Neuroscience and the Future of Human-Computer Interaction

Brad S. Minnery
The MITRE Corporation | bminnery@mitre.org

Michael S. Fine
The MITRE Corporation | mfine@mitre.org

If Carl Sagan had been a neuroscientist instead of an astronomer, he might have mused wondrously about the "billions and billions" of neurons that make up the human brain—approximately one hundred billion neurons with each neuron wired to communicate with thousands of neighbors. This massive mesh of computation gives rise to the impressive spectrum of human cognitive capabilities. To date, most HCI researchers have focused on readily observable behavioral metrics (for example, the speed of a keystroke or the accuracy of a mouse click) rather than on the mental machinery operating under the surface. Modern neuroscience offers HCI researchers a way to "lift the veil" on user cognition, greatly expanding the available tool kit for both research and design.

Neuroscience is the study of the brain and nervous system. Although the field concerns itself with the study of neurobiological systems at the smallest scales (molecules and genes), neuroscience also works to understand how the nervous system contributes to macro-level behaviors of interest to HCI researchers. Over the past 20 years, our understanding of brain function

has expanded dramatically—partially driven by advances in experimental methodology, but also enabled by a swell of research funding for the study of brain-related disorders like autism, Parkinson's disease, and traumatic brain injury. This growth is reflected in the scale and diversity of membership in the Society of Neuroscience. Founded in 1969, its ranks have doubled in the past 20 years, to more than 38,000 members. A quick tour of the society's annual meeting reveals the broad range of cognitive functions that neuroscientists are studying from a biological perspective—from perception to decision to action.

These advances in neuroscientific discovery are poised to have a profound impact on multiple facets of HCI research and system design. For starters, neuroscience enables us to build more accurate and robust models of human cognitive functions. These models may allow us to evaluate usability and predict user behavior through computation alone. In addition, neuroscience research methods will allow HCI researchers to answer questions that previously lay outside the reach of their methodologi-

cal toolkit—measuring hidden metrics like interest, affect, or satisfaction. Even further down the road, neuroscience offers the potential to truly close the gap between humans and computers through the development of devices that engage directly with the brain. The aim of this article is to describe these and other synergies between neuroscience and HCI and to make a case for greater collaboration between the two communities.

Building Better Models
The idea that one could use cognitive models in lieu of real humans as a way of reducing the time/costs associated with designing an interface and conducting usability studies is not in itself novel [1]. This so-called "engineering" approach employs sophisticated cognitive models—such as EPIC, SOAR, and ACT-R—to predict how a user (or class of users) might interact with a given interface to perform a specified task. These models represent the cumulative insights of decades of psychological and behavioral research, and their ability to replicate human behavior in some narrow domains has been remarkable. However, a common criticism leveled against

[1] Card, S., T. Moran, and A. Newell. *The Psychology of Human-Computer Interaction.* Hillsdale, NJ: Erlbaum, 1983.
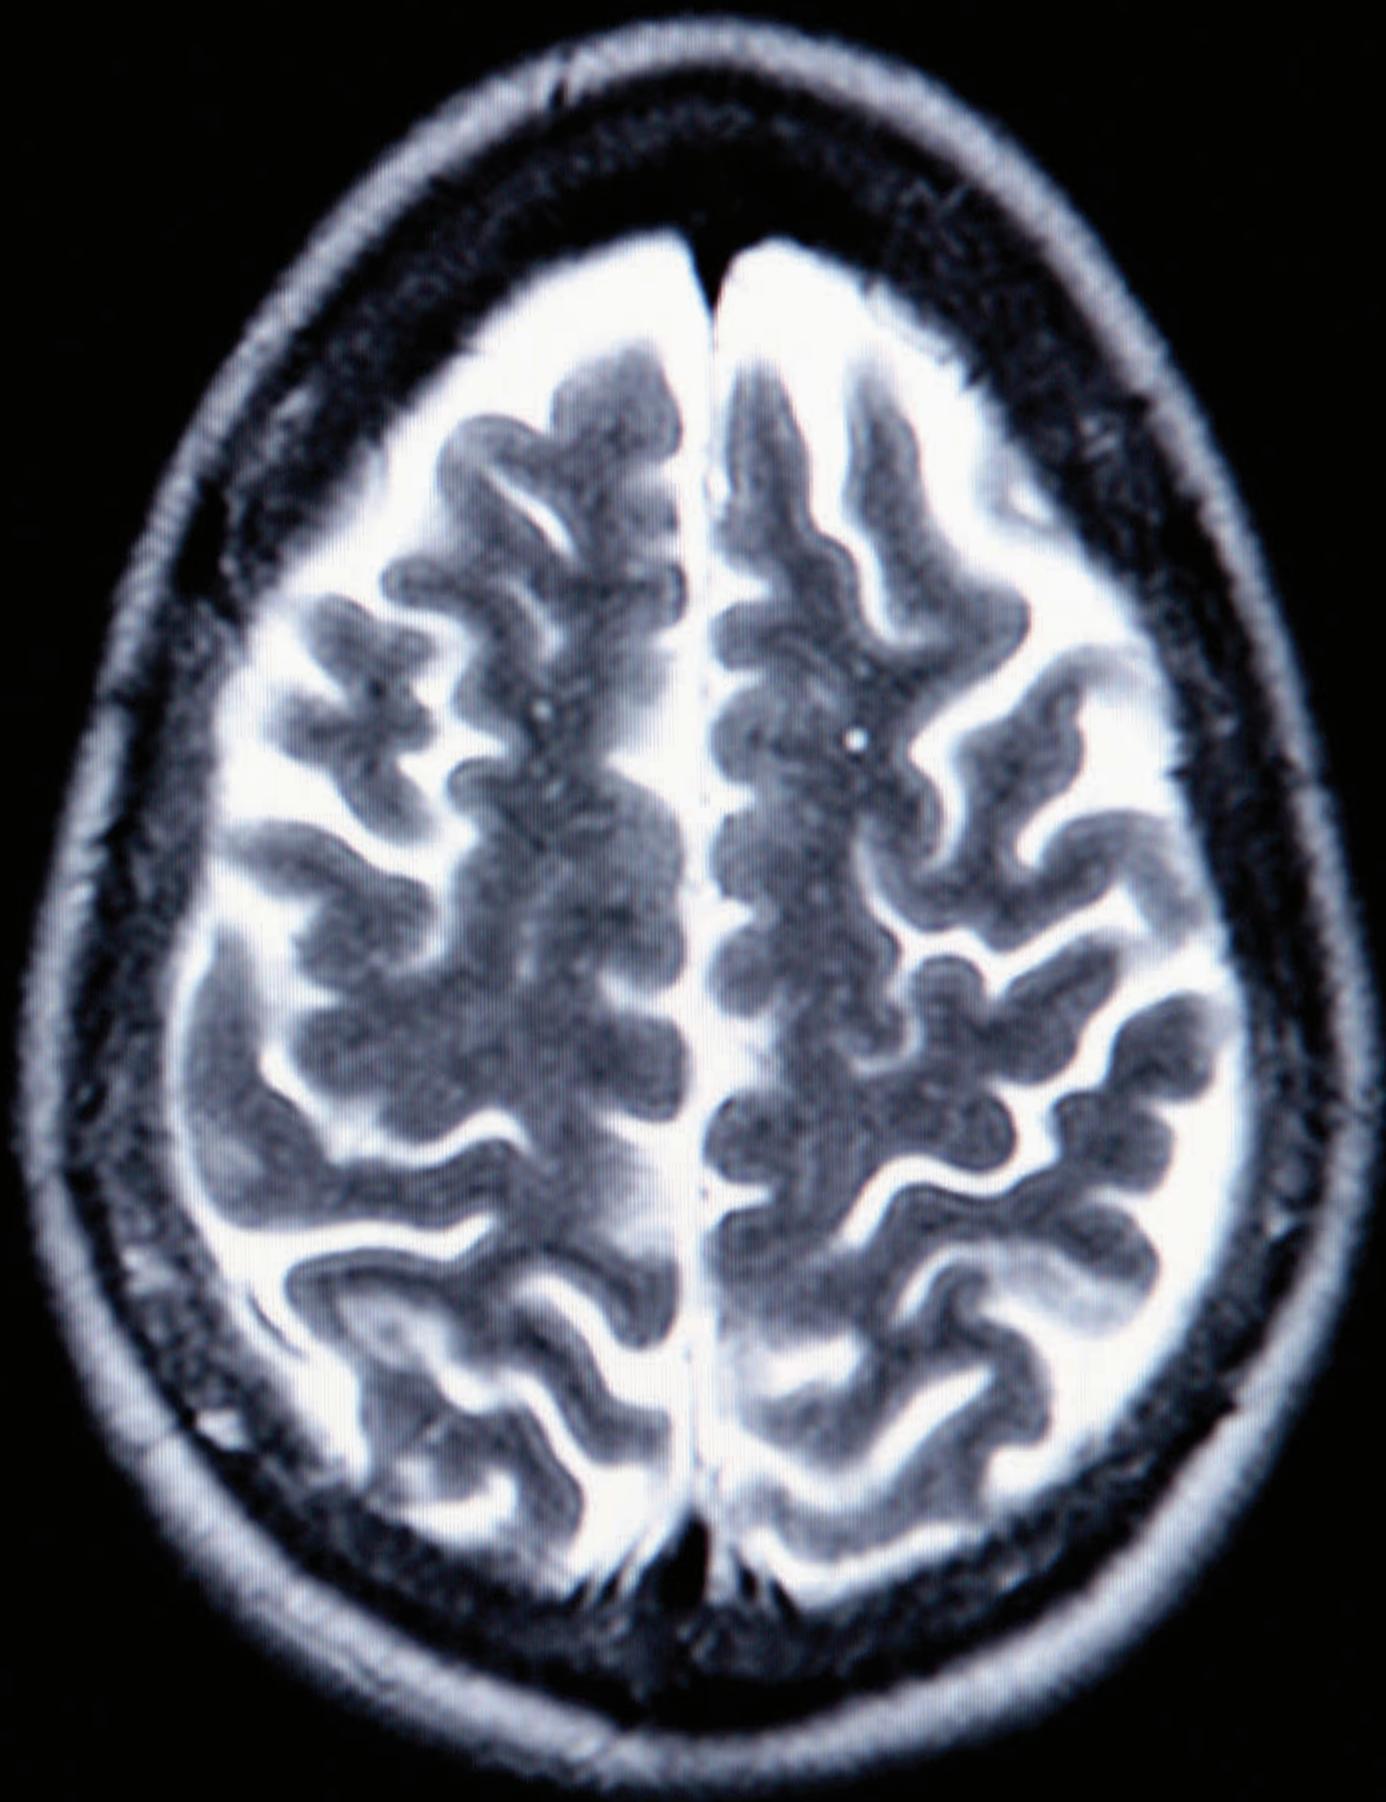
Neuroscience Tools

The tools with which neuroscientists examine brain function have multiplied over the past 15 years. While some of these tools remain fairly specialized, they are becoming more common within academia, research labs, and even commercial companies. Expertise with these and other methods is widespread throughout the neuroscience community, so the best option for the neuroscience-inclined HCI researcher (who isn't eager to wrestle with magnets or electrodes) would be to partner with an academic neuroscientist who shares an interest in the neuropsychological bases of human-computer interaction.

With the advent of Functional Magnetic Resonance Imagery (fMRI) in the 1990s, it became possible to "see" brain activity in subjects as they performed cognitive tasks. fMRI works by imaging changes in cerebral blood flow, which provides a proxy signal for neural activity throughout the whole brain. However, despite its advantages over more invasive recording techniques, fMRI is not without its drawbacks. For instance, human subjects must remain relatively motionless during an fMRI experiment, which limits the range of behaviors that can be studied. Furthermore, the cost and complexity of fMRI act as barriers to HCI researchers looking to incorporate neuroscientific approaches into their work.

There are alternatives, however, that may provide the "window into the brain" that HCI requires. Less capital-intensive methods like functional near infrared spectroscopy (fNIRS) also offer a noninvasive way to probe brain function. Like fMRI, the fNIRS signal reflects the dynamics of cerebral blood flow. However, fNIRS penetrates only a few millimeters below the brain's surface, and unlike fMRI, does not reveal the activity of deeper brain structures. Nonetheless, many of the higher cognitive functions of interest to HCI researchers —working memory, executive control, and visual-spatial processing—are localized within the outermost layer of brain tissue called the cerebral cortex, and thus are readily accessible to fNIRS.

In addition to fMRI and fNIRS, electroencephalography (EEG) —a stalwart tool of cognitive neuroscience researchers since the 1950s—offers a lightweight and low cost option. EEG setups typically consist of a web of electrodes that are carefully fitted over a person's scalp to record low-amplitude electrical brain activity at the surface of the skull. And although the drawbacks of EEG (poor spatial resolution, high susceptibility to electrical noise) limit its utility, recent advances in electrode design and signal processing techniques have expanded the range of applications for which EEG is suitable. EEG has been used to quantify cognitive workload in complex operating environments like air traffic control [2], and has also been used to quantify operator vigilance [3]; as well as to classify mental states like arousal and fatigue [4]. These findings may influence design decisions for features like adaptive automation (When should control be transferred to the machine?), and augmented cognition (In what situations does the user need assistance?). Despite these advances, EEG will need to be less obtrusive and less prone to contamination from outside sources to be useful outside the laboratory. But as the technology improves, EEG has the potential to greatly enhance our understanding of user cognitive function during the interaction experience.

this class of models is that they tend to reduce much of cognition to a collection of if-then rules, even though many, if not most cognitive functions (including perception, sensorimotor control, and some types of learning), are not neatly decomposable into a series of articulated statements. As a result, these models tend to be brittle and fail to capture the gamut of relevant user behaviors.

Computational neuroscience models, which aim not just to replicate cognitive functions but also to explain how such functions arise from underlying brain activity, represent a complementary approach. On the one hand, neuroscience-based approaches can help to improve the design of traditional cognitive models by providing a sort of biological ground truth against which to judge the plausibility of competing hypotheses and model architectures. More important, neural models may offer new functionality in domains where traditional approaches have been lacking. Visual perception, for instance, is one of the most extensively studied cognitive subsystems among neuroscientists, but traditional cognitive models have been notoriously poor at replicating the human ability to process raw visual data. As a result of neuroscience research, great headway has been made in understanding how humans make sense of their visual environment. Recent collaborative efforts between ACT-R modelers at Carnegie Mellon University and computational cognitive neuroscientists at the University of Colorado at

Boulder have led to the development of a cognitive architecture that combines the functionality of a visual neuroscience model with traditional rule-based elements [5]. Such hybrid architectures may represent the future of cognitive modeling approaches to usability analyses.

Inside the User's Head

Only a small percentage of current neuroscience research is explicitly aimed at understanding aspects of HCI. Nonetheless, some recent neuroimaging experiments point to ways in which experimental neuroscience methodology might be leveraged to measure facets of the user-interaction experience at a deeper level than can be achieved with other contemporary methods. For instance, modern neuroscience has begun to characterize the brain circuitry that governs reward-related behaviors, with fMRI experiments revealing that unexpected rewards elicit activation in areas of the human brain that utilize the chemical transmitter dopamine [6]. (A reward in these experiments is typically anything from a squirt of juice to a $10 bill.) These studies raise the intriguing possibility that neuroimaging techniques might someday be used to identify which aspects of the interaction experience a user finds pleasing.

Another example of applicable neuroscience research comes from a series of experiments examining how humans perceive computer-animated characters that vary in their degree of physical realism. One study showed that the tendency of a subject to perceive

a virtual character as realistic is correlated with activation in areas of the brain known to play a role "mentalizing" [7]. Mentalizing refers to our ability to place ourselves in the mind of another person and predict their intentions. It is fundamental to human social interaction. Research points to the possibility that neuroimaging methods might be used to assess the degree in which a user perceives a virtual entity (for instance, an avatar) as a fellow autonomous being, or merely as a non-sentient computer artifact.

The main challenge ahead will be to demonstrate that a neuroscientific approach to HCI adds value beyond what can be gleaned from behavioral studies alone. If other disciplines offer any indication, the outlook is promising. Consider the medical field, where tests that reveal what is going on "under the hood" (angiograms, throat cultures, or simple blood tests) are ordered precisely because they provide diagnostic value beyond what is available through observation of the patient's symptoms alone. And although the user (like the patient) possesses a unique awareness of what's happening inside his or her own brain (or body), and therefore can provide useful information simply by describing his or her own thought processes, an individual's ability to introspect is limited. In fact, a major thrust of modern psychological research is focused on mapping the extent of so-called implicit cognition—that vast chunk of the cognitive iceberg that floats beneath the surface of conscious thought but drives

behavior in powerful ways [8]. Neuroscience will likely make valuable contributions to the discipline of HCI by providing a richer account of user cognition than that which is obtained from any other source, including the user himself.

Current Research

As described here, vision is one of the brain's most extensively studied subsystems. In addition, the brain's memory circuits have also been the subject of intense research. Since visual perception and memory are key areas of study in HCI, neuroscience-based models of these functions may be particularly well poised to have an impact on HCI research. In our work here at the MITRE Corporation, we are exploiting models of visual attention and memory to predict how visual display properties influence perception and recall by users. As part of this study, we are implementing a neurocomputational model of visual attention developed by researchers

The main challenge ahead will be to demonstrate that a neuroscience approach to HCI adds value beyond what can be gleaned from behavioral studies alone.

[2] Wilson, G.F. and C.A. Russell. "Operator functional state classification using multiple psychophysiological features in an air traffic control task." *Human Factors* 45, no. 3 (2003): 381-389.

[3] Wilson, G.F. and C.A. Russell. "Performance enhancement in an uninhabited air vehicle task using psychophysiologically determined adaptive aiding." *Human Factors* 49, no. 6 (2007): 1005-1018.

[4] Muller, K.R., M. Tangermann, G. Dornheg, M. Krauledat, G. Curio, and B. Blankertz. "Machine learning for real-time single-trial EEG-analysis: from brain-computer interfacing to mental state monitoring." *Journal of Neuroscience Methods* 167, no.1 (2008): 82-90.

[5] Jilk, D.J., C. Lebiere, R.C. O'Reilly, and J.R. Anderson. "SAL: An explicitly pluralistic cognitive architecture." *Journal of Experimental & Theoretical Artificial Intelligence* 20, no. 3 (2008): 197-218.

[6] D'Ardenne, K., S.M. McClure, L.E. Nystrom, and J.D. Cohen. "BOLD Responses Reflecting Dopaminergic Signals In the Human Ventral Tegmental Area." *Science* 319, no. 5867 (2008): 1264-1267.

at the University of Southern California and the California Institute of Technology [9]. The model applies a series of feature-specific filters (color, intensity, and orientation) that emulate the processing that occurs in the retina and brain as a user views an image. For a given input image, the model produces a corresponding salience map (Figure 1) that quantitatively describes which regions of the image are most likely to draw the user's gaze—in other words, which regions "pop out" the most.

Our goal is to apply this model to investigate the relationship between the visual properties of an interface and an operator's "situational awareness." Situational awareness is the ability to perceive and understand a changing environment and predict probable future events. Memory facilitates situational awareness by enabling a user to maintain a continuously updated picture of his or her environment (playing back past events). A military commander, for example, needs situational awareness to keep track of assets and adversaries within the battle space over time. In an ongoing set of experiments, we are studying how visual salience affects memory. In particular, we are testing participants' ability to remember the location of icons on a 2-D map and examining whether greater icon salience correlates with lower spatial recall error. The broader scientific aim of our research is to examine how attention and memory subsystems interact within the brain. However, studies such as these, as well as

research by other groups [10], are laying the groundwork for a new class of smart interfaces that will be able to improve operator performance by monitoring—and by adaptively modifying—the contents and configuration of the current display.

Our own experiments are examples of research where neuroscience and HCI intersect. For instance, other research efforts for the defense and automotive industries seek to correlate neurophysiological measures of cognitive workload with the properties of a user interface, thereby providing a direct link between interface properties and brain activity. In fact, many of the questions that drive contemporary cognitive neuroscience research also speak to issues in interface design. Opportunities abound for HCI researchers to collaborate with neuroscientists to address these topics of common interest. In addition to the earlier examples, neuroscientists are also studying how the brain:

- manages attentional resources across multiple sensory channels,
- navigates through virtual as well as real environments,
- learns the most efficient procedures for performing a task,
- allocates trust in competitive and/or cooperative situations involving multiple agents or other users.

Brain-Machine Interface
The most extreme example of how neuroscience might change the trajectory of HCI comes from the nascent field of brain-

machine interface (BMI). BMI achieves a literal realization of the human computer interaction paradigm by physically connecting man and machine. Over the past several years, BMI research has led to the development of brain-implantable chips that can translate a user's neural impulses into a signal for controlling an external device, such as a robotic arm [11, 12]. Current state-of-the-art neuroprosthetic devices are far from the sleek biomorphic appendages featured in science fiction films, but are rather first-generation prototypes possessing minimal functionality. Although these technologies represent a remarkable step forward for amputees and other disabled persons, it is unlikely that healthy individuals will volunteer to undergo risky brain surgery simply for the potential interaction benefits. However, there is ongoing research to investigate the use of low-cost, noninvasive neural recording techniques—like electroencephalography (EEG)—as a basis for direct neural control of external devices [13]. These noninvasive devices may obviate the risks associated with implantable systems and provide a pathway for making BMI accessible to non-disabled users.

In fact, efforts to turn EEG into a sort of "BMI for the masses" are well under way, with at least two companies, Emotiv (www.emotiv.com) and Neurosky (www.neurosky.com), having developed EEG-based game controllers. It is unclear whether these stripped-down commercial systems (that feature far fewer electrodes than
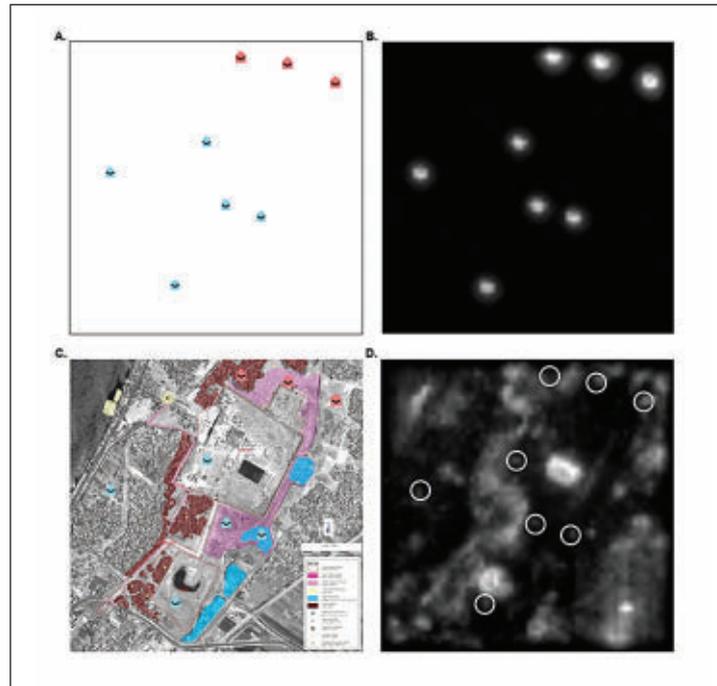
▶ Figure 1: An icon's salience, or "pop-out," is determined in part by the properties of the display. On a white background (A), each icon is readily found; this is quantitatively captured by the model-generated salience map (in B). When pasted on a map (C), the same icons are far less salient (D). In D., white circles denote the position of each icon.

a typical laboratory configuration) actually live up to the marketing campaign. Neural activity patterns recorded through EEG usually reflect slow changes in mental state, such as changing levels of attention and arousal. Without significant advances, it's unlikely that gamers will be able to execute a rapid sequence of actions (kick-punch-jump) with their thoughts alone. But as we've seen with hacks of the Wii controller, placing an EEG device in the hands of eager users may result in new innovative applications. As these and other neurally enabled technologies become more mainstream in the next decade, members of the HCI community should be ready to capitalize on their full potential.

ABOUT THE AUTHORS
Brad Minnery is the neurotechnology group leader in The MITRE Corporation's emerging technologies office. His interests included neurobiologically inspired intelligent systems and the cognitive neuroscience of human-machine interaction. He received his Ph.D. in neurobiology from the University of Pittsburgh.

Michael Fine is a senior artificial intelligence engineer at The MITRE Corporation, where he studies the interaction of visual attention with learning and memory. Fine has a Ph.D. in biomedical engineering from Washington University, where he developed computational models to predict human motor behavior in virtual environments.

[7] Chaminade, T., J. Hodgins, and M. Kawato. "Anthropomorphism influences perception of computer-animated characters' actions." *Social Cognitive and Affective Neuroscience* 2, no. 3 (2007): 206-216

[8] Bargh, J.A. and M.L. Ferguson. "Beyond Behaviorism: On the automaticity of higher mental processes." *Psychological Bulletin* 126, no.6 (2000): 925-945.

[9] Itti, L., C. Koch, and E. Niebur. "A model of saliency-based visual attention for rapid scene analysis." *IEEE Transactions on Pattern Analysis and Machine Intelligence* 20, no. 11 (1998): 1254-1259.

[10] Parasuraman, R. and G.F. Wilson. "Putting the brain to work: neuroergonomics past, present, and future." *Human Factors* 50, no. 3 (2008): 468-474.

[11] Schalk, G., K.J. Miller, N.R. Anderson, J.A. Wilson, M.D Smyth, J.G. Ojemann, D.W. Moran, J.R. Wolpaw, and E.C. Leuthardt. "Two-dimensional movement control using electrocorticographic signals in humans". *Journal of Neural Engineering* 5, no. 1 (2008): 75-84.

[12] Velliste, M., S. Perel, M.C. Spalding, A.S. Whitford, and A.B. Schwartz. "Cortical control of a prosthetic arm for self-feeding." *Nature* 453, no. 7198, (2008): 1098-1101.

[13] Wolpaw, J.R. and D.J. McFarland. "Control of a two-dimensional movement signal by a noninvasive brain-computer interface in humans." *Proceedings of the National Academy of Sciences of the United States of America* 101 (2004): 17,849-17,854.