is one way to facilitate a form of symmetrical HCI where information flows simultaneously from computer to user and vice versa (Hettinger, Branco, Encarnaco, & Bonato, 2003). The implications of this innovation are potentially profound. For example, ‘smart technology’ demonstrates a degree of intelligence by exhibiting sensitivity to task context and user intention without explicit information (Norman, 2007). Monitoring physiological data is a means of allowing a computer system to become aware of the user as a dynamic entity with the major advantage of being continuously available, even in the absence of any overt forms of behavior (Byrne & Parasuraman, 1996).

1.1 Categories of Physiological Computing

Physiological computing systems fall into one of two broad categories. The first are designed to extend the body schema, i.e., the system of sensory-motor functions that we use every time we tap a key or move a joystick. These functions are guided by a sense of agency, i.e., I am the one doing this. For example, the Brain-Computer Interface (BCI) offers an alternative mode of input control to extend the body schema (Allison, Wolpaw, & Wolpaw, 2007). BCIs capture electrocortical activity at source (e.g., the intention that precedes movement or selection) and offer a highly novel form of hands-free interaction that is capable of communicating with standard screen-based technologies as well as specialised devices such as prostheses. The same logic can be extended to monitoring muscle activity via electromyography (EMG). An EMG sensor on the forearm can detect patterns of gestures (Zhang et al., 2011) in a ubiquitous computing scenario. Similarly, muscle activity in the form of eye movements can be used for cursor control and other forms of input (Majaranta & Bulling, 2014).

The second category of physiological computing is relevant to perceptions of internal states related to psychological states. The body image has been defined as “a complex set of intentional states and dispositions… in which the intentional object is one’s own body” (Gallagher, 2005) (p.24. Physiological computing systems augment the body image by monitoring and responding in an adaptive fashion to spontaneous data originating from psycho-physiological interaction in the central nervous system. Biocybernetic Adaptation covers a range of systems designed to capture psychological states relating to performance and wellbeing (Allanson & Fairclough, 2004; Fairclough, 2009). These states include psychophysiological signatures of emotions, such as anger, frustration or fear, or changes in cognitive activity related to mental workload. For certain categories of software, such as games or auto-tutoring systems, we may be interested in changes that reflect elements of both cognition and emotion, i.e., when someone is mentally overloaded (too much information, not enough time), they also may experience anxiety or anger. This 'wiretapping' approach shares a number of features with biofeedback technology. Both the facilitate the process of self-regulation by providing feedback of physiological activity that can occur outside of conscious awareness. By interacting with data from the brain and body at an interface with technology, users can gain insight into states of cognition and emotion that can facilitate self-knowledge and associated strategies of self-regulation. There are two important facets of biocybernetic adaptation that distinguish this category of control from those BCI systems designed to extend the body schema. The system is designed to adapt to *spontaneous* changes in the psychological state of the user. If the person is frustrated, the software may offer help; if the user is overloaded, the software may filter the flow of incoming information. This is an *implicit* mode of HCI with no requirement for the user to exhibit intentional behaviour, which has also been termed passive BCI (Zander & Kothe, 2011).

Physiological computing systems are designed to extend the body schema or the body image by creating a dynamic quantification of those entities that may be accessed by a technological device. It is the quality of this dynamic quantification that will determine the efficacy of the interaction between user and system.

1.2 The Biocybernetic Loop

The biocybernetic loop (Pope, Bogart, & Bartolome, 1995) serves as a unifying concept for all physiological computing systems. This concept is derived from the cybernetic model of control and communication within a closed loop (Wiener, 1948).

There are three generic stages of data processing within this feedback loop: collection, analysis and translation. The first stage describes the collection of physiological data via sensor apparatus. A user at a desktop computer may wear a simple device to access heart rate or skin conductance level. The design of wearable, ambulatory sensors for data collection is a vital component of the data collection phase. The second stage of data analysis receives filtered data as an input and both quantifies the data in an appropriate way and identifies/corrects for the presence of artifacts. The analysis algorithm should be capable both quantifying incoming data in real-time and identifying those periods that include ‘bad’ data. It would be ideal if the analysis algorithm were capable of not only identifying sections of ‘bad’ data but also subjecting these periods to a correction algorithm, in order to preserve the integrity of the data stream. The analysis stage should yield an appropriate and accurate quantification of physiological data but this is a very loaded phrase; much depends on what particular aspect of psychology or behaviour is the target or focus of the biocybernetic loop. The final process of the loop is translation. This stage describes how physiological units of measurement are converted into a computer command to be executed at the human-computer interface.

The three stages are realised in different ways depending on the category of physiological computing system. For EMG-based interfaces and some categories of BCI (e.g. those where motor functions are captured at the cortex), the function of the biocybernetic loop is to translate patterns of physiological activity into a specific command. This act of translation may be representative and functionally equivalent in some cases. A system that translates vertical and horizontal eye movement (monitored via EOG) into vertical and horizontal movements of a cursor on a screen is characterised by one-to-one correspondence as eye movements are scaled into x and y coordinates on the screen.

Other categories of physiological computing depend on the accurate identification of spontaneous psychological states to inform system adaptation. The obvious examples fall into the category of biocybernetic adaptation, such as affective computing technologies designed to capture changes in emotional states. In these cases, the link between physiological activity and psychological processes is analogous rather than strictly representative.

The translation of real-time physiological data into computer control is achieved by a component called the adaptive controller. This is an element within the biocybernetic loop that incorporates the translational rules of the system, e.g., IF heart rate shows a rise of 30% THEN offer help, IF P300 amplitude is maximum for the 'delete' command THEN activate delete function. For pattern-matching algorithms that translate physiology into input control at the interface, adaptive control is relatively straightforward. Detection of the upward movement of gaze can be translated into vertical movement of the cursor with an emphasis on low-level dynamics, such as the gain between EOG activity and sensitivity of cursor response. Controlling the movement of an avatar via BCI would require the adaptive controller to recognise and respond to template patterns of EEG activity that represent left/right and forward/backwards. For those physiological computing systems that extend the body image, the purpose of the adaptive control is to translate physiological activity into an efficient (in terms of the rate of information transfer) and responsive mode of input control.

The controller serves a different function in the case of biocybernetic adaptation. These systems are designed to promote positive states and to prevent/ameliorate undesirable ones. Biocybernetic software serves as a dynamic mechanism for software to promote the same design goals. This is a very disruptive concept with respect to how people currently interact with technology. There is a shift of influence between user and system because biocybernetic control is designed to shape and manipulate the psychological state of the user. If an operator experiences a high level of mental workload, the system will intervene to reduce workload and preserve safety. If the player of a computer game is frustrated by the experience of repeated failure, the software may adapt to reduce challenge or offer help. The benign nature of these adaptations should not mask the fact that a working biocybernetic system is a machine with a prescribed agenda, a machine that deploys real-time adaptation as a means to achieve certain design goals, e.g., to prevent accidents, to promote positive affect, etc. This autonomy is achieved by incorporating a dynamic representation of the user, and by default the actual user, as an element in the control loop. The net result of the closed-loop design is an inevitable shift within the human-computer dyad towards the computer as a co-worker or team-player who actively is aware of the goals of the user as opposed to the dumb slave system that I'm using to type these words.

The biocybernetic loop may function at different levels of the human-computer interaction. With respect to muscle interfaces and BCI, the biocybernetic loop functions within the HCI dyad and is designed to explicitly communicate commands to the interface. The loop mediates the intentions of the user to move the cursor down or to make the avatar move forward. Biocybernetic adaptation tends to function at the meta-level of the HCI, adjusting the parameters of the interaction (e.g., altering game difficulty) or making dynamic interventions (e.g., offering help, activating automation). This type of adaptation may also be achieved without any conscious intention on the part of the user. It is even possible for biocybernetic adaptation to occur in ways that are sufficiently subtle to escape the conscious perception of the user. The purpose of the biocybernetic loop is to adjust settings and make interventions in order to shape the interaction in a desirable way.

The biocybernetic loop functions both as a model for information flow and a unifying concept for physiological computing systems. It encapsulates the basic properties of sensor design and signal processing that underpin all categories of system.

2. MEASUREMENT & CLASSIFICATION

The biocybernetic loop is the fundamental control unit of all physiological computing systems. The results of this data processing 'pipeline' inform the mechanism of software adaptation. The loop is responsible for the interpretation of raw physiological data into a coherent response from computer software. The process encapsulated with the biocybernetic loop is associated with a string of important caveats:

1. physiological measures must be a valid measures of psychological concepts
2. unobtrusive hardware must exist that is capable of capturing these measures in the field with sufficient fidelity
3. data must be analysed and categorised in near-real time in order to deliver a representation of the user to the system
4. changes in user representation must be translated into software control and adaptation that is both responsive and coherent.

These issues may be studied in isolation from one another (and often are), but from the perspective of an integrated system development, each part of the loop should be considered as mutually dependent on the others.

2.1 Inferring psychological meaning from physiological signals

The goal of measurement is the inference of the psychological or behavioural state of the user based on patterns of psychophysiological activity. One challenge for the designer of a physiological computing system is the identification of a distinct physiological pattern that is consistently associated with a target psychological concept. These patterns are used to trigger events at the interface and they must be salient, unique and reliably detected in real-world conditions.

One way to increase the consistency of the psychophysiological inference is to capture the responses to a stimulus with known properties. An array of letters where each item flashes in sequential fashion can be presented in order to capture both the presence and magnitude of an evoked response potential (ERP) from EEG activity. It is known that the magnitude of certain components of the ERP (e.g., P300) respond to attentional processing, hence the letter one wishes to select from this array will generally deliver a greater magnitude of response. This "probe" strategy employs temporal coincidence in order to link a specific physiological pattern with a particular stimulus event. By contrast, the 'wiretapping' approach that characterises biocybernetic adaptation seeks to capture psychophysiological "signatures" of emotions and cognitive states against a background of spontaneous activity. This type of system must work with a low signal-to-noise ratio as the size of physiological response to an emotional event is relatively small compared to the magnitude of change due to physical movements and other confounding factors. It is a mistake for any designer to assume that physiological measures are capable of directly capturing psychological states in a 'plug-and-play' fashion.

For those systems designed to extend the body schema, the accurate inference of intention is the critical issue. The 'wiretapping' systems that fall into a broad category of biocybernetic adaptation are designed to classify spontaneous activity into pertinent categories that capture 'target state' of user psychology. The process of inferring psychological events from physiological measures bears careful consideration in both cases. A penetrating analysis was provided by (Cacioppo & Tassinary, 1990) who scruntinised the specificity of association between concept and measure. The strongest category of psychophysiological inference for the development of a physiological computing system is the ‘one-to-one’ relationship where a particular physiological measure operates as a unique marker of a specific psychological construct across all contexts of measurement. This kind of relationship is optimal but is also the rarest level of psycho-physiological inference, particularly in the context of real-world testing.

Inference from psychophysiological measures can be a messy and inconclusive business, particularly in the context of everyday use outside of a laboratory. The most important starting point for the designer is a concrete understanding of the classification scheme that the act of measurement must deliver in order to make the system work, i.e., how many categories must be distinguished for the system to work (Fairclough & Gilleade, 2012). In all cases of designing and developing physiological computing systems, it is important to select sensors/measures to the adaptive or input capabilities of the system. A concrete notion of what the measures are meant to achieve, in what environment and with whom provides an essential context within which to select, test and validate the psycho-physiological inference.

2.2 Wearable sensors

The challenge for the design of sensors for physiological computing system is to create wearable sensors that maximise comfort, minimise intrusion and may be used in a public space without any embarrassment or self-consciousness – whilst maintaining a high fidelity of signal quality.

When designing sensors with low intrusiveness, there is a temptation to strip down the process of measurement. This simple strategy equates the number of sensors or measures directly to the comfort of the user (fewer sensors = greater comfort for the user) but this is a myopic approach. Ambulatory sensors are designed to be worn in the field, hence they must capture not only the signal of interest but also as many potential sources of noise as possible, in order to facilitate a process of artifact detection and correction. In the field, it is important to capture metabolic changes due to everyday activities using pedometers or accelerometers to quantify movement and to incorporate these variables into the diagnosis produced by the system. We may look to other data sources to provide additional context for our interpretation of spontaneous physiological activity, such as time of day or room temperature or background noise level. From the perspective of the user experience, the level of comfort associated with data capture is an overriding concern. As system designers, we need to move towards an invisible monitoring process wherein wearable sensors and the process of data capture, storage and analysis are rendered as unobtrusive as possible.

Remote sensing offers the best chance to achieve an “invisible” process of physiological monitoring (Poh, McDuff, & Picard, 2011), but like all camera-based systems, the sensor requires a stationary user in order to work. Wearable devices, such as chest or wrist straps or earplugs, have the advantage of being ambulatory in the sense that the user can move around, albeit with the disadvantage that sensors are relatively intrusive as they involve contact with the skin.

The provision of sensor apparatus capable of comfort and signal fidelity is an essential development if we are to realise the potential of physiological computing systems. If these devices are not available, physiological computing will never reach the vast majority of users. It is also important for sensor technology to come equipped with Software Development Kit (SDK) and the capability to use standard protocols for communication, such as Bluetooth. There is enormous potential for wearable sensors to interface with mobile devices via the specialised software niche of the app economy if the right combination of hardware and software tools can be brought to the market.

2.3 Signal analysis and classification

Once signals have been captured and filtered by the system, these data are subjected to a process designed to identify and classify significant patterns in the data. This combination of signal processing and diagnosis represents the operational crux of the biocybernetic loop whereby unique patterns of physiological activity are associated with psychological events.

There are several approaches to signal classification, which are applied to different types of categorisation 'problems' in physiological computing systems. The identification of a spontaneous psychological state falls in the domain of affective computing, where 'target states' such as frustration or excitement are operationalised as a pattern of physiological activity distinguishable from spontaneous responses that are associated with other states. The successful operation of a BCI requires accurate identification of physiological features associated with the initiation of input control, such as activity in the somatosensory cortex or a positive deflection of electrical activity that occurs in the same temporal window as a particular stimulus. Signal classification must be both fast and accurate to facilitate real-time input control. The classification of physiological signals along a unidimensional continuum, such as anxiety or mental workload, represents a different category of assessment where estimates of magnitude (low, medium, high) on a unidimensional scale are the focus of signal classification.

The purpose of signal classification is to create a literal interface between the human nervous system and a repertoire of software responses. But the data processing "pipeline" that underpins this interface must be carefully constructed in a bottom-up fashion. To use an analogy, if we were to design a system for language translation, the first design question would concern the size of the vocabulary; in other words, exactly how many different words must be recognised and translated by the system in order to function with a degree of utility? This is also a good starting point for the process of data classification in a physiological computing system. All systems are associated with a *functional vocabulary* that describes how many commands, subjective states or gradations of experience must be recognised in order to support successful operation. A system may function adequately on the basis of simple binary differentiation of two classes. Other systems may be equipped with repertoires of five or more adaptive responses, amplifying the challenge of accurate data classification. The number of items in the functional vocabulary defines the boundaries of the data classification problem inherent in the design of physiological computing systems, e.g., 2 classes or 5 classes or 10 classes. Without this kind of operational context, the specification of an optimal process of signal classification remains an abstract proposition.

The degree of similarity or overlap between target states is also important for classification accuracy. Making a two-category distinction between happiness and anger would generally yield more accurate classification than an attempt to differentiate fear from anger. This is logical because similar emotional states contain greater overlap in terms of how they impact on psychophysiological reactivity.

The application of machine learning algorithms is a common strategy for classification within the biocybernetic loop. The general methodology for the construction of a classifier is to generate a training set, which accurately represents the dimensional space associated with the functional vocabulary of the system. This database will subsequently be used to train a classifier and represents the template for all subsequent acts of categorisation conducted by the system, therefore it must provide a stable and well-defined mapping of physiological measures onto psychophysiological space. The first obstacle for classification is deriving and defining an optimal set of training data for the machine-learning algorithm. Data from the brain and body, particularly in the field, has poor signal-to-noise ratio, i.e., contaminated by physical artifacts that may account for outliers in each “class” to be identified.

Signal classification is generally based around the creation of features (or vectors) that are derived from multiple measures of psychophysiology. There is a good deal of research literature where single data streams are analysed in myriad ways; for example, measures of heart rate can be expressed in terms of descriptive statistics (mean, maximum, minimum, standard deviation) or subjected to further analysis to yield power in low and high bands, then further expressed as the ratio of both bands. The tendency to measure lots of variants from the same basic signal source creates the so-called ‘Curse of Dimensionality’ where the amount of data required to describe different classes increases exponentially with the dimensionality of those features that are used as input. The practical implication of this ‘curse’ is that the designer must acquire more training data when he or she adds new features as inputs to the classification process.

The act of classification may subsume both the process of discrimination and the implicit mapping onto psychological categories. It is hoped that the training set provides a good mapping in terms of a quantitative discrimination, but whether it provides the best possible mapping remains an open question. If we consider which factors may contribute to classification errors, three main sources are most likely (Lotte, Congedo, Lecuyer, Lamarche, & Arnaldi, 2007):

1. influence of noise from non-psychological sources, as stated earlier, noise is a ‘fact of life’ as far as ambulatory psychophysiology is concerned

2. degree of divergence between the estimated mapping provided by the training set and the best mapping possible. This second factor is determined purely by the representativeness of the training set, which is defined as the capacity of the training set to generalise across all instances of the target state or pattern. For instance, is the pattern of psychophysiological reactivity associated with a specific target state during the training set sensitive to all other instances of the target state that may be encountered by the system. The degree of this divergence between what is measured by the system now and what was measured during training is called Bias.
3. degree of sensitivity of the classifier to the training set. It is important to note that different approaches to signal classification differ with respect to their susceptibility to specific and idiosyncratic qualities of the training dataset. The degree of sensitivity to the training set exhibited by the classifier is called Variance.

This summary provides an overview of the challenges facing the design of a classification system using live data from brain or body as an input and producing real-time or near-real-time outputs. Noise is an irreducible component of the measurement process and may be dealt with by filtering the signal. If bias and variance are both low, we would expect classification errors to be minimal because the mapping is good and sensitivity to the training set is low. Obviously, a poor mapping can originate from several sources, such as multiple categories that are not clearly discriminated or a substantial overlap between the psychophysiological ‘signatures’ of different target states. If the training set is situation-specific, then classification errors will rise due to increased variance.

A good training set represents an essential prerequisite to enhance the accuracy of a classification algorithm. “Good” training data may be defined according to a number of different criteria. It is important to choose a psychophysiological measure (or collection of measures) that is sensitive to changes in psychological state or an intentional act. This measure must be sufficiently robust to exhibit sensitivity under the working conditions of the system. A training set that captures the range of physiological responses under realistic conditions will reduce the level of bias and variance in the classification system. Ecological validity is one key aspect of the training set that captures the ability of data to represent psychophysiological measurement as it would occur under real-world conditions.

A number of classification techniques have been utilised in the context of both BCI applications and the biocybernetic loop. The first category of classifiers may be described as *generative* and are designed to compute the likelihood of each class. This generative class includes Bayesian network approaches, which is a probabilistic model that uses a "maximum a posteriori" rule to assign a vector to the class with the highest posteriori probability. *Static* classifiers, such as artificial neural nets (ANN), represent a sophisticated approach to network-based classification. ANNs are arranged in layers (e.g., multilayered perceptron) where each node or neuron receives a number of inputs in order to calculate the cumulative activation of that particular neuron, this output is relayed to the next layer of the network and so on. One disadvantage of the ANN approach is that the network is very sensitive to overtraining, particularly with noisy psychophysiological data as a set of inputs. The static approach to classification is contrasted with *dynamic* techniques, such as Hidden Markov Models (HMM), which incorporate temporal features into the process of discrimination. HMM have evolved to calculate the probability of observing a particular sequence of feature vectors; this approach has been adopted in some BCI systems but is rarely used for biocybernetic adaptation.

Nearest-neighbour approaches fall into the category of *discriminative* classifiers where the distance (e.g. Euclidean) between each input and a feature vector is calculated based upon the training space. Support Vector Machines (SVMs) represent another form of discriminative classifier and has been widely used in both BCI and biocybernetic systems. SVMs are capable of creating nonlinear decision boundaries and whilst they may be slow (computationally), they have good generalisation and tend to be less sensitive to overtraining as well as the curse of dimensionality. *Stable* classifiers, such as Linear Discriminant Analysis (LDA), are characterised by their relative simplicity and insensitivity to small variations in the training data. This is an advantage in the sense that the classification system is sensitive to gross rather nuanced trends in the training data; however, as the name suggests, LDA performs poorly where the classification is based on complex (i.e., nonlinear) boundaries in contrast to SVMs or ANNs.

A number of reviews (Lotte et al., 2007; Novak, Mihelj, & Munih, 2012) have surveyed the prevalence of different classification techniques in the development of different categories of physiological computing systems. With respect to synchronous BCI, classification is mainly characterised by the use of SVMs, dynamic classification (HMM) and ensemble classification, particularly majority voting and boosting. The application of ANN to classification in BCI accounted for approximately a quarter of those systems in the review. Those systems designed for biocybernetic adaptation were characterised by a mixture of SVM, LDA and Classification Trees, although a small number did use an ANN approach.

The process of measurement and classification is fundamental to the integrity of the biocybernetic loop. The challenge for designers is to create dynamic user representations that: (a) are scientifically sound, (b) rely on comfortable and non-intrusive sensors that are capable of delivering good signal quality in everyday settings, and (c) can be classified into different categories within an adaptive controller in order to drive adaptation at the interface.

3. INTELLIGENT ADAPTATION

The prioritisation of current research on signal analysis and classification is understandable because the system cannot work without inputs and the act of classification is inextricably bound up with the quality of signal input. But the designer must make equally important decisions about how those categories are translated into a repertoire of software responses because action at the interface will ultimately determine user experience.

We must acknowledge the enormous scope for error that exists within the biocybernetic loop when data are collected in the field, or to phrase the statement in more precise terms, there is enormous scope for user perception of system error. A misclassification within the context of BCI interaction yields a response that is unintentional and obvious to the user. Spotting an error during interaction with a biocybernetic system is significantly more difficult for the user. Biocybernetic adaptation faces the significant hurdle of creating adaptation at the interface that resonates with the dynamic experience of the user. For systems designed to provide input control, the primary design issue is how to match the intentions with events at the interface within a small time window. For biocybernetic systems, the capacity of the adaptive loop to synchronise with user experience is complicated by the complexity, spontaneity and subjectivity of the latter.

The quality of adaptive response may be characterised by its accuracy, sensitivity and intuitiveness from the perspective of the user. The *accuracy* of a physiological computing system is an obvious starting point for this discussion. It is reasonable to assume that an accurate response is desirable from the perspective of the user, but what does that mean? Accuracy is defined as matching a predefined pattern of physiological response with a specific command from the functional vocabulary. If the system selects a response based on an erroneous act of classification, it is an error of commission and will be perceived as invalid by the user. However, inaccurate systems are also guilty of errors of omission - when the user expects a response from the system but does not receive one. It is important that both categories are minimised, errors do not just annoy and inconvenience the user but fundamentally undermine the development of trust in the technology.

Consider the relationship between the size of the functional vocabulary (the repertoire of possible system responses) and the criterion of accuracy. It could be argued that increasing functional vocabulary is beneficial for user experience, making a greater range of adaptive options available to the system should lead to a nuanced response at the interface that increases the perceived “intelligence” of the system as a whole. However, probability dictates that the errors of commission become increasingly frequent as items are added to the functional vocabulary and the designer of the system faces a trade-off between accuracy and adaptive capacity. In order to resolve this dilemma, the designer ought to consider the minimum number of adaptive responses necessary for the system to meet its operational requirements. It is reasonable to assume that accuracy will be maximised for a system with minimal functionality. In some cases, a physiological computing system may be capable of operating quite effectively with only two or three types of adaptive responses; much depends on the type of interactive experience that the system is designed to deliver.

If accuracy is concerned with the design of adaptive systems where responses are tailored to the limits of psychophysiological classification, the *sensitivity* of the system response describes the temporal relationship between physiological activity and events at the interface. Many users develop an intuitive heuristic through interaction with input control devices so that x amount of mouse movement is equated with y amount of cursor control on the screen. The same logic applies to BCI and eye control of cursors. The temporal relationship between changes in psychophysiological activity and events at the interface is described as gradation sensitivity. There are two aspects of this relationship: proportion and dynamics. In its simplest form, a small change in electrocortical activity will cause a small movement in a desired direction, e.g., an increase of EEG activity by one standard unit at a sensorimotor site causes an avatar to move forward 1m in virtual space, and larger changes in the EEG cause greater movement and so on. However, the ratio between signal and its output at the interface may be designed in different proportions, hence a 2:1 relationship may mean that avatar moves forward 1m in virtual space whenever sensorimotor activity increases by two standard units.

Gradation sensitivity can be weighted in different ways in order to create the largest amount of change at the interface for different levels of psychophysiological activity. The relationship between psychophysiological change and events at the interface may be conceived as an extension to Fitts Law (Fitts, 1954) where covert responses from the central nervous system are mapped onto the biocybernetic control loop. For the designer, creative adjustment of gradation sensitivity represents one mechanism to reinforce physiological self-regulation in order to improve the productivity and quality of the human-computer interaction.

The temporal dynamics between psychophysiology and software are perhaps less important in the case of biocybernetic adaptation. The sensitivity of biocybernetic adaptation is based mainly upon a perception of whether the system response is appropriate to a specific situation. The first criterion concerns the ability of the system to deliver the right response. If the user is extremely frustrated or very bored, there is an inherent expectation that a physiological computing system will intervene. If no response is forthcoming, the user perceives an error of omission. If the system makes an inappropriate response, e.g., offering help when the user is calm and relaxed, the insensitivity of the system revolves around errors of commission. Both categories of error create unique forms of insensitivity, both of which are equally damaging to the user experience.

The perceived accuracy from the perspective of the user is an important determinant of the quality of the interaction. A system that is perceived to be accurate in the short-term will create a positive impression that encourages further use. However, perceived accuracy and trust can also be affected by bias or be linked to the cost of errors to the user, i.e., errors may be more or less costly depending on the degree of annoyance or inconvenience experienced by the user as a direct result. The question of what is an acceptable level of accuracy for a physiological computing system has been addressed by several studies. Some simulated various levels of accuracy with respect to control of an input device and task difficulty respectively in order to explore levels of user acceptance and tolerance for system error (Novak, Nagle, & Riener, 2014; Van de Laar, Bos Plass-Oude, Reuderink, Poels, & Nijholt, 2013). One recent study assessed the ability of users to assess the classification accuracy of a physiological computing system designed to classify interest levels (high vs. low) in a series of movie trailers (Fairclough, Karran, & Gilleade, 2015). The authors reported that participants tended to over-estimate a mathematical accuracy score of 0.82-0.91 by approximately 5% - hence participants may be able to accurately assess classification accuracy of a system if they are provided with overt feedback (as they were in this study).

One design option to deal with the inherent uncertainty of matching adaptive responses to dynamic user states is to vary the autonomy of the system. It is generally assumed that triggers from the biocybernetic loop will action events at the interface in an all-or-nothing fashion without any input from the user. But total automation is not the only option available to the design of a physiological computing system. The Levels of Automation analysis (LOA) (Sheridan & Parasuraman, 2006) represents system automation as existing on a continuum where responses vary from 100% manual to 100% automated. There are a number of hybrid forms on this continuum, such as automated prompts (e.g., would you like the game to increase difficulty now?) and automated cues (e.g., highlighting appropriate functions or areas of the screen). Adaptive technology based on physiological computing could adjust the autonomy of responses at the interface based upon the level of confidence underpinning the episode of classification. If the system is highly confident, the response occurs at the interface without any consultation with the user. If confidence is not as high, the system may prompt the user for confirmation before taking action. Whilst biocybernetic adaptation is based upon correct classification, it is obvious that instances of greater physiological reactivity are easier to recognise than smaller ones. The physiological computing system would always identify a transition from low to high frustration but a change from low to medium frustration is harder to detect. Therefore, the system may again adopt a softer approach to automation, using prompts or cues, when the distance between categories is low and ambiguity of classification is high. This design option provides some insurance against both errors of commission and omission because the system makes a definite response but requires clarification from the user to resolve any uncertainty.

When a person interacts with a physiological computing system, the technology actions a series of adaptations in response to ‘live’ changes in brain activity or psychophysiology.

Intuition describes the intelligibility of those adaptations from the perspective of the user. This particular criterion focuses on the perceived meaning of adaptation at the interface as opposed to the process of measurement and classification supporting those adaptations. There is an analogy between adaptive interaction and a spoken conversation. Both activities are characterised in terms of turn-taking and mutual adaptation. Conversation is initiated in the biocybernetic loop when the user exhibits a particular pattern of psychophysiological activity that triggers an adaptive response. The interpretation of the system is communicated implicitly to the user via the specifics of adaptation at the interface: the presentation of calming music tacitly communicates an interpretation of stress or anxiety on behalf of the system, an assessment of the intention to move forward produces forward movement of an avatar in virtual space. This intuitive element informs users' understanding of how the system functions as an active agent working to a predetermined rationale. In order to engender user trust, it is important that the response from the system is understood with sufficient clarity to both inform and shape the human-computer interaction (Miller, 2005).

From the perspective of the user, intuition represents the extent to which an adaptation reflects feelings, thoughts and intentions that are active in consciousness at that particular time. The informative content of system adaptation may simply reinforce what is already known, when an angry user experiences a calming response from the system, feedback is simply confirming the contents of awareness. Alternatively, an adaptive response may grant insight into unacknowledged feelings or thoughts and the presentation of a calming response can increase awareness of increased frustration (Picard et al., 2004). A second-order aspect of intuition concerns users' understanding or assessment of the rationale underlying the adaptive response from the system. The detection of frustration may prompt a number of adaptive responses: offer help, suggest a rest break, play calming music, recommend a number of breathing exercises. Each adaptive response is unique but the underlying rationale is consistent and it is important for this rationale (e.g., to counteract frustration) to be clearly communicated to the user. Systems associated with biocybernetic adaptation are designed with a distinct rationale in mind: to reduce mental workload, to promote safety, to help the user, to challenge the player, etc. An understanding of this rationale will hopefully improve the intelligibility of system behaviour.

The criterion of intuitiveness also concerns the quality of the adaptive response as assessed by the user. If the system makes a response that accurately reflects the current state of the user, correct classification will count for little if the actual adaptation has no utility from the perspective of the user. This aspect of intuition concerns the extent to which an adaptation meets the needs and desires of the user. If the user is frustrated because they are behind schedule, an automated request to take a rest break is unlikely to be welcomed or be perceived as particularly helpful. There is a degree of overlap between these aspects as an intuitive system response is accurate, comprehensible, timely and useful from the perspective of the user.

The informative content of the adaptive response provides an explicit cue to the diagnosis of the user state. If the response from the system is opaque, the user may assess the system to be inert or unintelligible, both of which are undesirable. The informative element of adaptation is self-evident for input control systems where intention and feedback at the interface are closely coupled. For biocybernetic adaptation, the informative content of the adaptive response may be explicit and obvious to the user or implicit in the sense that the user may or may not notice a change at the interface. Explicit feedback concerns those categories of the adaptive response that are impossible for the user to miss: the appearance of an avatar offering help, an on-screen recommendation to take a break. This feedback represents an unambiguous statement of the design rationale underpinning the biocybernetic loop. There is a risk associated with explicit adaptation that errors of commission are obvious to the user and only a small number of high-profile errors are sufficient to damage trust in the system, but there are occasions, typically during extreme frustration or high mental workload, when the user might reasonably expect an explicit response from the system. Implicit adaptation may not be noticed by the user at the interface and have the advantage of making errors of commission without necessarily damaging trust because the user may not notice the error. The problem of implicit adaptation is that they may often exert a cumulative influence that takes time to achieve any tangible effect on the user.

For example, we constructed a biocybernetic version of the game Tetris (Fairclough & Gilleade, 2012) where the drop speed of the blocks were adapted to changes in real-time EEG using tiny adjustments that were barely noticeable to the player. The experiential effect of these implicit adaptations were double-edged; it took a period of several minutes for a clear trend of adjustment (to increase or decrease speed) to become apparent to the user and in the meantime the users generally perceived the system to be inert and creating errors of omission. The strengths and weaknesses of explicit/implicit adaptation may be designed into the interaction by using implicit adaptations for small magnitudes of change in the state of the user and relying on explicit adaptation when extremes are detected. See Fairclough (2009) for more detail on this topic.

The repertoire of adaptive responses available to the system designer is similarly finite but the process of design is crucial. The designer must decide which adaptive options to include within the functional vocabulary of the system and which to omit. If a specific adaptation is deemed worthy of inclusion, the magnitude of response or the levels of response must be clearly defined. The whole purpose of this process is the production of a system capable of intelligible interaction.

4. SUMMARY

Interaction with a physiological computing system represents one approach to the creation of a technology where control is achieved without touch and software responds to the psychological context of the user. The closed-loop logic of these systems describes how raw physiological data from the body and brain is translated into a series of dynamic control inputs and changes at the interface, which are conveyed directly to the user. This process of translation from raw physiology to input control contains a number of steps with significant hurdles, such as: the design of wearable sensors that deliver high quality data in an unobtrusive way, the process of inferring psychological states from physiological data in everyday life, the detection of artifacts and classification of data in real-time. These challenges of measurement and signal processing in this field are substantial but the design of the adaptive controller is central to the user experience. The adaptive controller represents the rationale of the closed-loop, which describes the way in which data is translated into adaptations and responses at the interface with the user. This component remains relatively unexplored compared to signal processing and classification, but it is the efficacy of the adaptive controller that will largely determine the user experience and the degree of 'intelligence' displayed by the system.

REFERENCES

Allanson, J., & Fairclough, S. H. (2004). A research agenda for physiological computing. *Interacting With Computers, 16*, 857-878.

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.

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

Cacioppo, J. T., & Tassinary, L. G. (1990). Inferring psychological significance from physiological signals. *American Psychologist, 45*(1), 16-28.

Fairclough, S. H. (2009). Fundamentals of Physiological Computing. *Interacting With Computers, 21*, 133-145.

Fairclough, S. H., & Gilleade, K. (2012). *Construction of the biocybernetic loop: a case study*. Paper presented at the 14th ACM International Conference on Multimodal Interaction, Santa Monica.

Fairclough, S. H., Karran, A., & Gilleade, K. (2015). *Classification accuracy from the perspective of the user: real-time interaction with physiological computing*. Paper presented at the CHI '15 Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, Seoul.

Fitts, P. M. (1954). The information capacity of the human motor system in controlling the amplitude of movement. *Journal of Experimental Psychology, 47*(6), 381-391. doi: 10.1037/h0055392

Gallagher, S. (2005). *How The Body Shapes The Mind*: Oxford University Press.

Hettinger, L. J., Branco, P., Encarnaco, L. M., & Bonato, P. (2003). Neuroadaptive technologies: applying neuroergonomics to the design of advanced interfaces. *Theoretical Issues in Ergonomic Science, 4*(1-2), 220-237.

Lotte, F., Congedo, M., Lecuyer, A., Lamarche, F., & Arnaldi, B. (2007). A review of classification algorithms for EEG-based Brain-Computer Interfaces. *Journal of Neural Engineering*, *4*, 1-24

Majaranta, P., & Bulling, A. (2014). Eye tracking and eye-based human-computer interaction. In S. H. Fairclough & K. Gilleade (Eds.), *Advances In Physiological Computing* (pp. 39-66): Springer.

Miller, C. A. (2005, July 22nd-27th). *Trust in adaptive automation: the role of etiquette in tuning trust via analogic and affective methods.* Paper presented at the First International Conference on Augmented Cognition, Las Vegas, NV.

Norman, D. A. (2007). *The Design of Future Things*. New York: Basic Books.

Novak, D., Mihelj, M., & Munih, M. (2012). A survey of methods for data fusion and system adaptation using autonomic nervous system responses in physiological computing. *Interacting With Computers, 25*, 154-172.

Novak, D., Nagle, A., & Riener, R. (2014). Linking recognition accuracy and user experience in an affective feedback loop. *IEEE Transactions of Affective Computing, 5*(2), 168-172.

Picard, R. W., Papert, S., Bender, W., Blumberg, B., Breazeal, C., Cavallo, D., . . . Strohecker, C. (2004). Affective learning - a manifesto. *BT Technology Journal, 22*(4), 253-269.

Poh, M.-Z., McDuff, D., & Picard, R. (2011). *A medical mirror for non-contact health monitoring*. Paper presented at the ACM SIGGRAPH Emerging Technologies, Vancouver.

Pope, A. T., Bogart, E. H., & Bartolome, D. S. (1995). Biocybernetic system evaluates indices of operator engagement in automated task. *Biological Psychology, 40*, 187-195.

Sheridan, T. B., & Parasuraman, R. (2006). Human-automation interaction. *Reviews of Human Factors and Ergonomics, 1*, 89-129.

Van de Laar, B., Bos Plass-Oude, D., Reuderink, B., Poels, M., & Nijholt, A. (2013). How much control is enough? Influence of unreliable input on user experience. *IEEE Transactions on Cybernetics, 43*(6), 1584-1592.

Wiener, N. (1948). *Cybernetics: Control and communication in the animal and the machine* (2 ed.). Cambridge, Mass.: M.I.T. Press.

Zander, T. O., & Kothe, C. (2011). Towards passive brain-computer interfaces: applying brain-computer interface technology to human-machine systems in general. *Journal of Neural Engineering, 8*, 1-5.

Zhang, X., Chen, X., Li, Y., Lantz, V., Wang, K., & Yang, J. (2011). A framework for hand gesture recognition based on accelerometer and EMG sensors. *IEEE Transactions on Systems, Man and Cybernetics - Part A: Systems and Humans, 41*(1064-1074).