Neurosecurity: Human Brain Electro-optical Signals as MASINT

by Dr. Matthew Canham and Dr. Ben D. Sawyer

INTRODUCTION

Applied neuroscience presently allows not only the scientific discovery-oriented probing of the inner workings of the mind, but increasingly the probing of individual minds toward gathering intelligence. Significant advances in neuroimaging, leveraging both active and passive electro-optical energy, can reveal specifics of information held in the mind even without cooperation (e.g., Lange et al., 2018; Sawyer et al., 2016a). The processes of the brain increasingly join many other energetic sources from which quantitative and qualitative data analysis may extract identifying features and other useful intelligence (Sawyer & Canham, 2019). Indeed, it is increasingly appropriate to discuss the human brain as a system which can be read from, written to, and the operations of which may therefore be collected for analysis or influenced (Sawyer & Canham, 2019). Indeed, we argue here that we are witnessing the end of the era in which human thought is generally accepted as an entirely private process, the starting point of an unquestionably remarkable transition. The collection of unintended emissions and byproducts toward intelligence fits well into the mold of Measurement and Signals Intelligence, and indeed Measurement and Signature Intelligence (both MASINT, Macartney, 2001), and so we believe this community within the Intelligence Community is well-suited to discuss these new realities of neurosecurity, as it helped shape many formative discussions surrounding cybersecurity. A MASINT perspective on biological, neural signatures comes with the need to discuss current capabilities, projected technological arc, practicalities, and potential abuses.

While these authors currently have no knowledge of remote monitoring of brain activity, multiple commercial entities are working toward this technology (Strickland, 2017) in various forms. Simultaneously, evidence of remote interference in normal brain functioning is in the news. Most recently, between December 2016 and October 2017, at least 21 employees stationed at the U.S. Embassy in Havana, Cuba, reported experiencing a constellation of symptoms usually associated with a concussion or traumatic brain injury (TBI). Eighteen of these employees reported a sudden onset of symptoms coinciding with an intense chirping or ringing sound similar to the Indies short-tailed cricket. Symptoms reported by employees included difficulty hearing, dizziness, headaches, cognitive difficulties, difficulties with balance, and intense brain pressure (Kirk, 2019). A clinical evaluation by researchers at the University of Pennsylvania found structural differences between exposed employees and healthy controls (Verma et al., 2019). While the clinical implications of this are currently unclear, it seems plausible that these employees were exposed to something that altered their neurological structures and cognitive functioning. The mystery continued to deepen in 2018 when an embassy employee stationed in Guangzhou, China, reported similar symptoms. While we stress that there is still considerable mystery surrounding these events, it does seem likely that these symptoms were (1) induced and (2) likely not the direct goal of whatever process produced the phenomenon. Initial examination of the victims suggests remote microwave energy, long known to affect temporal lobe function (Dyer, 2018). These phenomena provide potential evidence of the intentional targeting of neural architecture, potentially as an attack, potentially as a side effect to some other goal.

Less circumspect evidence also exists. Capability to monitor neural activity exists given direct physical proximity, and remote neural monitoring may be feasible. Recent advances have seen remote detection of other biosignals once considered only measurable from direct physical proximity. For example, NASA's Finding Individuals for Disaster and Emergency Response (FINDER) system uses low-power microwaves to detect heartbeats at great physical range (Liu et al., 2014). Core body temperature is now routinely monitored in crowds to identify individuals with infections (Ng, Kawb, & Chang, 2004). Moreover, two categories of neuroimaging technology are emerging with the promise to make remote brain access a near-term reality. Industry groups like Facebook and Open Water are working to advance near-infrared and holographic techniques for monitoring neural blood flow patterns in real time (Open Water, 2018).

Meanwhile, Neuralink, Kernel, and others are working to connect the electrical activity of the brain to intermediary electrodes, and then to the Internet. The success of either of these technologies, neuroimaging at range or Internet-connected electroencephalography, will open a new universe of possibilities for the realms of MASINT, SIGINT, and HUMINT alike.

CURRENT STATE OF THE ART

Before diving into the world of neuroimaging, we offer a brief introduction into what is currently known about how the brain functions. We begin with the neuron, the basic building block of the neural network that is our brain. A basic decision-making system, it takes in input from upstream neurons through receptors known as dendrites and, once a certain threshold of these signals is met, "fires" an action potential which travels down the long synapse to the synaptic gap which separates one neuron from another. Here, chemical signals take over, propagating further action potentials downstream to other neurons in spreading cascades of activity and activation. The process is a foundation for complex patterns of information being aggregated and processed. For example, while the earliest neurons to process visual information might only detect the presence or absence of an edge, neurons further downstream in visual cortex will aggregate the presence of an edge in a specific orientation or relative position and recognize this as the letter "K." Further downstream, neurons will respond more vigorously to the letter "K" when it is placed at the beginning of a word as opposed to the middle or end. In this way, information is aggregated and processed into meaningful coherence.

While there is still debate surrounding the validity of brain area specialization, and growing evidence for "network" approaches to understanding activity, at a coarse level, brain regions appear to be functionally specialized for different activities. Understanding this differential specialization allows for a limited, but growing, degree of reverse engineering of brain processes. A great deal of cognitive processing occurs in the neocortex, the outermost layer of the brain. Here, four "lobes," anatomical brain regions, have been linked by research to functional specializations (see Figure 1, Miller & Cummings, 2017). The occipital lobe, or visual cortex, is where much of visual processing takes place. The parietal lobe handles spatial awareness and somatosensory processes which feed the brain's sense of bodily positioning and stimulation. For example, tickling the hands or feet with a feather would activate somatosensory processing, which would occur primarily in the frontal parietal lobe. The temporal lobe also sits just forward of the occipital lobe and below the parietal lobe, usually just above one's ear. The temporal lobe (aka the auditory cortex) processes sound and often handles long-term memory processing as well. Finally, the frontal lobe is responsible for fine motor functioning, and actions known as executive functions: deliberate decision-making, inhibitory control, attention, and working memory. If you are intensely concentrating on a task, then there is a high likelihood that you are recruiting much of your frontal lobe's prefrontal cortex. This final example is especially significant from a MASINT perspective: it has been suggested that when deliberately trying to deceive someone, the deceiver relies on his/her frontal lobe to a greater degree than does someone who is not attempting to be deceptive **(**Zeki et al., 2004**)**. There is greater activation in the prefrontal cortex because the individual must inhibit the true version events and must hold two versions active simultaneously **(**Ofen et al., 2016**)**. Although there is still much debate on the validity of this assertion, as an example it illustrates how neural processing might be utilized in an intelligence-gathering capacity.

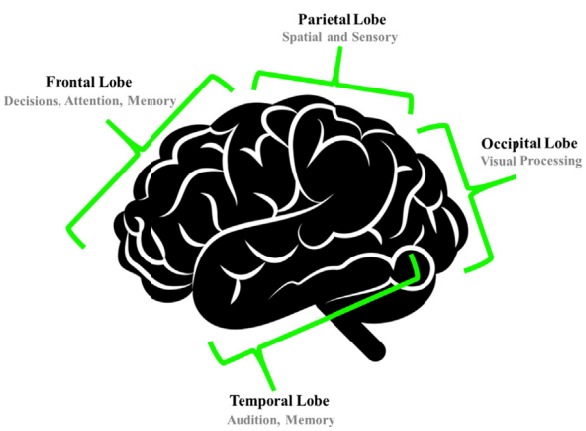

Figure 1: The neocortex or surface of the brain, disproportionately responsible for cognitive processing, is currently conceptualized as divided into functional regions. As with technical and social systems, useful MASINT consideration of these areas is in terms of intelligence and potential influence. Increasingly, it is possible to collect electro-optical energy emitted by the brain and, leveraging temporal and spatial dimensions, decode meaning and so acquire useful intelligence. Influence is also possible, and devices which project electrical force into the brain can disrupt or modify brain processes.

Detectable Signals – A discussion about neuro-imagining should first make the distinction between structural and functional imaging. Structural imaging provides a highly detailed static image of the neuro-

anatomical structures of an individual. When the researchers from the University of Pennsylvania examined the embassy employees and found differences in whole brain white matter, this difference was found through the analysis of static structural images (Verma et al., 2019). In contrast, functional imagery tends to be coarser but provides a dynamic series of snapshots that provide insight into the neural activity of an individual. While both techniques have relevance to MASINT applications, functional imaging will be the topic of focus here. Within the universe of functional imaging there are currently two types of signals, blood flow and electrical activity, that are detected to derive neural functioning.

Blood Flow Signals – When neurons are active, these cells consume sugar and oxygen and therefore require replenishment. This replenishment transpires through a process known as hemodynamic response. Termed a blood-oxygen-level-dependent (BOLD) signal, this difference between oxygenated and deoxygenated blood is detectable through various means such as magnetic manipulation or using infrared spectrum light. Examining this signal using magnetism usually involves a technology known as functional Magnetic Resonance Imaging (fMRI). fMRI technology witnessed an upshot in usage within brain research beginning in the early 1990s because it was considerably less intrusive than comparable imaging technologies available at the time. A major drawback in fMRI as a MASINT technique is the need to immobilize a subject and capture imagery over a long time period (from 45 minutes to a few hours), while secured to a table and loaded into a magnetic resonance tube. Movement during imaging is highly detrimental, meaning that only extremely compliant individuals can be imaged. Finally, high tesla (a measurement of magnetism strength) equipment capable of high spatial and temporal resolution imaging is extremely expensive and often requires a dedicated staff, making this technology largely confined to use within a dedicated laboratory. These inconveniences notwithstanding, several researchers have proposed methods of employing fMRI as a means of deception detection (Ganis et al., 2003; Kozel et al., 2005; Monteleone et al., 2009; Ganis et al., 2011). Continuing advances in the miniaturization of this technology suggest this could eventually be an approach moved out of the laboratory and into the field (see, for example, Cooley et al., 2015).

Other emerging techniques such as functional Near Infrared Spectroscopy (fNIRS) offer a window into more near-term workable solutions. Cheap, low-power, and portable, fNIRS utilizes the near infrared spectrum light to detect the BOLD signal. In the 700-900nm spectral range, bodily tissues are mostly transparent, allowing maximal detectability of the relative difference between oxygenated and deoxygenated hemoglobin. fNIRS utilizes a combination of infrared light emitters and receivers to parse out the BOLD signal through differences in infrared light intensity. These differences in light intensity can then be interpreted to detect and localize BOLD signals from specific brain regions to infer localized activity. One of the major advantages of fNIRS over fMRI from a MASINT perspective is the ease of use, and portability of these devices. Indeed, the technology is routinely held up as an excellent match for the demands of brain machine interface and field research (respective reviews are Naseer & Hong, 2015 and Quaresima and Ferrari, 2019). It is currently unclear what the ultimate detectable range using the infrared spectrum will be, but at present these signals are detected using a sensor cap worn by the subject which directly contacts the skin. This portability and ease of use would potentially allow for modern deployment in the debriefing of HUMINT assets by handlers or operational psychologists.

Electrical Activity Signals – While neuroimaging techniques dependent upon blood flow offer high spatial resolution and the capability of localizing neural activity, they lack the capability of detecting activity with a high temporal resolution because there is an inherent lag in the reuptake of oxygenated hemoglobin into active neural regions. This delay means that events which happen very quickly, such as visual recognition, can be missed by techniques reliant on BOLD signal. In these situations, techniques that detect electrical activity offer an advantage over those that detect signals related to blood flow. Electrical detection techniques have very high temporal resolution (on the order of milliseconds), but because electrical fields are distorted by the scalp, they lack the spatial resolution that blood flow-based imaging techniques have. Therefore, researchers often combine these techniques when studying neuro phenomena.

Techniques measuring electrical activity include deep brain electrodes, Electrocorticography (ECoG), and Electroencephalography (EEG), listed from most to least invasive. Brain-contact techniques utilize small probes (approximately 5 ìm thick) to directly connect to neurons to detect activity (Muthuswamy, 2012), and involve opening the skull to access the cortex. ECoG is somewhat less invasive, involving electrodes that rest upon the dura, a thin sheet of enervated tissue which contains the cerebrospinal fluid and the brain. Non-invasive techniques such as EEG detect voltage potential fluctuations deriving from the action potential activity within the neurons of the brain. Such measured "potentials" can be measured longitudinally over time, or measured relative to specific events, an approach which can identify specific patterns of brain activity known as event-related potentials (ERP). This connection between

outside events and brain activity is an excellent strategy to reverse engineer (to a limited degree) the brain activity as it relates to a specific stimulus. One of the most studied ERPs, the "P300" wave, is a distinctive positive fluctuation that occurs approximately 300 milliseconds after visual recognition of a stimulus. The P300 has therefore been proposed as a deception detection technique in "guilty knowledge tests." A subject wearing an EEG would, in such a test, be presented with visual stimuli in succession, and an amplified P300 of what occurred directly after any image recognized, and without the awareness or conscious control of the subject. Many other potentially useful ERPs exist, in the context of MASINT, and include error-related negativity (ERN, see Sawyer et al., 2016b), the P3 (see Rosenfeld et al., 1991), and ERN composite signals such as the multifaceted electroencephalographic response (MERMER, see Farwell & Smith, 2001), to name but a few. Indeed, while the present literature is focused upon individual signatures and their functional meeting, the overarching message here from a MASINT perspective is that electrical signals collected incidentally from brain activity can be used to provide actionable intelligence.

Directing Input into the Brain – Thus far our discussion has centered around reading activity from the brain, but electromagnetic energy can also be effectively used to input information into the brain. A delicate system, the brain can be influenced or disrupted by relatively small amounts of kinetic or electrical energy, and indeed is susceptible to informational patterns (Sawyer et al., 2016a; Sawyer & Hancock, 2018) Transcranial Magnetic Stimulation (TMS) is one such technology, and uses magnetic energy directed toward the neocortex either to excite or to suppress the underlying neural region. For example, an individual who has their visual cortex (occipital lobe, see Figure 1) may experience loss or aberration of vision. TMS has been used for decades in both clinical and research contexts. Recent applications of this technology are striking: for example, a research group at the University of Washington (Jiang et al., 2019) employed TMS as part of an "artificial telepathy" apparatus. In this experiment, two subjects (the senders) watched the orientation of Tetris-like pieces and focused on whether the piece should be rotated to align its placement. A third subject (the receiver), located in a different room and unable to see the pieces, was tasked with deciding whether to rotate the piece. The receiver performed well above chance (~81% accuracy) in deciding whether the piece needed to be rotated, based completely upon the signal he received from the senders. This suggests that beyond collecting actionable intelligence, there are presently ever-increasing opportunities for near engineering, potentially for influence or projecting force.

BRAIN MACHINE INTERFACES INTRODUCE NEW ATTACK SURFACES

Significant progress has been made in recent years in the development of both invasive and non-invasive Brain-Machine Interfaces (BMIs), allowing operators to communicate directly with machinery (computers, robotics, cars, artificial limbs, etc.) using only their thoughts (Roelfsema et al., 2018). A quick patent search reveals that over 3,800 patents were filed for such technology in 2018 (Google Patents, 2019). The intimate connection between the operator's brain and the controlled device opens an entirely new dimension of attack surfaces to be exploited by cyber threat actors. Information security primarily rests upon three pillars: Confidentiality (preventing unauthorized disclosure of information), Integrity (preventing unauthorized modification of information), and Availability (maintaining access to information), the so-called CIA Triangle (Wiley, 2008). Within the context of neuro-security a breach of Confidentiality could potentially allow unprecedented access to an individual's most private data, his/her thoughts. A breach of Integrity would mean that an attacker could inject commands into a neuro-device, or alternatively send false feedback to the brain from the device. A failure of Availability would prevent a user from being able to control the device or receive data from it. The failures of any of these pillars might seem to be purely within the realm of science fiction; however, proof of concept attacks have already been demonstrated for each.

Reaching into the uncooperative individual's mind to retrieve, or influence, information is increasingly a reality. Lange et al. (2018) were able to recover partial Personal Identification Numbers (PINs) from subjects' EEG (electroencephalogram) signal. Other research (Roelfsema et al., 2018) has demonstrated the ability to infer the words or concepts that an individual is thinking of, from EEG signals. Without the proper security, individuals using BMIs relying on similar signal processing would be subject to having their private thoughts exposed. Perhaps more disconcerting than breaching Confidentiality is a breach of Integrity; such a breach was demonstrated by Cusack et al., 2017 in a highly controlled environment. In this study, researchers conducted a Man-In-The-Middle attack against a BMI and a toy car and were able to intercept thought-based commands from the user's BMI and inject modified commands. In this case they substituted the command "turn left" with "turn right." If such an attack were launched against an artificial limb or a wheelchair (both of which can now be controlled with similar technology), an attacker could easily cause death or serious physical injury either the user or those around them. In a similar vein, Cusack et al. (2017) describe a simple modification to their integrity-focused attack of flooding the BMI connection with meaningless packets to disrupt the control channel and thereby deny the operator

access to the controlled device. This type of attack, properly timed, could lead to equally destructive results if the downstream device is the artificial limb or wheelchair mentioned above.

FUTURE DIRECTIONS AND NEUROSECURITY CONCERNS

The current state of the art in neuroimaging requires that sensors be placed in very close proximity to a subject's cranium, a state of affairs that many other energetic MASINT sources once shared. Could technology be someday (or presently) capable of detecting neural signals from a distance? While the signal detection difficulties of such a system are great, it is within the realm of possibility. Even such a technology with very limited range would have serious implications for the Intelligence Community, and open the door to covert neuro-surveillance. A few inches might allow an apparatus to be embedded in surfaces, such as seating. A few meters would allow for neuro-surveillance of an interview at a border crossing. More range comes with more interesting, and concerning, implications.

What about individuals who choose to use technology to project their neural information outward? Neuralink, and other industry actors, have this possibility as a direct piece of their value proposition. The idea of computer network-connected brains mirrors that of other computer network-connected sensors: surveillance becomes implicit in return for convenience. Indeed, it may be useful to consider the fact that surveillance capabilities of a covert microphone and a present generation household smart speaker are functionally very little. Covert or overt monitoring of neural activity holds many parallel possibilities, and Biafra mentioned remote neuroimaging is joined by technologies which will intentionally transmit neural information over the Internet, or other networks. It is extremely likely that industry and state actors, in the absence of legislative restraint, will find reason and avenues to collect and leverage such data. The rights of individuals to their own personal neural information, when transported through computer networks, is likely in the process of being decided presently by society and the courts, as rights to personal electronic information are a likely precedent.

Input, as discussed above, is another fascinating dimension of networked neural implants. The ideas are not radical, and indeed Apple and Nucleus, manufacturer of cochlear implants, recently made iOS the operating system connecting to more human implants than any other. These technologies join other apps which can be used to connect to a variety of human implants. In cochlear implants, for example, the intended mode of input is digital audio signals: it is better to listen to your phone call when beamed directly to your implant them through a microphone facing the phone speaker. However, these devices offer opportunities for MASINT, and for influence. Indeed, just as personal information and computer networks can be used for both surveillance and influence, it may be possible to manipulate overtly or covertly a target through an active neural, sensory nerves, or peripheral nervous system connection.

Consider a concerted effort to expose a subject to positive or negative stimulation in response to specific actions. Such a campaign would certainly result in some level of conditioning. We can, for example, imagine creating incentives not to enter a geo-fenced location, or not to leave one. Threat actors with the goal of rendering a target ineffective in their current occupation might leverage a cochlear implant to arrange for painful, annoying, disturbing, or other negative stimulus to be inflicted whenever the target entered their office. They could also simply degrade the quality of the function of the device. Because cochlear implants connect to the Internet through iPhones, this could be accomplished through the malicious employment of code. Note that such an attack would leverage intelligence about use location from the phone, and use the same phone to send negative stimuli to the target through the cochlear implants. Of course, cochlear implants in the United States presently all have removable external units, and could simply be removed. Submitting to deafness in order to remove the stimuli a denial in its own right, it is worth considering that such a scheme would work on other implants, each with its own uncomfortable set of possibilities.

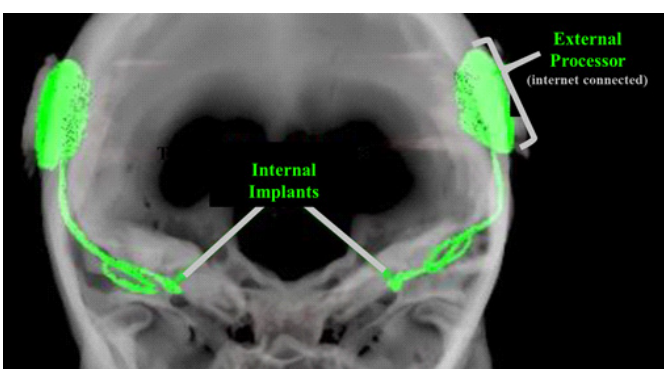

Figure 2: Modern cochlear implants are now compatible with Apple's iOS, which has therefore become a new and widely available attack surface for individuals with this type of sensory nerve-connected prosthesis. Neurosecurity questions exist regarding which central, sensory, or peripheral nervous system-connected devices will soon also be Internet-connected, and whether these have input or output capabilities.

The implications of direct and potential remote neuroimaging are, course, not limited to intelligence, nor to influence, nor to negative outcomes. Neuroimaging, especially remotely,

might prove a particularly robust new form of biometrics, through the recording of an individual's neural responses to specific stimuli, using amenable ERPs, for example. The possibilities for industry, health, and human computer interface are monumental. Interpersonal communication might be revolutionized, or at least improved. However, we believe that this hopeful narrative must be tempered with understanding of the implications to individual and aggregate security. Major questions exist, and at present there are no answers.

OPEN QUESTIONS AND CONCERNS SURROUNDING NEUROIMAGING AS A MASINT SOURCE

In sum, applied neuroscience techniques previously reserved for experts probing scientific questions are now increasingly amenable to MASINT. There are presently multiple scenarios in which intelligence can be gathered through passive monitoring of the electro-optical signals concurrent with brain activity (blood flow and neural discharge patterns), and in the near future such access may become available at greater physical distance. These opportunities are joined by rapid advancements in understanding of the functional organization and temporal signaling of the brain, coupled with rapid advancement in occupational power and machine learning technique quite familiar to the MASINT community. The result is the beginning of an era in which neural information, and the machinations of the human brain, are joining many other systems previously made amenable to MASINT information-gathering approaches. Indeed, the impacts of these combined advances are undoubtedly fueling scattered conversation and innovation in the public and classified spheres of many countries. While some outcomes will be undeniably positive, we feel that there are strong signs that a more focused conversation needs to be held.

Recently, several U.S. embassy workers stationed at Guangzhou, China, have reported symptoms like those reported by U.S. embassy workers stationed in Cuba. Again, there is much controversy surrounding these reports. One widely held assumption is that these are in fact the result of some type of "neuro-attack." Perplexing problems now arise. How could such an attack be detected? Every time your brain forms a new memory (which happens constantly), your brain changes in subtle and poorly understood ways. This constant change makes baselining incredibly challenging, and there remains some question as to whether this is even possible. Moreover, it seems likely that an "input"-based technology, as may be the cause, would be infinitely more detectable than a technology monitoring output. It seems evident that neuroimaging technology holds great potential for MASINT, and for this reason alone there is the likelihood that state-sponsored intelligence services will attempt to employ this technology as an intelligence-gathering technique. The high likelihood of this experimentation, and the relatively feasible nature of creating such a technology, should compel more research to be conducted on a variety of related neurosecurity topics.

It seems evident that neuroimaging technology holds great potential for MASINT, and for this reason alone there is the likelihood that state-sponsored intelligence services will attempt to employ this technology as an intelligence-gathering technique.

Beyond the fundamental question of whether neural tissue is amenable to gathering intelligence, or a likely target for projecting force, fundamental forensic questions which should be addressed by such a line of research are as follows:

How do we ensure neurosecurity? Just as cybersecurity was once poorly understood, so now is neurosecurity. We must understand which approaches are real threats, what their limitations are, and develop understanding as to how our own state, industry, and greater public population can be protected. We must also begin a dialogue in scientific, legislative, and public spheres to address how best to integrate these coming realities into our society. How do we safeguard freedom and security when the information between our ears is no longer inherently our own?

How do we detect attacks? In terms of information-gathering attacks, neurosecurity is likely to suffer from many of the same challenges as cybersecurity; by definition, a well-executed attack need leave no trace (see Hancock, Hancock & Sawyer, 2015). In terms of influence, the more difficult question is one of trust. What is possible in terms of influence, and how can we detect it? Indeed, this is the challenge of cyber-compromised computer systems which serve new masters, or have their cycles turned toward threat actor goals. How do we know when an individual has been attacked? One of the greatest challenges in the "Havana Syndrome" has been establishing whether something in fact occurred. Subjectively, patient reports align very closely (sudden onset, hearing a high-pitched chirping or ringing, difficulty concentrating and maintaining balance), but there is thus far no way to establish exposure conclusively.

Is it possible to develop a baseline? In cybersecurity, understanding of the original state of the system is vital for understanding whether an intrusion has occurred, and how the system is compromised. If a method for detecting a

neuroattack is developed, it will likely involve establishing an analogous neural baselining. The clinical evaluation of Havana Syndrome victims by researchers at the University of Pennsylvania found structural differences between exposed employees and healthy controls. Specifically, structural imaging indicated significantly decreased levels of whole brain white matter, differences in regional gray and white matter volumes, cerebellar microstructural integrity, and functional connectivity in the visuospatial and auditory subnetworks (Verma et al., 2019). While this study found differences between the exposed population and healthy controls, it was unable to demonstrate differences within patients before and after the time of exposure because there was no baseline created prior to their deployment. Another limitation of this study was that it focused on the structural aspects of the patients' neural architectures, but not their cognitive functioning. Baselining to detect a neuro-attack will likely necessitate a cognitive functioning component, perhaps involving rapid response to various stimuli. Developing a baseline of cognitive functioning will likely utilize neuroimaging, for example EEG to measure patient responses to stimuli over time. One of the greatest challenges to this will be understanding whether such baselining is even possible. The brain is incredibly plastic and changes constantly. In fact, every new memory formed causes changes within the brain. An unanswered question is what does "normal" change look like compared to "abnormal" change, and can these differences be detected? If they can be detected, is EEG the right technique, and are EEG responses to stimuli consistent over time? The few answers that presently exist come from vastly different domains in the neurosecurity threat to come.

CONCLUSION: TOWARD A MASINT UNDERSTANDING OF THE BRAIN

MASINT has existed for long enough that the community has witnessed many energetic signals moving from non-useful to pivotal. We here predict that the energetic emissions of the human brain will follow that pattern. Understanding the time frame of that change is difficult. It may take the entirety of our coming careers. It may have already happened. The cause of the Havana Syndrome remains a mystery at the time of this writing. It is also unclear whether Havana Syndrome is specifically the result of a neuro-weapon, or something entirely different. It does, however, provide the opportunity for a timely thought experiment, as the world will witness the effects of neuro-weapons in the foreseeable future. It is critical that tools and techniques be developed to detect the effects of these weapons, and to guard against them. We believe that the framework of MASINT, and the broader Intelligence Community which has such implicit interest in these ongoing developments, is an excellent place to begin this critical work.

References

Cooley, C. Z., Stockmann, J. P., Armstrong, B. D., Sarracanie, M., Lev, M. H., Rosen, M. S., & Wald, L. L. (2015). Two dimensional imaging in a lightweight portable MRI scanner without gradient coils. *Magnetic Resonance in Medicine*, *73*(2), 872-883.

Cusack, B., Sundararajan, K., & Khaleghparast, R. (2017). Neurosecurity for brainware devices.

Dyer, O. (2018). Microwave weapon caused syndrome in diplomats in Cuba, US medical team believes. *Bmj*, 362, k3848-k3848.

Farwell, L. A., & Smith, S. S. (2001). Using brain MERMER testing to detect knowledge despite efforts to conceal. *Journal of Forensic Science*, 46(1), 135-143.

Ganis, G., Kosslyn, S. M., Stose, S., Thompson, W. L., & Yurgelun-Todd, D. A. (2003). Neural correlates of different types of deception: An fMRI investigation. Cerebral cortex, 13(8), 830-836.

Ganis, G., Rosenfeld, J. P., Meixner, J., Kievit, R. A., & Schendan, H. E. (2011). Lying in the scanner: covert countermeasures disrupt deception detection by functional magnetic resonance imaging. *Neuroimage*, 55(1), 312-319.

Google Patents, (2019). Retrieved from https://patents.google.com.

Hancock, P. A., Hancock, G. and Sawyer, B. D., 2015, Cybernomics and the implications of cyber-deception. *The Ergonomist*, 537, 12-14.

Jiang, L., Stocco, A., Losey, D. M., Abernethy, J. A., Prat, C. S., & Rao, R. P. (2019). BrainNet: a multi-person brain-to-brain interface for direct collaboration between brains. *Scientific Reports*, 9(1), 6115.

Kozel, F. A., Johnson, K. A., Mu, Q., Grenesko, E. L., Laken, S. J., & George, M. S. (2005). Detecting deception using functional magnetic resonance imaging. *Biological Psychiatry*, 58(8), 605-613.

Lange, J., Massart, C., Mouraux, A., & Standaert, F. X. (2018). Side-channel attacks against the human brain: The PIN code case study (extended version). *Brain Informatics*, 5(2), 12.

Liu, L., Liu, Z., Xie, H., Barrowes, B., & Bagtzoglou, A. C. (2014). Numerical simulation of UWB impulse radar vital sign detection at an earthquake disaster site. *Ad Hoc Networks*, 13, 34-41.

Macartney, J. D. (2001). John, how should we explain MASINT? *Intelligence and the National Security Strategist: Enduring Issues and Challenges*, 170-171.

Miller, B. L., & Cummings, J. L. (eds.). (2017). *The human frontal lobes: Functions and disorders*. Guilford Publications.

Monteleone, G. T., Phan, K. L., Nusbaum, H. C., Fitzgerald, D., Irick, J. S., Fienberg, S. E., & Cacioppo, J. T. (2009). Detection of deception using fMRI: better than chance, but well below perfection. *Social Neuroscience*, 4(6), 528-538.

Muthuswamy, J., Sridharan, A., & Okandan, M. (2016). MEMS Neural Probes. *Encyclopedia of Nanotechnology*, 1993-2009.

Naseer, N., & Hong, K. S. (2015). fNIRS-based brain-computer interfaces: a review. *Frontiers in Human Neuroscience*, 9, 3.

Ng, E. Y., Kawb, G. J. L., & Chang, W. M. (2004). Analysis of IR thermal imager for mass blind fever screening. *Microvascular Research*, 68(2), 104-109.

Ofen, N., Whitfield-Gabrieli, S., Chai, X. J., Schwarzlose, R. F., & Gabrieli, J. D. (2016). Neural correlates of deception: Lying about past events and personal beliefs. Social cognitive and affective neuroscience, 12(1), 116-127.

Open Water (2018). At https://www.openwater.cc/technology. Poulsen K. (29 March 2008). Hackers assault epilepsy patients via computer. Retrieved 23 September 2019 from https://www.wired.com/2008/03/hackers-assault-epilepsy-patients-via-computer/.

Quaresima, V., & Ferrari, M. (2019). Functional near-infrared spectroscopy (fNIRS) for assessing cerebral cortex function during human behavior in natural/social situations: a concise review. *Organizational Research Methods*, 22(1), 46-68.

Roelfsema, P. R., Denys, D., & Klink, P. C. (2018). Mind reading and writing: the future of neurotechnology. *Trends in cognitive sciences*.

Rosenfeld, J. P., Angell, A., Johnson, M., & Qian, J. H. (1991). An ERP based, control question lie detector analog: Algorithms for discriminating effects within individuals' average waveforms. *Psychophysiology*, 28(3), 319-335.

Sawyer, B. D., & Canham, M., (2019) Neurosecurity: Infosec meets Brain-machine Interface. *B-Sides Las Vegas 2019*. Video at https://t.co/dGIJD9UIH4?amp=1.

Sawyer, B. D., Finomore, V. S., Funke, G., Warm, J. S., Matthews, G and Hancock, P. A., 2016a, Cyber vigilance: The human factor. *American Intelligence Journal*, 32(2), 157-165.

Sawyer, B. D., Karwowski, W., Xanthopoulos, P. and Hancock, P. A., (2016b), Detection of error-related negativity in complex visual stimuli: A new neuroergonomic arrow in the practitioner's quiver. *Ergonomics*, 1-7.

Strickland, E. (2017). Silicon valley's latest craze: Brain tech [News]. *IEEE Spectrum*, 54(7), 8-9.

Tosini, G., Ferguson, I., & Tsubota, K. (2016). Effects of blue light on the circadian system and eye physiology. *Molecular Vision*, 22, 61-68.

Wiley, J. (2008). Security Engineering: A Guide to Building Dependable Distributed Systems. 2nd edition, 239-274.

Yazdani, P. (2017). *Assessing seizure susceptibility using visual psychophysical tests* (doctoral dissertation, Newcastle University).

Dr. Matthew Canham, a psychologist and neuroscientist, is Research Professor of Cybersecurity for the Institute of Simulation and Training at the University of Central Florida in Orlando. His PhD degree in Cognition, Perception, and Cognitive Neuroscience is from the University of California, Santa Barbara. His research focuses on the human aspects of privacy and cybersecurity. Previously, Dr. Canham was a Supervisory Special Agent with the FBI, where he served in the field performing investigations of cyber-breaches, intellectual property theft, and other federal violations. Later, as manager of the Emerging Technologies Program with the Operational Technology Division based in Quantico, VA, he co-authored the FBI briefing for the incoming U.S. Presidential Cabinet on "Technology Based Threats to Law Enforcement in the 21st Century" and provided subject matter expertise to the FBI's Cyber Behavioral Analysis.

Dr. Ben D. Sawyer is Assistant Professor within Industrial Engineering and Management Systems at the University of Central Florida, and Director of the Laboratory for Autonomy-Brain Exchange (LabX). His PhD degree in Applied Experimental Psychology and MS in Industrial Engineering are from the University of Central Florida. At the Massachusetts Institute of Technology, he leveraged biosignals, big data, and machine learning to engineer models of human performance and behavior for the use of machine counterparts. At the Air Force Research Laboratory 711th Human Performance Wing's Applied Neuroscience and BATMAN Divisions, he built mathematical models of human performance for special operations, including cyber operations. His present research, and laboratory, seek to accelerate information transfer between human and autonomy. Dr. Sawyer's design recommendations are leveraged by Fortune 500 companies. His work has been covered by Forbes*, Reuters, Fast Company, and the BBC, and published in leading scientific journals.*