Brain-Computer Interfaces:
Applying our Minds to Human-Computer Interaction

Desney S Tan[†]

Microsoft Research

One Microsoft Way, Redmond, WA 98052, USA

desney@microsoft.com

ABSTRACT

Recent advances in cognitive neuroscience and brain imaging technologies provide us with the unprecedented ability to interface directly with activity in the human brain. Beyond traditional neuroscience work, researchers have begun to explore brain imaging as a novel input mechanism. This work has largely been targeted at allowing users to bypass the need for motor movement and to directly control computers by explicitly manipulating their brain signals. In our work, we use brain imaging to passively sense and model the user's state as they perform their tasks. In this workshop, we will discuss how we are using brain imaging to explore human cognition in the real world, to evaluate interface design, and to build interfaces that adapt based on user state. We will ground our discussions in user studies we have conducted to perform task classification while users perform various tasks. We will also provide lessons learned as well as a general methodology so that other HCI researchers can utilize these technologies in their work.

INTRODUCTION

Human-computer interaction researchers continually work to increase the communication bandwidth and quality between humans and computers. We have explored visualizations and multimodal presentations so that computers may use as many sensory channels as possible to send information to a human. Similarly, we have devised hardware and software innovations to increase the information a human can quickly input into the computer. Since we have traditionally interacted with the external world only through our physical bodies, these input mechanisms have all required performing some motor action, be it moving a mouse, hitting buttons, or speaking.

Recent advances in cognitive neuroscience and brain imaging technologies provide us with the unprecedented ability to interface directly with activity in the brain. Driven by the growing societal recognition for the needs of people with physical disabilities, researchers have begun to use these technologies to build brain-computer interfaces, communication systems that do not depend on the brain's normal output pathways of peripheral nerves and muscles [9]. In these systems, users explicitly manipulate their brain activity instead of using motor movements in order to produce signals that are measured and used to control computers.

While removing the need for motor movements in computer interfaces is challenging and rewarding, we believe that the full potential of brain imaging as an input mechanism lies in the extremely rich information it provides about the user. In our work, we will use brain imaging to passively monitor users while they perform their tasks in order to acquire a more direct measure of their state. We will use this state either as feedback to the user, as awareness information for other users, or as supplementary input to the computer so that it can mediate its interactions accordingly. This communication channel is unlike any we have had and has the potential not only to change the way we view ourselves acting within the environment and to augment traditional interaction processes, but also to transform the field of human-computer interaction.
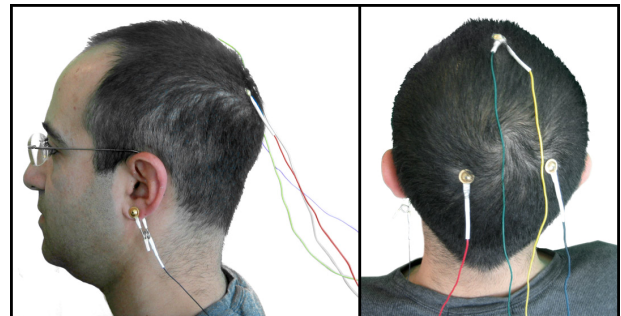

Figure 1. EEG Electrode placements in our experimental setup.

BACKGROUND

The human brain is a dense network of about 100 billion nerves cells called neurons. Each neuron communicates with thousands of others in order to regulate physical processes and to produce thought. Neurons communicate either by sending electrical signals to other neurons through physical connections or by exchanging chemicals called neurotransmitters (for a more detailed discussion of brain function, see [1]). Advances in brain sensing technologies allow us to observe the electrical, chemical, and physical changes that occur as the brain processes information or responds to various stimuli.

With our current understanding, brain imaging allows us only to sense general processes and not the full semantics of our thoughts. Brain imaging is not mind reading. For exam-

---

ple, although we can probably tell if a user is processing language, we cannot determine the semantics of the content. We hope that the resolution at which we are able to decipher thoughts grows as we increase our understanding of the human brain and abstract thought, but none of our work is predicated on these improvements happening.

Brain Imaging Technologies

While there are a myriad of brain imaging technologies, which we will review briefly in the workshop, we believe that only two, Functional Near Infrared (fNIR) imaging and Electroencephalography (EEG), present opportunities for inexpensive, portable, and safe devices, properties that are important for brain-computer interface applications within HCI research.

fNIR technology works by projecting near infrared light into the brain from the surface of the scalp and measuring optical scattering that represent localized blood volume and oxygenation changes (for a detailed discussion of fNIR, see [4]). These changes can be used to build detailed functional maps of brain activity. While we have begun exploring this technology, most of our current work has utilized EEG.

EEG uses electrodes placed directly on the scalp to measure the weak (5-100 µV) electrical potentials generated by activity in the brain (for a detailed discussion of EEG, see [10]). Because of the fluid, bone, and skin that separate the electrodes from the actual electrical activity, signals tend to be smoothed and rather noisy. Additionally, EEG is also susceptible to non-cognitive signals, or "noise", which have been problematic in traditional EEG work. However, we show how we can actually exploit this noise for accurate task classification within HCI applications.

USING BRAIN-SENSING SIGNALS

User State as an Evaluation Metric

One use of user state derived from brain imaging could be as an evaluation metric for either the user or for computer systems. Since we might be able to measure the intensity of cognitive activity as a user performs certain tasks, we could potentially use brain imaging to assess cognitive aptitude based on how hard someone has to work on specific tasks. With proper task and cognitive models, we could potentially use these results to generalize performance predictions in a much broader range of tasks and scenarios.

Rather than evaluating the human, a large part of human-computer interaction research is centered on the ability to evaluate computer hardware or software interfaces. This allows us not only to measure the effectiveness of these interfaces, but more importantly to understand how users and computers interact so that we can improve our computing systems. Thus far, we have been relatively successful in learning from performance metrics such as task completion times and error rates. We have also used behavioral and physiological measures to infer cognitive processes, such as mouse movement and eye gaze as measures of attention, or heart rate and galvanic skin response as measures of arousal and fatigue.

However, there remain cognitive phenomena such as cognitive workloads or particular cognitive strategies, which are hard to measure externally. For these, we typically resort to clever experimental design or subjective questionnaires which give us indirect metrics for specific cognitive phenomena. In our work, we explore brain imaging as a measure that more directly quantifies the cognitive utility of our interfaces. This could potentially provide powerful measures that either corroborate external measures, or more interestingly, shed light on the interactions that we would have never derived from external measures alone.

Adaptive Interfaces based on User State

If we take this idea to the limit and tighten the iteration between measurement, evaluation, and redesign, we could design interfaces that automatically adapt depending on the cognitive state of the user. Interfaces that adapt themselves to available resources in order to provide pleasant and optimal user experiences are not a new concept. In fact, we have put quite a bit of thought into dynamically adapting interfaces to best utilize such things as display space, available input mechanisms, device processing capabilities, and even user task or context.

In our work, we assert that adapting to users' limited cognitive resources is at least as important as adapting to specific computing affordances. One simple way in which interfaces may adapt based on user state is to adjust information flow. For example, verbal and spatial tasks are processed by different areas of the brain, and cognitive psychologists have shown that processing capabilities in each of these areas is largely independent [1]. Hence, even though a person may be verbally overloaded and not able to attend to more verbal information, their spatial modules might be capable of processing more data. Sensory processes such as hearing and seeing, have similar loosely independent capabilities.

Using brain imaging, the system could know approximately how the user's attentional and cognitive resources are allocated, and could tailor information presentation to attain the largest communication bandwidth possible. For example, if the user is verbally overloaded, additional information could be transformed and presented in a spatial modality, and vice versa. Alternatively, if the user is completely cognitively overloaded while they work on a task or tasks, the system could present less information or choose not to interrupt the user with irrelevant content until the user has free brain cycles to better deal with more information. This is true even if the user is staring blankly at the wall and there are no external cues that allow the system to easily differentiate between deep thought and no thought.

Finally, if we can sense higher level cognitive events like confusion and frustration or satisfaction and realization (the "aha" moment), we could tailor interfaces that provide feedback or guidance on task focus and strategy usage in

training scenarios. This could lead to interfaces that drastically increase information understanding and retention.

INITIAL RESULTS

For a BCI technology to be useful as a communication device, the system must be capable of discriminating at least two different states within the user. A computer can then translate the transitions between states or the persistence of a state into a form that is appropriate for a particular application [6].

Our initial (as yet unpublished) results represent out first steps in exploring how BCI technology can be applied to HCI research. The contributions of this work are twofold. First, we demonstrate that entry into the field does not necessarily require expensive high-end equipment, as is commonly assumed. Results from a user study we conducted show that we were able to attain 84.0% accuracy classifying three different cognitive tasks using relatively low cost ($1500) off-the-shelf EEG equipment (see Figure 1 for setup). Second, we present a novel approach to performing task classification utilizing both cognitive and non-cognitive artifacts measured by our brain sensing device as features for our classification algorithms. We present a second user study showing 92.4% accuracy classifying three tasks within a more ecologically valid setting, determining various user states while playing a computer game.

User Study 1

In our first study, we used machine learning approaches in an attempt to take only EEG data and classify which mental task a user was performing. We adopted the general experimental design presented by Keirn and Aunon [6] in an effort to extend their results using a low-cost EEG device.

Based on pilot studies with our system, we chose three tasks to classify: 1) *Rest*, in which we instructed participants to relax and to try not to focus on anything in particular; 2) *Mental Arithmetic*, in which we instructed participants to perform a mental multiplication of a single digit number by a three digit number, such as "7 × 836"; 3) *Mental Rotation*, in which we cued participants with a particular object, such as "peacock", and instructed them to imagine the object in as much detail as possible while rotating it.

For data collection, we used the Brainmaster AT W2.5, a PC-based 2-channel EEG system. This device retails for approximately USD$1500, comparable to the cost of a laptop computer. Eight (3 female) users, ranging from 29 to 58 years of age, volunteered for this study. All were cognitively and neurologically healthy, and all were right handed, except for one participant who had a slight nerve injury in his right hand and who had trained himself to depend more on his left hand.

In order to classify the signals measured using our EEG device, we first performed some basic processing to transform the signal into a time-independent data set. We then computed simple base features, which we mathematically

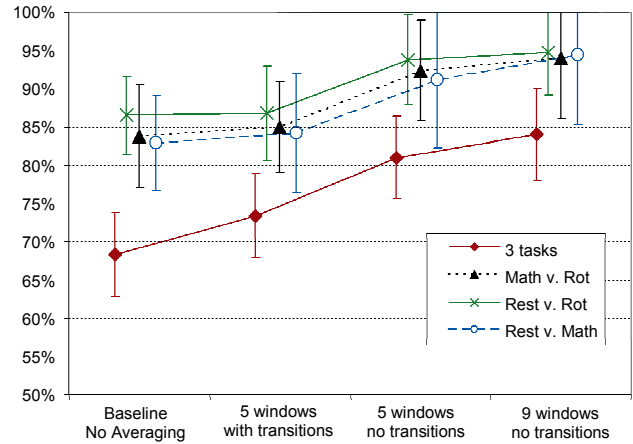Mean Classification Accuracy vs. Averaging Scenarios (Mental Tasks)

Figure 2. Plot of overall classification accuracies for three mental tasks in study 1 under various averaging scenarios. Error bars represent standard deviation.

combined to generate a much larger set of features. Next, we ran these features through a selection process, pruning the set and keeping only those that added the most useful information to the classification. This pruning step is important as it eliminates features not useful for classification and prevents a statistical phenomenon known as overfitting. Our feature generation and selection process is similar to that used by Fogarty et al. in their work on modeling task engagement to predict interruptability [8]. We then used the pruned set of features to construct Bayesian Network models for task classification.

The Bayesian Network classifiers for these three mental tasks yielded classification accuracies of between 59.3 and 77.6% (μ=68.3, σ=5.5), depending on the user. The prior for these classifications, or the expected result of a random classifier, was 33.3%. It should be noted that we would expect a human observer to perform only as well as a random classifier. The pair-wise classifiers had a prior of 50% and yielded accuracies of between 68.5 and 93.8% (μ=84.4, σ=6.0). Using various averaging schemes, we attained improvements of up to 21% (going from 67.9% to 88.9% for participant 1), with average improvements ranging between 5.1% and 15.7% for the 3 task classifier. See Figure 2 for a summary of these results.

We will present this entire methodology in more detail at the workshop as well as discuss how averaging may be used to enhance the classification accuracies, leading us to our final results. Interestingly, this methodology can be generalized for use with time sequence data coming from other sources, such as physiological or environmental sensors.

User Study 2

The classification accuracies we achieve in the first study suggest that we can indeed reliably measure and classify performance of our three mental tasks. However, we cannot be certain that the phenomena providing the classification power is entirely generated by neuronal firings in the brain.

We believe that various mental tasks are involuntarily coupled with physiological responses [7] and that it is difficult, if not impossible, to isolate cognitive activity using EEG in neurologically healthy individuals. This is problematic for researchers ultimately aiming to apply the technology to disabled individuals, as they have to guarantee that the features of interest are generated solely by the brain. For this reason, many researchers have conducted extensive work to remove the 'confounds' introduced by physical artifacts before classification [e.g. 5] or have limited their data collection to include only participants who suffer from the same disabilities as those as their target users.

However, since we are aiming to apply this to a generally healthy population, we only need to determine the reliability of the features in predicting the task. This concept was briefly explored by Chen and Vertegaal for modeling mental load in their physiologically attentive user interfaces [3]. In fact, if non-cognitive artifacts are highly correlated with different types of tasks or engagement, we should fully exploit those artifacts to improve our classification power, even though the neuroscience community has spent large efforts to reduce and remove them in their recordings.

We conducted a second user study to explore using both cognitive and non-cognitive artifacts to classify tasks in a more realistic setting. The tasks we chose involved playing a PC-based video game. This experiment serves as a demonstration of a general approach that could be applied to a much broader range of potential applications.

The game we selected for this experiment was Halo, a PC-based first person shooter game produced by Microsoft Game Studios. The tasks we tested within the game were: 1) *Rest*, in which participants were instructed to relax and fixate their eyes on the crosshairs located at the center of the screen; 2) *Solo*, in which participants navigated the environment, shooting static objects or collecting ammunition scattered throughout the scene; 3) *Play*, in which participants navigated the environment and engaged an enemy controlled by an expert Halo player.

Using the same processing methodology as before, the baseline classification accuracies were between 65.2 and 92.7% (μ=78.2, σ=8.4) for the 3-task classifiers and 68.9 and 100% (μ=90.2, σ=8.5) for the pair-wise classifiers. After averaging, we were able to achieve 83.3 to 100% accuracies (μ=92.4, σ=6.4) for 3-task and 83.3 to 100% (μ=97.6, σ=5.1) for pair-wise comparisons. While shifted slightly higher, the graph of the results looks similar to that of user study 1.

While it would have been possible to achieve reasonable classification accuracy by hooking into the game state, or even by monitoring the keyboard and mouse, our goal in running this experiment was not to show the most ideal way of discriminating these tasks but to demonstrate the impact of non-cognitive artifacts on EEG-based task classification in a realistic computing scenario. We assert that EEG shows interesting potential as a general physiological input sensor for distinguishing between tasks in a wide variety of computing applications.

CONCLUSION

The idea of using brain imaging in HCI work has been previously proposed [11] and we believe that the technology is now ripe for us to re-articulate this unexplored area of inquiry as well as the challenges that have to be addressed in order for this research project to be a success. We believe that our initial methodology and results will quickly get other researchers caught up with the topic and look forward to inspiring discussion that follows.

REFERENCES

1. Baddeley, A. D. (1986). Working memory. New York: Oxford University Press.
2. Carey, J (ed.). (2002). Brain facts: A primer on the brain and nervous system (4th edition). Washington DC, USA: Society for Neuroscience.
3. Chen, D., & Vertegaal, R. (2004). Using mental load for managing interruptions in physiologically attentive user interfaces. *SIGCHI 2004*, 1513-1516.
4. Coyle, S., Ward, T., Markham, C., & McDarby, G. (2004). On the suitability of near-infrared (NIR) systems for next-generation brain-computer interfaces. *Physiological Measurement, 25*, 815-822.
5. Gevins, A., Leong, H., Du, R., Smith, M.E., Le, J., DuRousseau, D., Zhang, J., & Libove, J. (1995). Towards measurement of brain function in operational environments. *Biological Physiology*, 40, 169-186.
6. Keirn, Z.A., & Aunon, J.I. (1990). A new mode of communication between man and his surroundings. *IEEE Transactions on Biomedical Engineering, 37(12)*, 1209-1214.
7. Kramer, A.F. (1991). Physiological metrics of mental workload: A review of recent progress. In Multiple Task Performance (ed. Damos, D.L.), 279-328.
8. Fogarty, J., Ko, A.J., Aung, H.H., Golden, E., Tang, K.P., & Hudson, S.E. (2005). Examining task engagement in sensor-based statistical models of human interruptibility. *SIGCHI 2005*, 331-340.
9. Millan, J. Adaptive brain interfaces. *Communications of the ACM, 46(3)*, 74-80.
10. Smith, R.C. (2004). Electroencephalograph based brain computer interfaces. *Thesis for Master of Engineering Science*, University College Dublin.
11. Velichkovsky, B., & Hansen, J.P. (1996). New technological windows into mind: There is more in eyes and brains for human-computer interaction. *SIGCHI 1996*, 496-503.