Your Brain, Your Computer, and You
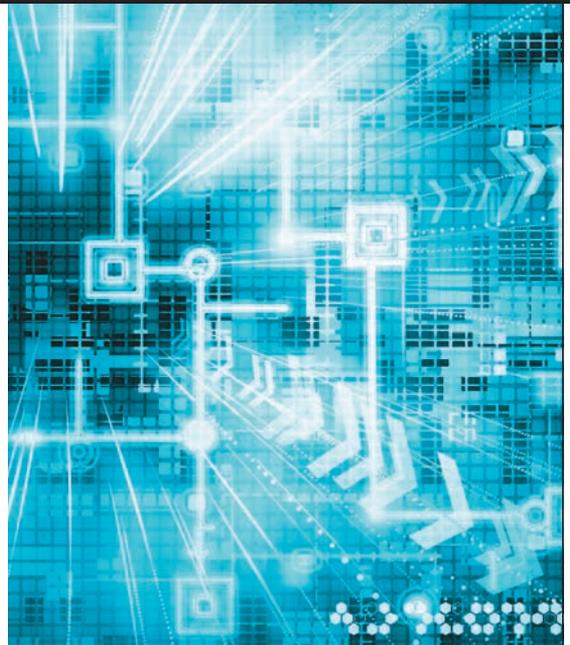
→ **Evan M. Peck, Erin Treacy Solovey, Krysta Chauncey, Angelo Sassaroli, Sergio Fantini, and Robert J.K. Jacob,** *Tufts University*

→ **Audrey Girouard,** *Queen's University*

→ **Leanne M. Hirshfield,** *Hamilton College*

Passive brain-computer interfaces increase the bandwidth from user to computer in new and uniquely powerful ways.

Every day, we experience a disconnect with our computer. There's so much we can't tell it, and so much it can't respond to. For example, it doesn't know when we're tired or distracted, and it can't recognize that we're cranky because we just spent an hour in traffic driving to the office.

At our workspaces, most of us are within arm's reach of a desktop computer as well as other devices such as a laptop computer and smartphone. We might be writing a paper, checking e-mail, monitoring the news, and glancing at Twitter all within a five-minute period. But there's a problem: with each device or application that expands the bandwidth of available information, the computer's understanding of us remains unchanged.

What if we consumed food in the same way that we consume information? What if we overate as often as we cognitively overloaded ourselves, or stopped what we were doing to snack as frequently as we delayed work to check our e-mail?

It's not really a fair comparison: our brain has evolved to deal with rich, complex environments, and we can digest information much more than three times a day. And yet our current environment is more heavily saturated with data than at any moment in human history, just as food is more plentiful in the developed world than at any time in history.

Professional athletes don't eat junk food before events because they know that their eating habits impact their abilities. They know what, when, and how much to eat to perform at their highest level. However, as rising rates of obesity and related health disorders such as diabetes attest, most of the rest of us aren't so good at managing our diet.

Are we likewise digesting information in a way that is optimal for our cognitive performance? Do we know what, when, and how much information our brains should consume? What do we do if we're digesting too much or the wrong type of information?

These are the types of questions we ask at Tufts University's Human-Computer Interaction (HCI) Lab (http://hci.cs.tufts.edu). During the past few years, we have wrestled with the myriad challenges of integrating brain-computer interfaces (BCIs) into everyday situations. At the technical level, this involves addressing two basic problems: how to measure and identify brain states, and how to respond or adapt to those brain states.

FROM BRAINS TO BYTES

To learn how to provide computers with more information about the user's state, we've turned to the field of cognitive neuroscience, which examines how brain activity shapes and reflects behavior. Researchers in this field often use functional magnetic resonance imaging (fMRI) as a window into this brain activity.

Anyone who has participated in an fMRI experiment knows that the methodology is not conducive to a pleasant or unobtrusive experience. Subjects must lie completely motionless in a large, loud chamber surrounded by powerful, spinning magnets and their support infrastructure. These scanners are also

Figure 1. The arrangement of detectors and light sources on an optical probe used for functional near-infrared spectroscopy can differ depending on the probe's design. This probe has five light sources and one optical detector.

remarkably expensive, affordable to only the top universities and hospitals, and require legions of technical staff and entire purpose-built buildings.

Although you won't see an fMRI device docked on your car's dashboard or embedded in your workstation anytime soon, neuroimaging technologies can plausibly be integrated into an office setting as long as they aren't invasive, inflict minimal physical discomfort, and are small enough to be portable. Two technologies that meet these criteria, electroencephalography (EEG) and functional near-infrared spectroscopy (fNIRS), could bring us a step closer to brain imaging in everyday life.

EEG measures the electrical activity on the scalp that is created by neurons firing in the brain. A set of electrodes placed on the participant's head, often using a conductive gel and cap, records this electrical activity. While EEG can accurately reflect changes very quickly—on a millisecond-to-millisecond basis—researchers can make little to no judgment about the underlying neural regions that might be contributing to the signal.

In contrast, fNIRS uses near-infrared light to measure the concentration and oxygenation of blood in tissue at depths of 1-3 cm on the brain's surface. An optical probe like that shown in Figure 1 delivers light into the forehead, where it is diffusely reflected by the tissues of the cortex. Because oxygenated and deoxygenated hemoglobin are the primary absorbers of light at near-infrared wavelengths, researchers can calculate the blood oxygenation, as well as the concentration of blood in the tissue, by examining the amount of light returning to a detector on the probe.

Changes in blood flow take 5-7 seconds to register after initial brain activation, so fNIRS is sensitive to slower signals than those recorded with EEG. However, because probe placement corresponds to the region of the brain being measured, fNIRS has better spatial accuracy and enables more localized observations about brain signals' neural underpinnings.

The HCI Lab primarily relies on fNIRS, as it is generally less prone to movement artifacts and less cumbersome to set up. In 2009, we tested the robustness of fNIRS signals associated with common computer user tasks such as forehead movement, minor and major head movements, typing, clicking, respiration, ambient light exposure, eye blinking, and eye movement. We found that only forehead movement and major head movement were extremely problematic to the signal, demonstrating that fNIRS is a reasonable tool in everyday situations (E.T. Solovey et al., "Using fNIRS Brain Sensing in Realistic HCI Settings: Experiments and Guidelines," *Proc. 22nd Ann. ACM Symp. User Interface Software and Technology,* ACM Press, 2009, pp. 157-166).

The fNIRS measurements gleaned from the brain vary from person to person. Even if we can distinguish between two brain states in one person, the cognitive "footprint" may differ slightly between subjects. This cross-subject instability requires us to draw heavily from machine-learning literature to classify brain states. Using classification algorithms in a real-time environment, we can move closer to an adaptive BCI.

TWO WAYS OF THINKING

Once we've measured and identified brain states, we can use this information to try to enhance users' experiences or cater to their cogni-
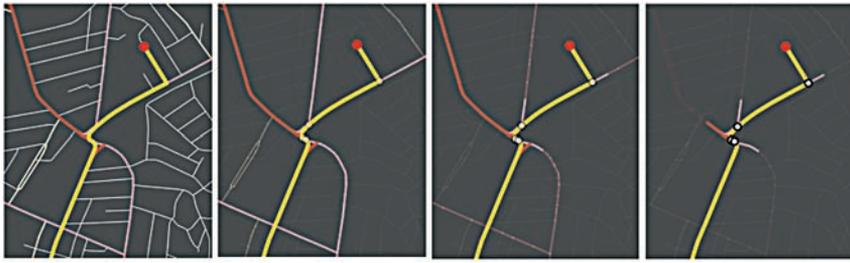
Figure 2. One possible application of brain-computer interfaces is adjusting the visual representation of data to fit users' needs or cognitive strengths. This example shows a vehicle GPS system slowly fading noncritical information in and out—in this case, surrounding streets—depending on the driver's workload.

tive strengths. Most BCI research has focused on the disabled, but the general public could also benefit.

There's an important distinction between using brain data as *passive input* and using it as *explicit input* to an interface. The latter relies on a one-to-one mapping of brain states to explicit actions—for example, the user thinks about raising his or her left arm to move the mouse cursor left. However, identifying a wide range of explicit thoughts in a complex environment is very difficult, and we're likely many years away from direct-control interfaces that contain natural and intuitive mappings for everyday use. In the near future, applications will use explicit brain-data input only if it provides users with an ability that justifies the interface's unintuitive controls.

Instead of trying to understand the effort required to control an explicit interface, we use passive brain-data input, leaving as much of the user's cognitive ability as possible to the task itself. Our focus is on looking for and identifying brain states that occur naturally during a task. Rather than identify specific thoughts, we make general observations about the brain.

For example, in one early experiment we had participants try to memorize the number of different-colored sections on a rotating cube. The number of colors subjects had to juggle in memory corresponded to varying levels of working memory. By analyzing the fNIRS signals

during performance of these tasks, we identified three working-memory conditions (A. Sassaroli et al., "Discrimination of Mental Workload Levels in Human Subjects with Functional Near-Infrared Spectroscopy," *J. Innovative Optical Health Services,* vol. 1, no. 2, 2008, pp. 227-237).

In another experiment suggesting that fNIRS could be useful for interface evaluation, we were able to separate the cognitive workload associated with using an interface's semantic and syntactic portions as well as distinguish spatial from verbal workload (L.M. Hirshfield et al., "Brain Measurement for Usability Testing and Adaptive Interfaces: An Example of Uncovering Syntactic Workload with Functional Near Infrared Spectroscopy," *Proc. 27th Int'l Conf. Human Factors in Computing Systems,* ACM Press, 2009, pp. 2185-2194).

Finally, in the closest approximation of a real-world scenario, we used fNIRS signals to discriminate between two general difficulty states experienced by users as they played Pac-Man at varying difficulty levels (A. Girouard et al., "Distinguishing Difficulty Levels with Non-Invasive Brain Activity Measurements," *Proc. 12th IFIP TC 13th Int'l Conf. Human-Computer Interaction,* Part 1, ACM Press, 2009, pp. 440-452).

FROM BYTES TO BRAINS

In imagining the types of interfaces that might adapt to cognitive load, it's important to consider what

types of changes are appropriate for passive input. We lean toward adaptations that are as unobtrusive as possible, are dependent on gentle visual manipulations (perhaps using focus versus context, or varying levels of summarization), or involve behind-the-scenes changes that are not distracting or immediately noticeable to the user (such as filtering or prioritizing data). Most importantly, we aim to stay far away from the Midas Touch scenario, in which the user unintentionally triggers dramatic adaptations with every whim and fancy.

One BCI application that could become commonplace in the future is adjusting the visual representation of data to fit users' needs or cognitive strengths. Visualizations are becoming increasingly important in interpreting complex datasets, but they provide only a limited snapshot view of data. Adapting the type of visualization to different types of users or to certain times, mental states, or tasks could result in more effective results. For example, as Figure 2 shows, a vehicle GPS system could slowly fade noncritical information in and out depending on the driver's workload.

BCIs could also help identify weak links in current human-robot interaction techniques. For example, a robot might not understand when we're stressed out, cognitively overloaded, or just plain uncomfortable dealing with it, and a BCI could aid the robot in understanding these cues and acting in a more socially appropriate manner.

Returning to the office scenario, where our attention is divided among multiple devices, imagine a BCI that monitors the cognitive effort we're spending on work and accordingly adjusts the other applications we're using. For example, if you're expending a lot of your mental resources while writing an article, the system alerts you to only the most important e-mail messages or news items; perhaps it moves alerts to your

phone, where they're potentially less distracting.

Passive BCIs increase the bandwidth from user to computer in new and uniquely powerful ways. Someday they might help us to not only better understand the information on our computer, but also help our computer better understand us. HCI Lab researchers continue to imagine and prototype a wide array of BCIs as well as to develop ways to evaluate them experimentally. **C**

Evan M. Peck is a PhD student in the Human-Computer Interaction (HCI) Group within the Department of Computer Science at Tufts University. Contact him at evan.peck@tufts.edu.

Erin Treacy Solovey is a PhD student in the HCI Group at Tufts University. Contact her erin.solovey@tufts.edu.

Krysta Chauncey is a postdoctoral researcher in the HCI Group at Tufts University. Contact her at krysta.chauncey@tufts.edu.

Angelo Sassaroli is a research assistant professor in the Department of Biomedical Engineering at Tufts University. Contact him at angelo.sassaroli@tufts.edu.

Sergio Fantini is a professor in the Department of Biomedical Engineering at Tufts University. Contact him at sergio.fantini@tufts.edu.

Robert J.K. Jacob is a professor of computer science and leads the HCI Group at Tufts University. Contact him at Jacob@cs.tufts.edu.

Audrey Girouard received a PhD in computer science from Tufts University as part of the HCI Group and is currently a postdoctoral Fellow in the Human Media Lab at Queen's University, Kingston, Ontario, Canada. Contact her at audrey@cs.queensu.ca.

Leanne M. Hirshfield received a PhD in computer science from Tufts University as part of the HCI Group and is currently a research professor at Hamilton College, Clinton, New York. Contact her at leannehirshfield@gmail.com.
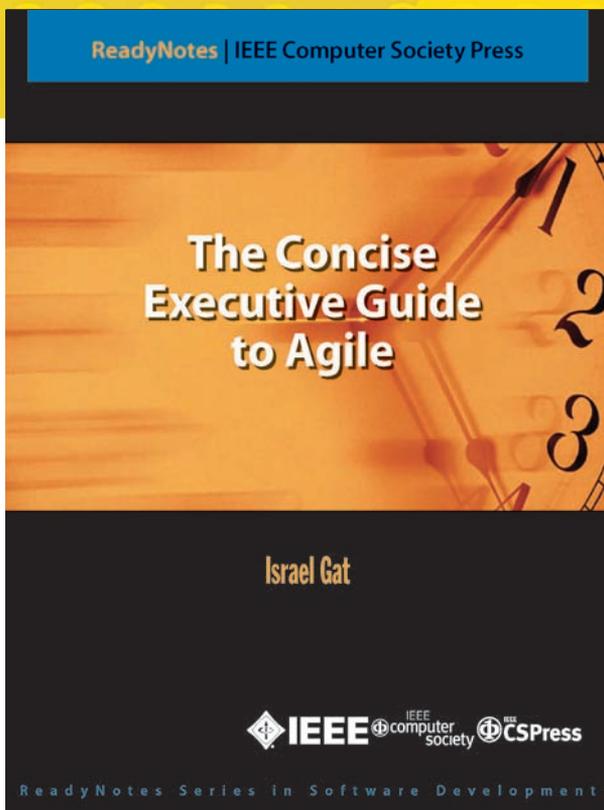
This article was featured in

computing|now

ACCESS | DISCOVER | ENGAGE

For access to more content from the IEEE Computer Society,
see computingnow.computer.org.

◆IEEE   IEEE⊕ computer society

Top articles, podcasts, and more.

computingnow.computer.org