

# Advancing Information Superiority Through Applied Neuroscience

R. Jacob Vogelstein

**I**nformation superiority, defined as the capability to collect, process, and disseminate an uninterrupted flow of information, is the pillar upon which the United States will build its future military and intelligence dominance. But the foundation upon which that pillar will be constructed is neuroscience superiority: the capability to develop new technologies based on our understanding of the brain. Neuroscience superiority is needed because although the United States' ability to collect and disseminate information has dramatically increased in recent years, its ability to process information has remained roughly constant, limited by the bandwidth of human sensory perception. Advances in neuroscience and neurotechnology afford the opportunity to correct this imbalance. Here we provide an overview of our efforts in applied neuroscience research and development and highlight some of the ways in which these advances can provide critical contributions to the nation's critical national security challenges. In particular, we review technologies that provide alternative broadband communication channels into and out of the human mind, create facsimiles of neural circuits in silico, and identify expert performers through assays of the brain.

## INTRODUCTION

Over the past few decades, neuroscientists have amassed an incredible amount of information about the brain, both in terms of its physical configuration as well as its dynamic operation. Historically, most neural research has focused on basic science or therapeutics.

With this focus, there have been significant advances in understanding the mechanisms behind specific neural functions or phenomena and in restoring or replacing neural functions lost to injury or disease. However, a new subdiscipline, which we call applied neuroscience,

has recently emerged that uses knowledge of brain structure and function to complement and augment human performance. At APL, we use applied neuroscience to create novel solutions to critical challenges faced by analysts, warfighters, and others who serve our country.

Although applications of neuroscience can be found throughout the military and intelligence community, these applications are perhaps most easily understood in the context of information superiority. As described in the U.S. Secretary of Defense's Annual Report to the President and the Congress<sup>1</sup> in 1999,

"Information superiority is the capability to collect, process, and disseminate an uninterrupted flow of information while denying an adversary's ability to do the same. It is the backbone of the Revolution in Military Affairs and provides comprehensive knowledge of the status and intentions of both adversary and friendly forces across the air, land, sea, and space components of the battlespace."

Indeed, for over a decade, the U.S. Armed Forces have invested heavily in information superiority, fielding an ever-increasing number of persistent sensors in the air, on land, at sea, and in space. The advantage of all of these sensors is that the U.S. military and intelligence communities have an unsurpassed window into the actions of our adversaries. The disadvantage is that there is now, or will soon be, much more data than there is time to analyze those data. In 2009, projecting just a few years ahead, Lt. Gen. David A. Deptula, Air Force deputy chief of staff for intelligence, surveillance and reconnaissance, said, "We're going to find ourselves in the not too distant future swimming in sensors and drowning in data."<sup>2</sup> Today, that not-too-distant future has arrived.

Put in quantitative terms, the scenario that Lt. Gen. Deptula described arose because the rate of information collection and dissemination increased faster than the rate of information interpretation. This imbalance results from the fact that the information transmission rate is limited only by available telecommunications bandwidth, whereas the information interpretation rate is limited by the bandwidth of human analysis. Because no one wants to decrease the rate of information collection, the only way to reverse this trend is to increase the efficiency with which information is processed. At APL, one of the ways that we are approaching this problem is to use applications of neuroscience.

In the context of information superiority, applied neuroscience at APL offers three main solution domains to alleviate the human information-processing bottleneck described above:

1. Brain-computer interfaces (BCIs)
2. Neuromimetic computers
3. Neural predictors of expert performance

BCIs are devices that provide high-bandwidth communication channels between man and machine, typically using noninvasive neural recording (or stimulation) technologies such as electroencephalography (EEG) or near-infrared imaging. Such devices can be used to allow humans to process data more quickly by eliminating any physical manifestation of information processing (e.g., spoken or typed words) and substituting for it a direct reading of neural activity. Neuromimetic computers attempt to emulate the structure and function of neural circuits in the brain with silicon hardware or specialized software, allowing human-like cognitive processes to be executed at computer-like speeds. Finally, neural predictors of expert performance are static or dynamic brain markers that allow for quantitative assessment of an individual's cognitive strengths for the purposes of enhancing training or optimizing tasking. Example solutions within each of these domains will be described in the following sections.

## BRAIN-COMPUTER INTERFACES

The field of BCIs has exploded in the past decade, evolving from a niche university laboratory activity to a major focus of forward-looking research agencies throughout the U.S. government<sup>3</sup> and the commercial sector (including the gaming, advertising,<sup>4</sup> and biomedical device industries). For the purposes of information superiority, the appeal of a BCI is clear: the process of ingesting, interpreting, and reporting on a piece of data is limited in many cases by the physical activities involved—reading words on a computer screen or typing on a computer keyboard. Technologies that can replace the act of reading or typing with a direct link between one's brain and a computer can significantly decrease the amount of time required to perform a given information-processing task. At APL, we are developing practical technologies and applications to achieve this vision, largely through noninvasive neural interfaces such as EEG.

The operational principle of EEG is simple: conductive electrodes in contact with the scalp transduce microvolt-scale electrical potentials created by large, simultaneously active populations of neurons in the brain.<sup>5</sup> These potentials are then recorded and processed by a computer.

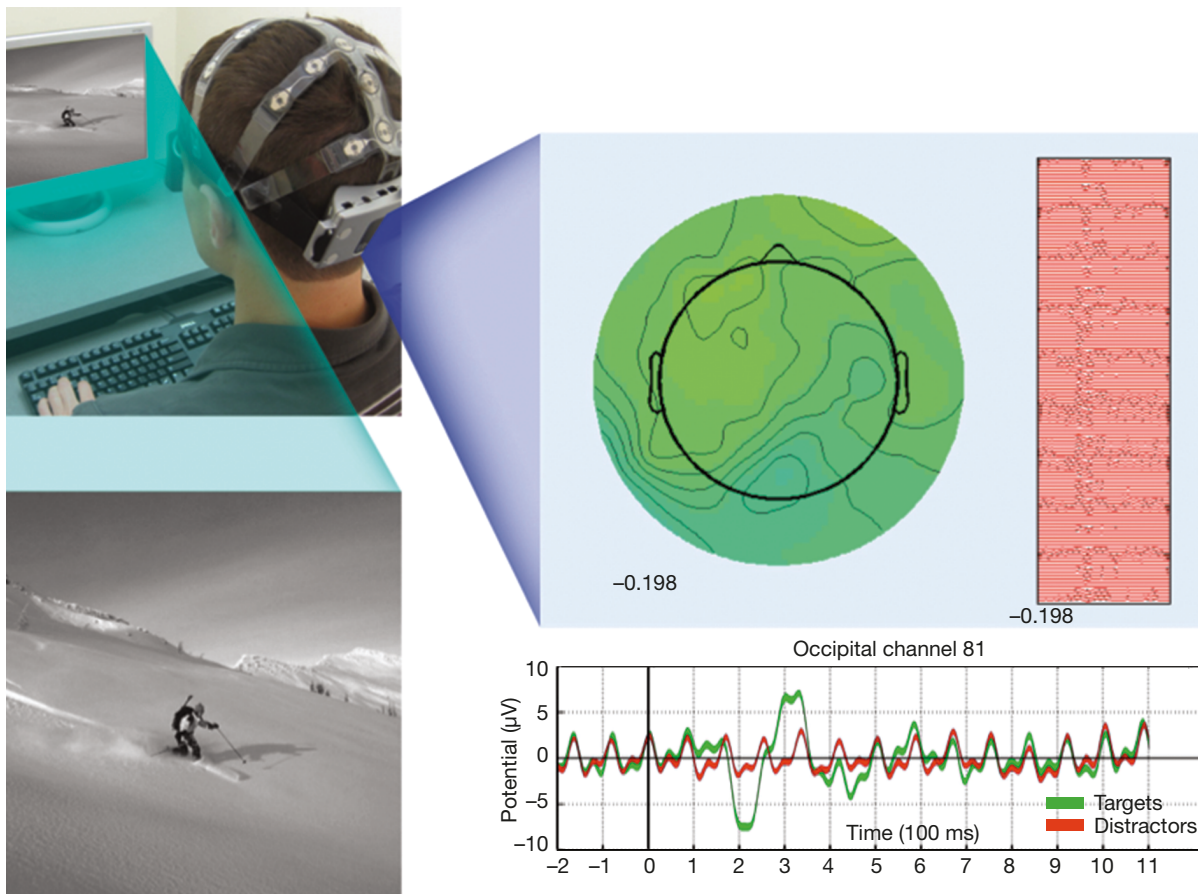
Various types of cognitive, motor, and sensory phenomena can be sensed with EEG, ranging from the brain's current "state" (e.g., alertness)<sup>6</sup> to the brain's response to a particular sensory stimulus.<sup>7</sup> Most of these phenomena have been studied with EEG for many decades; only recently has EEG been used as a real-time communication channel between a brain and a computer.<sup>3</sup> Clever designs for user interfaces, advances in machine learning algorithms, and the ready availability of high-performance computers have all contributed

to EEG’s transformation from an observational to an interactive technology. Furthermore, advances in the human factors aspects of EEG-based neural interfaces, namely the creation of relatively attractive, low-profile, and easily donned headgear (such as that from Emotiv), have made EEG the focus of many mission-relevant BCI designs.<sup>3</sup>

A particularly striking example of how EEG can be used to advance information superiority is adapted from a method called rapid serial visual presentation (RSVP).<sup>8</sup> In RSVP, a user is presented with a series of visual images (pictures) in rapid succession while the user’s neural activity is monitored by EEG. During passive visual processing in this context, the EEG signal shows a large periodic component at the fundamental frequency of the image presentations. However, if the user is instructed or chooses to look for a particular type of visual content within the image stream, the brain signals following the presentation of the “target” content will differ from their baseline representation. This difference—commonly called a “P300” due to its positive polarity and 300-ms latency relative to image presentation—is reliably detectable through

advanced signal-processing techniques and can be traced back with high precision to the particular image that induced the activity. More importantly for the purposes of BCI, because the P300 is a product of the preconscious processing of visual content, it is manifested in the absence of any physical response and can be elicited at image presentation rates far exceeding those to which a human could physically respond, up to 50 images per second.<sup>9</sup>

To put RSVP in an appropriate mission context, consider an analyst tasked with sorting through a cache of images (Fig. 1). Most of the images are likely benign pictures of family and friends, but a small fraction may have some intelligence value. Using traditional computer software, it would take this analyst a second or two to categorize each image as a “target” or “nontarget,” which quickly adds up to hours of work for thousands of images. What neuroscience tells us is that most of this time is spent on physical activities (e.g., hitting a key on the keyboard or clicking the mouse)—the brain actually consumes visual information and processes it over the course of a few hundred milliseconds,<sup>9</sup> so the rest is wasted time from an information processing perspec-



**Figure 1.** Mock-up of an image analyst using an RSVP-based BCI to rapidly triage images with potential intelligence value. [ERRATUM: The elements and data in the upper and lower right quadrants of Fig. 1 were produced by J. G. Martin and M. Riesenhuber at Georgetown University Medical Center (used with permission).]

tive. In contrast to the standard approach, by using a 50-Hz RSVP image-processing paradigm with a BCI, an analyst can “process” 3000 images per minute. That is, by using a computer to decode the activity from one’s brain (as measured by EEG) in real time, it is possible to achieve a 100× increase<sup>8,9</sup> in the rate of human information processing!

We are currently developing RSVP-based BCI technologies for a variety of potential customers ranging from signals intelligence analysts to cargo screeners at border crossings. Although we do not expect these systems to provide a 100× increase in throughput in typical work conditions, even a 2× increase could halve the time or number of personnel required to perform a particular task, providing significant savings in time and/or money. And RSVP is not the only BCI technology that we are developing—we are also actively pursuing neural interfaces for monitoring and enhancing situational awareness, controlling head-mounted displays, and creating hybrid man–machine classifiers for speech and language, among other things. Each of these systems has the potential to revolutionize the way that humans and computers exchange information and therefore revolutionize the capabilities of analysts, warfighters, and others who serve our country.

## NEUROMIMETIC COMPUTERS

In contrast to BCIs, in which the human brain is used as the information-processing engine for a computing system, neuromimetic computers attempt to replicate the function of the human brain *in silico* through hardware and/or software simulations.<sup>10</sup> The overall motivation for designing computers that operate like the brain is that despite significant advances in algorithm design, humans still outperform computers in a large variety of tasks, especially those that require interpretation of complex sensory data. Thus, if a computer can perform operations sufficiently similar to those performed in the brain, it is possible that the resulting system could have the best of both worlds: human-like cognition at computer-like speeds. At APL, we are striving to achieve this goal by developing a number of neuromimetic algorithms and neuromorphic computer hardware systems, primarily for the purposes of automated image interpretation.

Image interpretation is a traditionally hard problem for computers.<sup>11</sup> Indeed, even the most basic image-interpretation task—segmenting and labeling objects depicted in an image—can in many cases be performed more effectively by a toddler than a supercomputer. For this reason, there has long been interest in better understanding the neural pathways and mechanisms supporting human visual processing, with the hope that this understanding could be translated into algorithms for machine use.

To date, the most prominent synthesis of the human visual processing pathways in the brain is the so-called HMAX model of human vision.<sup>12</sup> HMAX uses a hierarchical array of matched filters to generate a biologically based feature set into which all images are decomposed; classification of image content is based on the output from a classifier that accepts this feature set as an input. After “training” the algorithm by presenting it with tens or hundreds of examples of images containing a particular target object, an HMAX instance can effectively identify other exemplars of the same object type at various positions, scales, and rotations in novel (previously unseen) images.

For the past few years, we have been making improvements and embellishments to the baseline HMAX model described in the literature and evaluating its performance on mission-relevant datasets for our sponsors in the intelligence community. The basic concept of operation for HMAX (or any object-recognition algorithm in this context) is simple and similar to that described above for RSVP-based BCI systems. Briefly, given a large database of unlabeled images with unknown intelligence value, we deploy one or more instances of neuromimetic object-recognition algorithms and automatically tag (i.e., label) any image containing a user-defined “target” object or scene. An analyst can then query the database based on the applied tags and selectively inspect those images containing objects of interest. Although it is difficult to conclude that any one algorithm is uniformly better than all the others in all conditions, in our experience (and in the experience of others<sup>13</sup>), HMAX consistently outperforms state-of-the-art conventional computer vision algorithms.

Despite our success with HMAX and other neuromimetic algorithms (the algorithm called Map-Seeking Circuits, or MSC,<sup>14</sup> has been particularly effective on a specific subset of object-recognition problems), none of the algorithms that we have developed or tested to date achieves human-like performance in complex image-interpretation tasks. This is undoubtedly caused by a number of factors, not least of which is the fact that all existing neuromimetic algorithms are abstract representations of the biological circuits that they intend to mimic. This abstraction is in part due to mathematical convenience but mostly due to a fundamental lack of information about the organization of human neural circuits at the level of individual neurons. Presumably, if our software and hardware more closely resembled our native wetware (i.e., our biological processing systems), we would observe more equitable performance between the physiological and artificial systems. To this end, we have recently begun an effort to reverse-engineer neural circuits in the visual cortex.

Historically, most of our knowledge about biological neural networks is extrapolated from small numbers



of observations in a limited area of the brain. This is primarily due to technological limitations—individual connections between neurons are small (submicron), but a single neuron can easily make thousands of connections across many millimeters or centimeters, posing a significant measurement challenge. However, within the past few years, techniques based on serial electron microscopy have proven that high-resolution (nanometer-scale) measurements can effectively be made over large (millimeter-scale) regions of cortex, opening the door for a revolution in understanding brain function.<sup>13</sup> At APL, we are at the forefront of this revolution, working with academic collaborators across the country to interpret the terabytes of data produced by this method and virtually reconstruct the complex neural networks captured by the electron microscopy images.<sup>15,16</sup> Although these efforts are nascent, we expect this work to pay significant dividends in the future as we advance the biological fidelity of our neuromimetic designs and come closer to replicating human neural functions *in silico*.

## NEURAL PREDICTORS OF EXPERT PERFORMANCE

The previous two sections described two ways in which applications of neuroscience can advance our capabilities in information superiority: (i) by increasing the rate at which a given person can process information by creating a direct link to the brain, and (ii) by increasing the ability of a computer to process information by mimicking neural functions. In this section we describe a third way to apply neuroscience to challenges in information superiority, which is to identify functional and structural characteristics of the brain associated with expert performance on information-processing tasks and then use this knowledge to advance training or selection of personnel who work on information superiority tasks. In other words, whereas the previous two sections are about improving existing human or physical computing resources, this section is about identifying and training new human resources.

It is generally accepted that cognitive functions such as intelligence and creativity emerge from the activity of distributed neural networks in the brain. Until recently, high-resolution measurement of this activity has been impossible, so cognitive assessments are most often done using proxy measures such as performance on standardized tests. In the military and intelligence communities, these test scores are used to assign roles and responsibilities, though it has often been shown that standardized test scores are poorly correlated with success at any particular task.<sup>17</sup> With the advent of high-resolution structural, functional, and diffusion magnetic resonance (MR) imaging of the brain, a new opportunity has emerged to create quantitative metrics of cognitive


capabilities through direct measurement of the neural substrate. At APL, we are using these measurements to predict multiple parameters of an individual's cognitive capabilities (or susceptibilities), with the ultimate goal of improving return on investment in training and more effective use of human resources.

The specific metrics that we are developing are derived from an individual's MR connectome, a comprehensive description of the neural networks in one's brain as measured by MR imaging.<sup>18</sup> We postulate that individual differences in MR connectomes underlie differences in cognitive capabilities, and, similarly, that individual differences in task performance can be predicted from one's MR connectome. These predictions are based on similarity scores between the observed structural and functional neural networks in an individual's brain and the network motifs expressed most commonly in individuals who are perceived as experts in a particular domain. That is, by measuring the MR connectomes from expert performers (and less competent performers, for comparison), we can assign performance probabilities to prospective performers based on similarities to the observed templates. These probabilities can then be used to optimally align personnel selection and training with capabilities.

To date, we have developed two enabling technologies toward the goal of realizing MR connectome-based cognitive assessments. The first is an automated pipeline that converts multimodal MR imaging data into a mathematical representation of a connectome.<sup>19</sup> The second is a set of statistical graph theory-based techniques for assigning MR connectomes a class value according to their similarity (or difference) to two distinct cognitive classes.<sup>20</sup>

The Magnetic Resonance Connectome Automated Pipeline (MRCAP; Fig. 2) takes as input a combination of diffusion-weighted MR images and structural MR images to generate an MR connectome derived from connectivity measurements between anatomically defined cortical regions.<sup>19</sup> The connectome is quantified by a connectivity matrix suitable for input to graph theoretic or statistical algorithms that can infer meaning from the data (Fig. 2, middle). In this representation, the rows and columns represent different regions of the brain (e.g., left medial orbital frontal cortex, inferior temporal cortex, etc.), and entries in the matrix represent the strength of connection between the regions indicated by each row and column combination. Each brain is summarized in a single matrix, and brains with similar network connectivity patterns will have similar matrices.

Because of the high dimensionality of a brain connectivity matrix (2415 dimensions in our formulation), it is difficult to infer meaning by manual inspection. Consequently, we have developed a set of graph theoretic algorithms that can effectively characterize the




**Figure 2.** Graphical illustration of MRCAP. MRCAP combines diffusion-weighted images (DTI) with structural MR images (MPRAGE) to generate an MR connectome derived from connectivity measurements between anatomically defined cortical regions. The latest stable release of MRCAP is available for download from NITRC.<sup>19</sup>

key differences between two groups of brain matrices and then determine whether a given brain matrix is more similar to the first group or the second.<sup>20</sup> For example, Fig. 3 illustrates a particular set of 38 connections between brain regions that has proven useful at discriminating sex based on brain connectivity. The most interesting feature of this finding is that no single connection is as predictive as the combination of multiple connections—this implies that traditional univariate analyses are underestimating the amount of information available in the data, and that using a richer representation (e.g., representing brain networks as graphs) could be valuable for myriad investigations of brain function.

The combination of an MR image-processing pipeline and sophisticated graph-based classification algorithms forms a comprehensive set of tools to predict brain function from brain structure. The tools have been validated on test data and are currently being evaluated for their performance in discriminating between multiple mission-relevant cohorts such as individuals who are likely to achieve high proficiency in foreign languages and those who are not. Moreover, these same techniques can be used for clinical applications such as predicting susceptibility to psychological impairments such as post-traumatic stress disorder or prescribing treatment for otherwise poorly characterized psychological “spectrum” disorders.<sup>21</sup>

## CONCLUSION

An increasing ability to acquire and disseminate information promises to provide the United States with a strategic advantage in current and future military conflicts. However, information alone is insufficient to confer this benefit—what is needed is a concurrent increase in our ability to process and interpret the information collected. Although great strides are being made in automated machine learning and information fusion, it seems unlikely that the next few years will yield man-made solutions that exceed the performance of the human brain, which has been optimized for information processing over the course of millions of years of evolution. The promise of applied neuroscience is that wholly artificial solutions need not be the only solution space—there are now alternatives based on innovative ways to interact with, optimize the use of, or replicate the function of the human brain. At APL, we are beginning to realize these alternate approaches by creating novel BCIs that provide high-bandwidth communication channels between man and machine, by implementing detailed models of cortical circuits *in silico*, and by designing new brain-based assays to predict cognitive performance and optimize training and selection pro-



**Figure 3.** The 38 most significant edges that discriminate male from female brains in our test data. Within this set, there are 28 interhemispheric and 10 intrahemispheric connections. The 43 brain regions from which these connections originate and terminate are colored according to their frequency of occurrence; the left hemisphere's lingual gyrus is most prevalent. Colors progress from blue (gyrus not used) to red (gyrus used frequently) in both panels, although the colormap scales are different to best visualize the underlying information. In the left panel, red indicates a maximum value of five occurrences; in the right panel, red indicates a maximum value of one occurrence.

cesses. Our hope is that these and other innovative neurotechnology solutions will allow us to take full advantage of the myriad sensors deployed throughout physical space and cyberspace and to stay afloat on the deluge of data rather than drowning in it.

#### REFERENCES

<sup>1</sup>Cohen, W. S., *Annual Report to the President and the Congress*, Government Printing Office, Washington DC, [http://www.fas.org/man/docs/adr\\_00/index.html](http://www.fas.org/man/docs/adr_00/index.html) (1999).

<sup>2</sup>Magnuson, S., "Military 'Swimming in Sensors and Drowning in Data,'" *National Defense*, <http://www.nationaldefensemagazine.org/archive/2010/January/Pages/Military'SwimmingInSensorsandDrowninginData.aspx> (Jan 2010).

<sup>3</sup>Brunner, P., Bianchi, L., Guger, C., Cincotti, F., and Schalk, G., "Current Trends in Hardware and Software for Brain-Computer Interfaces (BCIs)," *J. Neural Eng.* 8(2), 025001 (2011).

<sup>4</sup>Ariely, D., and Berns, G. S., "Neuromarketing: the Hope and Hype of Neuroimaging in Business," *Nat. Rev. Neurosci.* 11(4), 284-292 (2010).

<sup>5</sup>Nunez, P. L., and Srinivasan, R., *Electric Fields of the Brain: The Neurophysics of EEG*, 2nd Ed., Oxford University Press, New York (2005).

- <sup>6</sup>John, M. S., Kobus, D. A., Morrison, J. G., and Schmorow, D., "Overview of the DARPA Augmented Cognition Technical Integration Experiment," *Proc. First International Conf. on Augmented Cognition*, Las Vegas, NV, pp. 131–149 (2005).
- <sup>7</sup>Bressler, S. L., and Ding, M., "Event-Related Potentials," in *Wiley Encyclopedia of Biomedical Engineering*, John Wiley & Sons, Inc. (2006).
- <sup>8</sup>Gerson, A. D., Parra, L. C., and Sajda, P., "Cortically Coupled Computer Vision for Rapid Image Search," *IEEE Trans. Neural Syst. Rehabil. Eng.*, **14**(2), 174–179 (2006).
- <sup>9</sup>Johnson, J. S., and Olshausen, B., "Timecourse of Neural Signatures of Object Recognition," *J. Vis.* **3**(7), 499–512 (2003).
- <sup>10</sup>Mead, C., *Analog VLSI and Neural Systems*, Addison Wesley Publishing Company (1989).
- <sup>11</sup>Caelli, T., and Bischof, W. F., *Machine Learning and Image Interpretation (Advances in Computer Vision and Machine Intelligence)*, Plenum Press, New York (1997).
- <sup>12</sup>Serre, T., Wolf, L., Bileschi, S., Riesenhuber, M., and Poggio, T., "Robust Object Recognition with Cortex-like Mechanisms," *IEEE Trans. Pattern Anal. Mach. Intell.* **29**(3), 411–426 (2007).
- <sup>13</sup>Bock, D. D., Lee, W. C., Kerlin, A. M., Andermann, M. L., Hood, G., et al., "Network Anatomy and In Vivo Physiology of Visual Cortical Neurons," *Nature* **471**(7337), 177–182 (2011).
- <sup>14</sup>Arathorn, D. W., *Map-Seeking Circuits in Visual Cognition: A Computational Mechanism for Biological and Machine Vision*, Stanford University Press, Stanford, CA (2002).
- <sup>15</sup>The Open Connectome Project website, <http://www.openconnectomeproject.org> (accessed 11 Aug 2011).
- <sup>16</sup>Lichtman, J. W., Livet, J., and Sanes, J. R., "A Technicolour Approach to the Connectome," *Nat. Rev. Neurosci.* **9**(6), 417–422 (2008).
- <sup>17</sup>Tufflash, M., Roring, R. W., and Ericsson, K. A., "Expert Performance in SCRABBLE: Implications for the Study of the Structure and Acquisition of Complex Skills," *J. Exp. Psychol. Appl.* **13**(3), 270–304 (2007).
- <sup>18</sup>Hagmann, P., Cammoun, L., Gigandet, X., Gerhard, S., Grant, P. E., et al., "MR Connectomics: Principles and Challenges," *J. Neurosci. Methods* **194**(1), 34–45 (2010).
- <sup>19</sup>MR Connectome Automated Pipeline (MRCAP) project page, <http://www.nitrc.org/projects/mrcap/> (accessed 11 Aug 2011).
- <sup>20</sup>Vogelstein, J. T., Gray, W. R., Vogelstein, R. J., and Priebe, C. E., "Graph Classification using Signal Subgraphs: Applications in Statistical Connectomics," doi: arXiv:1108.1427v1, <http://arxiv.org/pdf/1108.1427v2.pdf> (2011).
- <sup>21</sup>Bullmore, E. T., and Bassett, D. S., "Brain Graphs: Graphical Models of the Human Brain Connectome," *Annu. Rev. Clin. Psychol.* **7**, 113–140 (2010).

# The Author

**R. Jacob Vogelstein** serves as a Program Manager in the Research and Exploratory Development Department at APL, where he oversees the applied neuroscience program area. Dr. Vogelstein's research focuses on neuromimetic computing, BCIs, neuroprosthetics, and other neural technologies that can help to solve critical challenges in the DoD and intelligence community. His e-mail address is [jacob.vogelstein@jhupl.edu](mailto:jacob.vogelstein@jhupl.edu).