(U) The Human Side of the Net

March 2010

Not Releasable to the Defense Technical Information Center per DoD Directive 3200.12.

DISTRIBUTION LIMITED TO U.S. GOVERNMENT AGENCIES ONLY. THIS DOCUMENT CONTAINS NSA INFORMATION (MARCH 2010). REQUESTS FOR THIS DOCUMENT MUST BE REFERRED TO THE DIRECTOR, NSA

(b) (6)
OGA

OFFICIAL USE ONLY
May be exempt from public release under the Freedom of Information Act (5 U.S.C. 552), exemption number and category: Exemption 5h—Protection of Contractor Submitted Proposals, 10 U.S.C. 2305 (g).

Name/Org [redacted] Manager Emergent Threats Department
Date: 25 March 2010
Guidance (if applicable): DoD Directive 3200.12

(b)(7)(E)
OGA

(U//~~FOUO~~) [redacted]

Sandia Systems Assessment and Research Center Document

(U) The Human Side of the Net

(b) (6)
OGA

Prepared by: [redacted]

Sandia National Laboratories
Albuquerque, New Mexico

(b) (7) (E)
OGA

(U//~~FOUO~~) [redacted]

(U) Contents

(U) System Virtualization Technology Forecast 6

(U) Neuroscience–New Game–Changing Enabler 6

(U) Forecasting the Human-Network Synergy: What it Might Mean for People to be "More than Metaphorically" a Part of the Network 13

- (U) The Current Human-Network Interaction 15
- (U) Direct Input from the Human Brain to the Network 16
  - (U) Biometric Key Encryptions 18
  - (U) Identification: Unique Signals as a "Fingerprint" 18
  - (U) Reynard Problem: A method for Determining the Delta Between What Individuals Say and What They Really Think or Who They Really Are 18
  - (U) Early Spotting and Assessment 19
  - (U) "Brain Onion Routing" 19
  - (U) Outsourcing the Recognition Events: People as SETI Processors 21
- (U) Two-Way Interaction Between the Human Brain and the Network 21
  - (U) Altering Environmental Consequences Via Two-Way Electrodes 22

(U) Key Technical Hurdles for Neurotechnology Surprise 23

(U) Appendix A: Enabling Technologies 24

(U) Appendix B: Enabling Applications 53

(U) Acronyms and Abbreviations 89

(U) References 91

(U) Figures

Figure 1. (U//~~FOUO~~) Non-Cumulative Number of Peer-Reviewed Publications by Year that Involve Neuroscience or Cognitive Neuroscience Work 6

Figure 2. (U//~~FOUO~~) The Number of Peer-Reviewed Publications for the First 38 Years and for the Last 39 Years of Sustained Scientific Study in the Neuroscience and Cognitive Neuroscience Fields 7

(b) (7) (E)

Figure 3. (U//~~FOUO~~) [redacted] 8

Figure 4. (U//~~FOUO~~) [redacted] 9

Figure 5. (U//~~FOUO~~) [redacted] 10

Figure 6. (U//~~FOUO~~) [redacted] 11

Figure 7. (U//~~FOUO~~) Plot of Cost, Time Resolution, and Sensor Size 15

Figure 8. (U//~~FOUO~~) Information Available from the Network to the Computer....................16
Figure 9. (U//~~FOUO~~) The Amount of Information Transfer to the Network about that Individual Would Increase If the Human Brain Were Tied to the Network............17
Figure 10. (U//~~FOUO~~) Illustration of "Brain Onion Routing" Concept.................................20
Figure 11. (U//~~FOUO~~) A Two-Way Direct Connection Between the Brain and the Network..22

(U) Tables

(b) (7) (E)

Table 1. (U//~~FOUO~~) [redacted]..12
Table 2. (U//~~FOUO~~) Neuroimaging and Neurosensing Techniques..14

(U) System Virtualization Technology Forecast

(U) Neuroscience–New Game–Changing Enabler

(U//~~FOUO~~) While scientific exploration of the links between brain and behavior has been underway for the better part of a century, the general field of neuroscience (including cognitive neuroscience) has experienced an explosion in terms of the number of researchers engaging in the area and the amount of money being dedicated to the area has been underway for the last 30 or so years. [redacted] Figure 1 shows the number of articles from 1960–2008, measured approximately 10 years apart, that had at least one of the following terms as keywords:

(b) (7) (E)

- fMRI (functional magnetic resonance imaging)
- EEG (electroencephalography)
- MEG (magnetoencephalography)
- fNIRS (functional near-infrared spectroscopy)
- PET (positron emission tomography)
- TDCS (transcranial direct current stimulation)
- TMS (transcranial magnetic stimulation)
- Brain
- Neruoimaging
- Neuroscience

UNCLASSIFIED//~~FOR OFFICIAL USE ONLY~~

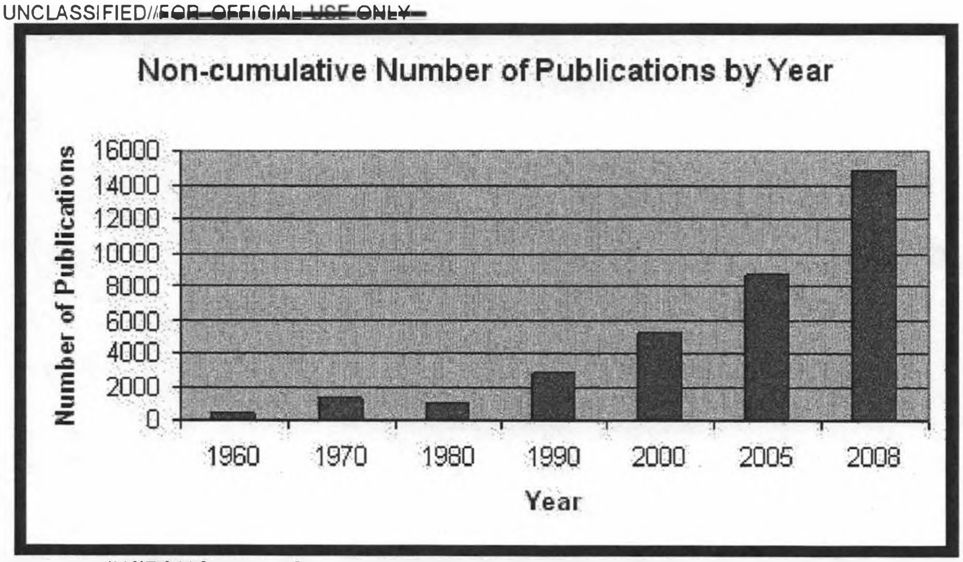

Figure 1. (U//~~FOUO~~) Non-Cumulative Number of Peer-Reviewed Publications by Year that Involve Neuroscience or Cognitive Neuroscience Work[a]

(U//~~FOUO~~) Figure 2 shows the number of references from 1932–1970 and from 1971–2010 (including papers in press) that contain one of the above keywords.

[a] (U//~~FOUO~~) Taken from the American Psychological Association's PsycInfo database using the above-listed keywords.

UNCLASSIFIED//~~FOR OFFICIAL USE ONLY~~

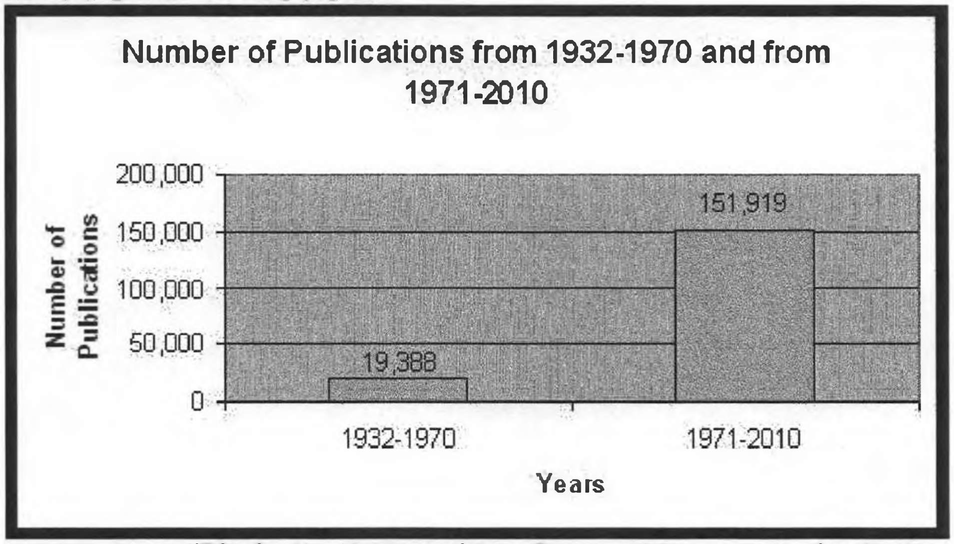

Figure 2. (U//~~FOUO~~) The Number of Peer-Reviewed Publications for the First 38 Years and for the Last 39 Years of Sustained Scientific Study in the Neuroscience and Cognitive Neuroscience Fields[b]

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) [redacted]

(b) (7) (E)

[b] (U//~~FOUO~~) Taken from the American Psychological Association's PsycInfo database using the above-listed keywords.

[redacted]

(b)(7)(E)

[c] (U//~~FOUO~~) [redacted]

[redacted]

(b) (7) (E)

[d] (U//~~FOUO~~) [redacted]

[redacted]

(b) (7) (E)

e (U//~~FOUO~~) [redacted]

[redacted]

(b)(7)(E)

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) [redacted]

[f] (U//~~FOUO~~) [redacted]

(b) (7) (E)

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) Thus, the current forecast attempted to survey this broad field in order to determine both where there might be opportunity for technical surprise and where the science does not support current capability claims. [redacted]

(b) (7) (E)

(U//~~FOUO~~) Therefore, this report is organized as follows: [redacted] After the main body of the report, *appendix A* contains summaries of enabling technologies for this forecast, each of which is treated as a separate chapter. *Appendix B* contains a more detailed review of the existing state of the art in the four topic areas that form the basis for this work—Deception Detection, Cognitive State Detection, Cognitive State Management, and Direct Control of Devices.

(U) Forecasting the Human-Network Synergy: What it Might Mean for People to be "More than Metaphorically" a Part of the Network

(U//~~FOUO~~) As a part of the forecast, we presented the results of the broad-brush overview to the customer and asked for specific direction for a final deep-dive investigation that would be of particular interest. The area that held particular interest was the issue: "how might humans become, more than metaphorically, a part of the network?" We start this section with a high-level view of the technologies that might make this possible, and then we discuss possible scenarios for humans to become a part of the network while supplementing and/or replacing the existing infrastructure, for example, using the human brain as a supplementary processor.

(U//~~FOUO~~) [redacted]

(b) (7) (E)

[g] (U//~~FOUO~~) EEG is electroencephalography.

Table 2. (U//~~FOUO~~) Neuroimaging and Neurosensing Techniques

Technique	Underlying Technology	Current Cost ($US)	Current Bottleneck	Needed Advancements	Current Viability
fMRI	Blood molecules' spin relaxation detection		Large homogeneous magnets, Cryogenically-cooled; Poor temporal resolution	Portable, large magnetic fields with acceptable field homogeneity functioning at/near room-temperatures	No
PET	Radioactive tracer detection		Radiation, poor temporal resolution. Paired with CT scan for anatomical information	Technique is self-limited due to tracers and their half-lives	No
MEG	Emanating magnetic field detection		Cryogenically-cooled detectors, Magnetically shielded rooms	Portable room-temperature detectors, active noise cancellation	Yes, lab environment [redacted]
EEG	Emanating electric field detection	50K	Reliable contacts, live signal processing	Improve ease-of-use, fast processing algorithms, improve contact impedance	Yes, COTS (e.g. Emotiv, Neurosky)
fNIRS	Optical signal (scattered and altered by the cortex) detection		Weak optical signal requiring close proximity for detectors	Sensitive room-temperature IR detectors,	Yes, lab environment [redacted]
Eye Movements	Eye movement and fixation tracking (optical methods)		Non-obstructed line-of-sight to one's eyes (no glasses).	Reliable electromagnetic methods not prone to artifacts (head movement, obstructions, etc)	Yes, lab environment (e.g. academia)

(b) (7) (E)

UNCLASSIFIED//~~FOR OFFICIAL USE ONLY~~

(U//~~FOUO~~) Figure 7 displays these techniques in three dimensions—temporal resolution, spatial resolution (probe size), and financial. If the techniques depicted in the graph below are treated equally (yield the same information) the desired tool would have the smallest probe size, would be the cheapest and would be very inexpensive. Unfortunately, to date no single technique can satisfy all three requirements. Moreover, the information obtained from the above techniques is often complimentary and not substitutive; there is no one ideal neuroimaging technique.

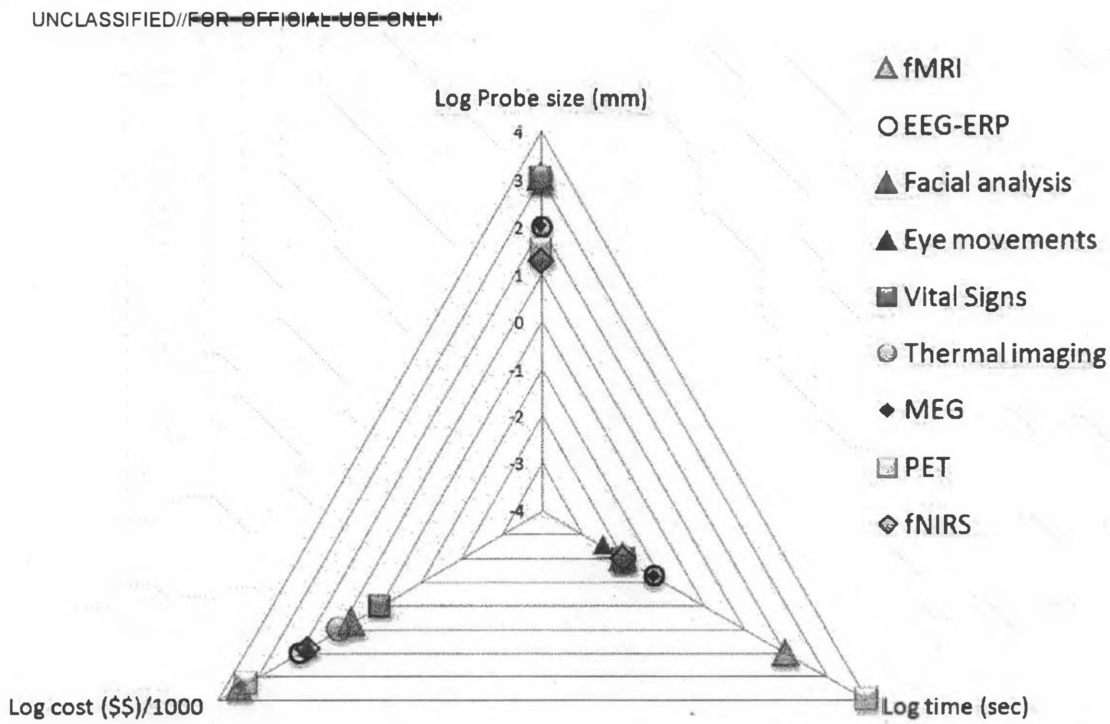

Figure 7. (U//~~FOUO~~) Plot of Cost, Time Resolution, and Sensor Size

(U) The Current Human-Network Interaction

(U//~~FOUO~~) In the discussion of the human-network interaction, we will take into account EEG as well as some of the other capabilities (for example, fNIRS and MEG) in determining how humans might become more than metaphorically a part of the network.

(U//~~FOUO~~) Figure 8 illustrates the current nature of the human interaction with the network. Information available from the network to the computer is tremendous in comparison to the information provided to the human, and back out again to the network, as represented by arrows of different size. Because the human only sees and attends to a small amount of information that is actually available on the network, and because he filters the information he provides back out to the network, the amount of information lost in the current human-network interaction is huge.

(U//~~FOUO~~) The second loss of information happens in between the input the human receives and the output he sends back onto the network. Specifically, his brain responds to the input through retrieval of huge amounts of information from long-term memory, including emotional responses, tangentially related memories, and unique or insightful links he makes between the input and other experiences. However, the amount of information that he then transmits back out onto the network is very sparse compared to the actual psychological/physiological response elicited by the input. Part of this filtering happens explicitly by way of what he decides to share (that is, what to type), part of this filtering happens as a function of the input devices he has access to—keyboard and mouse—which require him to translate the information he does want to share into mouse clicks, movements, and linguistic/symbolic input from the keyboard. While quite a bit of information can still be gleaned from this input, it pales in comparison to the information that could potentially be gleaned if his brain were tied directly to the CPU or to the network.

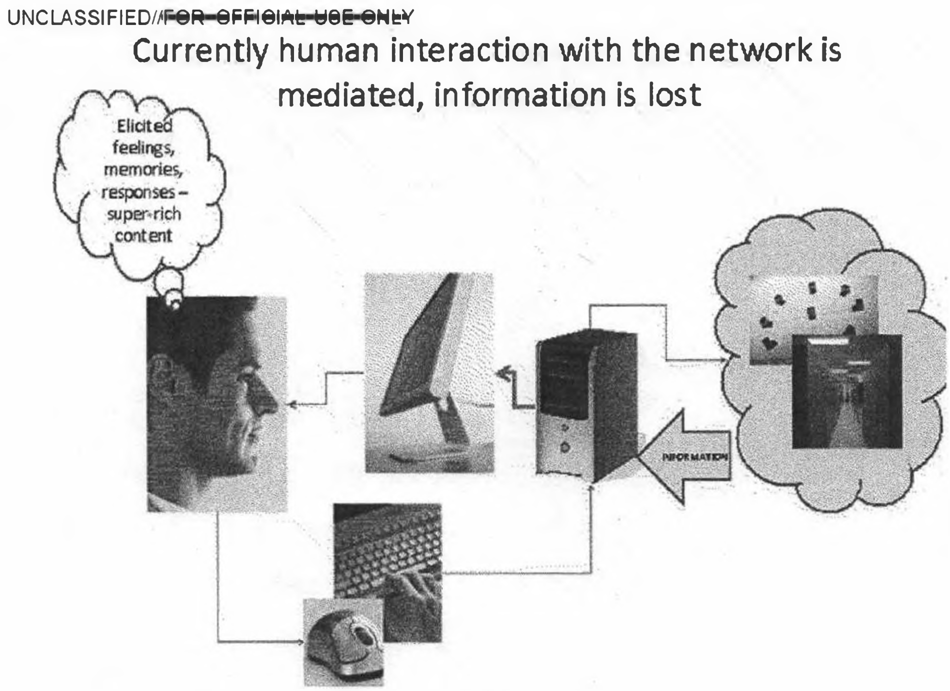

Figure 8. (U//~~FOUO~~) Information Available from the Network to the Computer

(U) Direct Input from the Human Brain to the Network

(U//~~FOUO~~) Figure 9 illustrates what the human-network interaction might be like if neural activity were transmitted more directly to the network and not mediated by keyboard and mouse. In this diagram, the user is still providing input to the network via keyboard and mouse, but this information simply augments the tremendous amount of information available via a direct, brain-machine interface connection to the network. Thus, the complexity of the human brain directly contributes information to the complexity of the network.

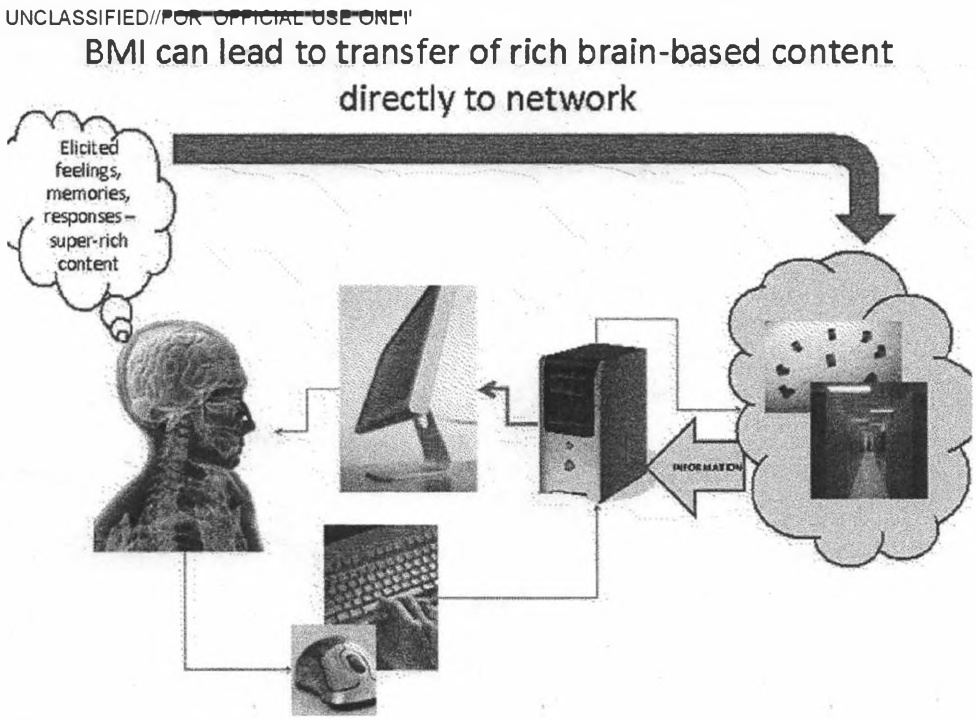

Figure 9. (U//~~FOUO~~) The Amount of Information Transfer to the Network about that Individual Would Increase If the Human Brain Were Tied to the Network

(U//~~FOUO~~) A number of applications could emerge from this kind of human-network interaction. [redacted]

(U//~~FOUO~~) One class of "more than metaphorical" applications involves direct access to "unprocessed"-unsmoothed EEG data. We envision these situations enabling someone on the network to tap into more than just what a user says—but to be able to tap into how they are responding to whatever is being shown on their local monitor. Several of these ideas leverage a common paradigm in cognitive psychology and cognitive neuroscience called rapid serial visual presentation (RSVP) in which images are shown to subjects for very short periods of time (for example, 100–200 milliseconds). When using RSVP, subjects are not always aware of having seen the image, but their brain (particularly visual cortex) still responds. That is, while they don't explicitly process the stimulus, their brain still processes it subconsciously. Importantly, a number of different researchers have demonstrated that the signals generated by visual cortex differ based on subjects' experience with the stimulus.[1,2] These ideas include the following, which we will discuss in more detail:

- Biometric key encryptions—using "brain waives" as a method for encryption.
- Identification—unique signal as a "fingerprint".

[redacted]

(b) (7) (E)

- Outsourcing the recognition of images—People as SETI

(U) Biometric Key Encryptions

(U//~~FOUO~~) In this possible application we explore a possibility of encrypting the information using innate, biometric signals. A combination of sign als (for example, EEG and fNIRS) emanating and detected as a result of one's thought process is unique in it's nature since it is attributed to an individual. This implies a need for an *a priori* "calibration". A calibration routine may be performed and shipped to the destination in a covert manner later serving as a decryption key to an incoming process. For example, a combination of signals describing each letter of alphabet is stored on a local (source) and remote (handler/HQ) computers. Moreover, due to brain plasticity (changes over time) this combination key is updated on a regular basis. The typed cable is then converted into a combination of waveforms by the computer and transmitted over the network to the destination (handler/HQ). Using the biometric key the transmission is now decrypted.

(b) (7) (E)
OGA

(U) [redacted]

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) [redacted]

(U) [redacted]

(U//~~FOUO~~) One interesting area of research in the online world is in the area of identity—how do individuals identify themselves in otherwise anonymous environments? This research originally surfaced in online classroom environments and found that while many people portray themselves fairly accurately, some chose wildly different personas—taking on dialects, appearances, and cultural markers that did not match their self-presentation in the real world. This relatively anonymous environment has given rise to the problem of online predators and online "children" who are really law enforcement agents acting like vulnerable teens and pre-teens.

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) [redacted]

(b)(7)(E)
OGA

(U) [redacted]

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) [redacted]

(U) [redacted]

(U//~~FOUO~~) [redacted]

[redacted]

(U//~~FOUO~~) [redacted]

[redacted]

[redacted]

(b)(7)(E)
OGA

(b) (7) (E)
OGA

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) [redacted]

(U) Outsourcing the Recognition Events: People as SETI Processors

(U//~~FOUO~~) One possible application that involves using the brain's processing power is to essentially turn individuals into implicit processors. An analogy is the use of individual personal computers to process information collected by the Search for Extraterrestrial Intelligence (SETI) program. In that program, individuals signed up to allow SETI to use their home computers to process data when their computers were not being used. The idea here is that the human brain—especially primary sensory cortices—process much more information than ever makes it into conscious awareness. Thus, a human brain that is tied more directly to the network (for example, via a COTS device) could be used to process sensory stimuli (such as imagery and maybe auditory stimuli) at an implicit level without significant disruption to the explicit processes in which they are engaging.

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) [redacted]

(U) Two-Way Interaction Between the Human Brain and the Network

(U//~~FOUO~~) In addition to a one-way feed of complex brain-based signals to the network, we can envision a scenario in which the network can provide direct input to the human brain. Considering the COTS EEG devices, one clear question to ask is: can these electrodes be used for transcranial direct current stimulation (TDCS)? That is, could they become two-way electrodes—both sensors and stimulators?

(U/~~FOUO~~) [redacted]

UNCLASSIFIED//~~FOR OFFICIAL USE ONLY~~

This could lead to a two-way implicit feedback loop

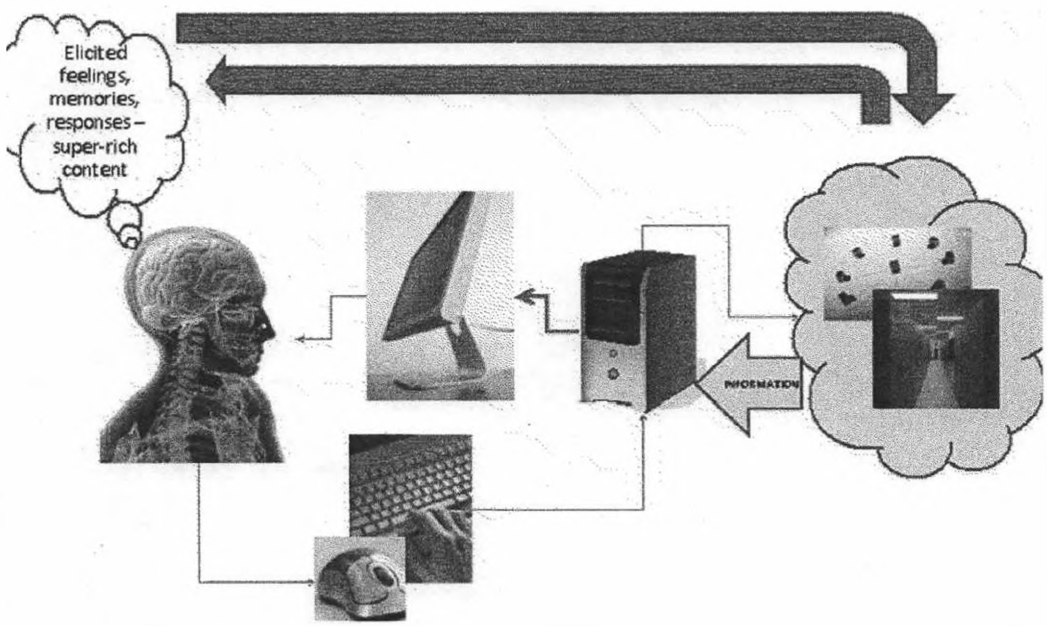

(b) (7) (E)
OGA

Figure 11. (U//~~FOUO~~) A Two-Way Direct Connection Between the Brain and the Network

(U//~~FOUO~~) As with the one-way flow of information discussed above, there are a number of applications that might emerge from a two-way brain-network relationship. [redacted]

(U) Altering Environmental Consequences Via Two-Way Electrodes

(U//~~FOUO~~) Much of human behavior is shaped by the relationship between behavior and the consequences of that behavior. Originally codified by B.F. Skinner, operant mechanisms of behavior are a critical factor in understanding and modifying the way people act and interact, either in person or in a virtual world. Behaviors that yield pleasurable outcomes (for example, getting money, recognition, approval) increase the likelihood of that behavior occurring again whereas behaviors which yield aversive outcomes (for example, pain, hunger, rejection) reduce the likelihood of that behavior occurring again. In part, the neural circuits for this environmental "teaching signal" are known.[8]

(b)(7)(E)

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) [redacted]

(U) [redacted]

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) [redacted]

(b) (7) (E)

(U) Appendix A: Enabling Technologies

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) [redacted]

(U) Neuroimaging

(U) Goals

(U//~~FOUO~~) Neuroimaging is the use of one or more sensors that are designed to detect neural activity while the subject is performing some kind of task, most often cognitive, although neuroimaging is also used for diagnosis of brain injury and disorders. There are a number of approaches to neuroimaging, and each has its advantages and disadvantages. Thus, there is no one best tool for measuring cognitive activity, the best tool depends on the problem one is trying to solve, and the constraints (for example, need for portability).

(U//~~FOUO~~) The framework of this section is to describe how a technology works and its relative advantages and disadvantages. Then, ways in which each technology might be improved are examined. Improvements to a component of the system will start with italics, while a new idea will simply be set off with a new paragraph. Each idea will have a subjective plausibility and confidence rating, where the confidence rating (from 1 to 10, with 10 being most confident) combines an estimate of how likely it is that the technology would work and be helpful with a time component; high scores are plausible ideas that we believe may be feasible soon.

(U) Functional Magnetic Resonance Imaging (fMRI)

(U//~~FOUO~~) Functional magnetic resonance imaging (fMRI) has a number of current advantages, including its relative availability (and thus, proliferation of its use in the literature). This imaging method provides relatively fine-grained spatial resolution (on the order of cubic millimeters), but has fairly poor temporal resolution (on the order of seconds). Thus, it can be used to localize selective neural activity throughout the volume of the brain. But, because cognitive events happen on the order of tens to hundreds of milliseconds, fMRI does not provide information on the temporal relationships between different areas of the cortex or subcortical structures that are selectively recruited for specific tasks. Thus, while a network of structures can be identified, how they interact with one another cannot be determined. Furthermore, if two tasks seem to recruit similar networks, one cannot determine if there are differences in how specific areas of those networks interact with one another as a function of task.

(U//~~FOUO~~) Finally, and this is true of all noninvasive neuroimaging methods, imaging at the level of cubic millimeters is still a relatively coarse measure of neural activity as there may be thousands of neurons, and thus hundreds of neural circuits, contained in a cubic millimeter of brain tissue. Thus, noninvasive neuroimaging methods are not able to image at the resolution of local neural circuits or cortical columns, even though these are hypothesized to be the functional units of the brain.

(U) *How It Works*

(U//~~FOUO~~) A subject is placed in a large, homogenous magnetic field, usually 1.5 or 3 Teslas in strength (although there are a few locales with stronger fields). The magnetic field has the effect of aligning the spins of susceptible particles, most importantly water (for structural imaging) and oxygenated/deoxygenated blood (they have different magnetic properties, and so can be detected for functional imaging). In order to create a strong, homogenous field, the magnet is a liquid-helium-cooled superconducting electromagnet. Other resistive electromagnets, sometimes in the same machinery as the primary magnet and sometimes in separate coils, create a small gradient in the magnetic field which allows for localization of signal. Finally, at specified times a radio frequency (RF) pulse is sent into the magnet (and the subject) to disrupt the particles from their orientation. As the spins relax back into alignment with the primary magnet they send off an RF signal, which is detected and analyzed to create the MRI/fMRI image. MRI has excellent spatial resolution but poor temporal resolution because it measures the hemodynamic response (blood flow) instead of direct neural events.

(U) Primary Magnet

(U//~~FOUO~~) Because the magnetic field must be strong and homogenous, a superconducting electromagnet is generally used. However, this leads to a large, non-portable system. There are a couple of possibilities for making the primary magnet portable:

(U//~~FOUO~~) High temperature superconductivity: If an electromagnet could become superconductive at a higher temperature, the liquid helium and associated machinery would be unnecessary. Currently the "warmest" temperature for a (consistently tested) superconductor is around 130° Kelvin or -140° Celsius. Obviously this is far away from room temperature superconductivity, which is optimal from an operational perspective. These high temperature superconducting materials have the advantage that they can be cooled by liquid nitrogen or something else cheaper and more available than liquid helium, but the disadvantage that they are typically ceramics that are expensive and not easily formed into wires. If a magnet could be created with a high temperature superconductor, it is unclear how much of a reduction in size or increase in portability would result. Judgment: plausible, but unlikely to be overly helpful from an operational perspective. Confidence: 4

(U//~~FOUO~~) Alternative magnets: Instead of using a superconducting electromagnet, a different magnetic field generator could be used. This would lead to a smaller system since cooling equipment would not be necessary. Prior work in MRI has used non-superconducting electromagnets or permanent magnets. Both of these methods suffer from a lack of field strength and homogeneity and are not currently used. However, if portability is a key concern, and the other elements of the machinery could be incorporated (gradient field generator, RF transmitter and receiver), then a trade-off between signal to noise ratio and size may be worthwhile. Judgment: plausible, but unlikely to be helpful. Confidence: 3

(U) Gradient Magnets

(U//~~FOUO~~) These magnets produce a small gradient, on the order of tens of milliteslas, across the primary field to enable spatial localization. These are typically normal (resistive) electromagnets. There is no particularly attractive way to improve this portion of fMRI.

(U) Radio Frequency Coils

(U//~~FOUO~~) A radio frequency pulse is sent from transmitters usually incorporated into the magnet machinery to "knock over" the magnetic spin of particles in the subject. The resulting radio frequency signal that comes back as the particles relax back into orientation with the magnetic field is picked up by detectors. More accurate imaging comes from receiver coils placed in the area of interest (for example, a "birdcage" coil placed over the head of the subject for fMRI) as opposed to receivers in the machinery itself. [redacted]

(b) (7) (E)
OGA

- (U//~~FOUO~~) [redacted]
- (U//~~FOUO~~) [redacted]

(U//~~FOUO~~) Thus, while fMRI has some significant advantages for laboratory work with cooperative subjects, the likelihood that it is an easily fieldable technology that could be ripe for technology surprise is, at this time, somewhat low. [redacted]

(U) Positron Emission Tomography (PET)

(U//~~FOUO~~) Positron emission tomography (PET) is another method for imaging the location of activity throughout the brain. However, as with fMRI, it currently images hemodynamic activity, and thus provides a "slow" signal—that is, low temporal fidelity. It also requires the use of radioactive tracers, and as such, should be considered mildly invasive. Finally, since the development of fMRI methods, the use of PET for neuroimaging in cognitive neuroscience has radically decreased.

(U) How It Works

(U//~~FOUO~~) A variety of substances can be used as markers of brain function; oxygenated blood, for example (as in fMRI), but also water and various neurotransmitters. A radioactive tracer can be created that mimics one of these substances and thus acts as a proxy for it; fluorodeoxyglucose is a popular one that follows the use of glucose in the brain. As the tracer releases positrons, they interact with electrons in nearby tissue and annihilate, releasing two gamma photons. These photons are detected by machinery placed around the head and the locations and counts are used to determine the amount of "activity" with good spatial accuracy. However, PET can only be used so many times on the same subject due to radiation concerns, and the use of radiotracers with relatively short half-lives requires a cyclotron nearby. Furthermore, no anatomical data is collected by a PET scan, so a CT scan is usually incorporated into the procedure to allow for registration of the data. The temporal resolution of PET is very poor due both to the hemodynamic-like nature of the measure as well as the decay properties of the radioactive tracer. Between the radioactivity, need for a separate anatomical scanner, and the need for a cyclotron, it is not worth trying to upgrade PET; current fMRI technology is sufficient to replace PET in most aspects.

(U) Magnetoencephalography (MEG)

(U//~~FOUO~~) In contrast to fMRI and PET, magnetoencephalography (MEG) actually senses the magnetic field emitted when neurons fire, and thus provide relatively fine temporal resolution (on the order of tens of milliseconds). However, because of the fact that the emitter of this signal is being sensed at the scalp, the location of that emitter must be inferred rather than being directly measured, as in fMRI. Thus, the spatial localization available for MEG is not as reliable as it is for fMRI.

(U) How It Works

(U//~~FOUO~~) When a neuron fires, a current essentially flows through the cell body down the axon to a neighboring cell. When enough neurons in the same area and same general direction fire in synchrony, the current becomes sufficient to produce a magnetic field detectable at the scalp. This signal is still many orders of magnitude smaller than the ambient magnetic noise around a person, however, so recordings are done in a magnetically shielded room with very sensitive superconducting quantum interference devices (SQUIDs) operating at cryogenic temperatures. Thus, the requirement for magnetic shielding and the large cryogen container make current MEG not feasible for fielded applications. This as also resulted in a smaller amount of the published literature using MEG relative to fMRI.

(U//~~FOUO~~) Due to the drop-off of magnetic strength with distance, it is more difficult to image deep structures in the brain, although it can be done. Current data analysis techniques allow for excellent temporal and pretty good spatial resolution; better spatial resolution compared to EEG (electroencephaography; described below) comes from the fact that the skin and scalp smear electrical signals but not magnetic signals (at least, not as much). Like PET, MEG contains no anatomical data, so a separate MRI is usually conducted with markers that allow the data to be coregistered. One method for making MEG more operationally relevant is to improve the sensors such that they don't require supercooling.

(U) Superconducting Quantum Interference Device (SQUID)

(U//~~FOUO~~) Much of the size of the MEG apparatus could be eliminated with portable, room temperature SQUIDs or other magnetometers. It sounds as if this may be possible soon. One issue with high-temperature superconducting SQUIDs is that they are typically less sensitive than typical superconducting SQUIDs; testing would be necessary to ensure that the signal to noise ratio is acceptable for brain imaging. However, there are alkali magnetometers that have been developed by researchers at Princeton, and that are currently being developed with neuroimaging in mind at Sandia National Laboratories. While these alkali magnetometers still require a magnetically shielded room, they are currently 1 cubic inch in size, and are thus significantly more portable than production-quality MEG systems. Judgment: Very plausible. Confidence: 8

(b) (7) (E)
OGA

(U) Magnetic Contrast Agents

(U//~~FOUO~~) Similar to the idea described under fMRI, a nanoparticle could be created that produces a magnetic field that would effectively boost the signal detected by the MEG. A nanoparticle that could align itself with a neuron would produce an accompanying magnetic field in the same direction, essentially becoming a second neuron and increasing the signal. [redacted]

(U//~~FOUO~~) Thus, the possibility that MEG systems could be fielded in operational environments in the relatively near future is much more realistic than fMRI. However, significant work will need to be done on the signal processing side in order to make MEG signals usable in real or near to real time. This need is true for other imaging methods as well, and represents one of the key technical hurdles identified in the main body of this report.

(U) Electroencaphalography (EEG)

(U//~~FOUO~~) Electroencaphalography (EEG) is a complimentary measure to MEG, and EEG data can actually be collected at the same time as a MEG (unlike with fMRI because of the disruption the magnetic field caused by EEG electrodes). Similar to MEG, EEG provides very fine temporal resolution, but it has the disadvantage of the inverse problem, like MEG, as well as the smearing of the electrical signal at the scalp. However, there are a significant number of researchers who are working to perfect algorithms to allow accurate inference of dipole (emitter) localization in EEG. Unlike MEG, however, EEG is significantly more portable, and as will be discussed in more detail below, commercial off-the-shelf EEG headsets are currently being used in a variety of applications.

(U) How It Works

(U//~~FOUO~~) As noted for MEG, neural firing is essentially a small current (at least within the cell). When sufficient neurons fire in synchrony in the same direction, the current becomes large enough to detect at the skull. Also similar to MEG, the electrical potential drops off with distance, so it is harder to detect activity in deeper brain regions. Because the electrical signal is smeared by the skull and scalp, spatial resolution is poor. However, temporal resolution is very good. The signal is recorded at the scalp by an array of electrodes, typically ranging from as few as 16 to as many as 256, and then amplified and compared to a reference. The raw EEG signal can be analyzed for the power of different frequency bands, which correspond to different cognitive states, or it can be "cut up" and averaged together by locking the EEG to the onset of a number of events, creating an event-related potential (ERP). ERPs are analyzed for the amplitude and latency of various components which correspond to different cognitive processes. In either case, the signal undergoes filtering and artifact correction to take out frequencies due to heart beat, breathing, power lines, and eye movements/blinks, to name some common sources.

(U//~~FOUO~~) Despite its relative portability and low cost, there are still a number of ways that EEG could be improved to make it more attractive for operational environments, and therefore more likely as a technology surprise area.

(U) Electrodes

(U//~~FOUO~~) To lower resistance and enhance the electrical signal, a subject is prepped for recording by having some skin abrasion at electrode sites to remove dead skin cells and a conductive gel is applied which ends up throughout the hair. The electrodes are also attached to wires which send the signal to amplifiers and then on for analysis. It would be preferable to have electrodes that don't require preparation and can transmit the signal wirelessly. A few companies have created such dry EEG systems, including Emotiv and NeuroSky (although the Emotiv system actually requires moistening felt pads with saline solution, so they aren't really dry). The downside to a dry, wireless EEG system is that the amount of noise must be greater than that in the research setting because of the lack of preparation. It is unclear how the companies have dealt with this issue, but to the extent that the product works they must have accounted for it. Another alternative would be to create electrodes that can release their own gel locally, perhaps by keeping limited contact with the scalp via suction cups, which would enhance conductance while limiting the undesired goopy gel everywhere. Judgment: doable, if not already done. Confidence: 9

(U) Remote Sensing

(U//~~FOUO~~) The main current limitation to EEG besides spatial resolution is the ability to read the brain's electrical signal from further away than the scalp. Given the miniscule size of the signal, and thus a very poor signal to noise ratio, this is highly implausible. Creating a nanoparticle to boost the electrical signal in the brain is likely a mistake as injecting electrical current throughout the brain can have unknown and undesired effects like inhibited cognition and seizures. As a completely ridiculous idea, however, if a method could be developed for creating quantum entanglement between particles in the brain and particles outside the brain, measurements could be made of those particles outside the brain, allowing for remote sensing. Judgment: totally implausible. Confidence: 0.

(U//~~FOUO~~) Thus, even though EEG does have a number of limitations, we believe that due to its portability and to the amount of cognitive neuroscience research that has used it as a tool, EEG actually is an area that is ripe for technology surprise. This assertion is discussed in the final section of this report.

(U) Functional Near-Infrared Spectroscopy (fNIRS)/Optical Imaging

(U//~~FOUO~~) Systems such as Functional near-infrared spectroscopy (fNIRS) have the capability to detect two types of signals—a slow, hemodynamic signal and a fast optical signal. The fNRIS optical signal, called the *event-related optical signal*, or EROS has the notable advantage of both high temporal and spatial resolution—something no other neuroimaging method offers. fNIRS is also highly portable and robust against artifacts that plague EEG. The disadvantage of such systems currently is that first, they can only measure very shallow activity – penetration of near-infrared light into the head is only around 2 to 3 centimeters, and second, current systems are actually blood oximeters and have thus not been optimized for cortical imaging. EROS data also have a very poor signal to noise ratio. Finally, such systems are still relatively expensive ($500K) relative to EEG systems, which run around $50K or so.

(U) How It Works

(U//~~FOUO~~) There are a number of changes in optical properties related to brain function. For one, oxygenated and deoxygenated blood have different absorption properties. By passing near infrared light through the scalp and skull into the brain then detecting it as it scatters back to the scalp, this difference in absorption can be measured and used to detect the amount in a particular region of brain. This is essentially analogous to how fMRI and PET work and is called the slow optical signal. Absorption can also be used to determine the amount of other substances in the brain.

(U//~~FOUO~~) It is also true that when neurons fire, they change shape (swelling in size) due to the influx of ions. This change in size can be measured via changes in the scattering of the light that passes through the neurons. This is called the fast optical signal (EROS). The slow optical signal has the temporal resolution issues of fMRI whereas the fast optical signal has good temporal and spatial resolution. In either case, however, imaging is limited to the top 2 or 3 centimeters of cortex in the best case because the light can only be detected if it scatters back to the surface; if it scatters more deeply or is absorbed, the information is lost. Thus the major limitation in fNIRS is the portion of the brain that can be imaged. While certain types of fNIRS (different measurements of the optical signal are possible) can be conducted with LEDs and simple detectors to allow small, portable systems, these are typically the less informative measures. The fast signal, at least currently, requires lasers and a wired system.

(U//~~FOUO~~) Given that the major limitation of fNIRS is related to how deep into the brain it can image, the primary desire would be to enhance the ability to detect the signal from deeper in the brain.

(U//~~FOUO~~) One method of accomplishing this would be by transmitting the light back to the surface in some manner, perhaps via nanoparticles. However, this is likely impossible because the light would continue to be absorbed and scattered on the way out, and so what is desired is a method of taking a low energy photon and converting it into a higher energy photon, which would violate conservation of energy. This could potentially be circumvented if the nanoparticle were additionally powered by something else, perhaps stray electrical energy from action potentials, but this seems unlikely. Judgment: not likely. Confidence: 0

(U//~~FOUO~~) Another option would be to have either photovoltaic or photomagnetic nanoparticles throughout the brain that would respond differentially to different characteristics of the light incident on them. Instead of having light detectors on the scalp, EEG or MEG would be recorded around the lasers/LEDs. When light sent through the brain scatters onto a nanoparticle, it would respond and emit an electrical or magnetic signal. This would enable information from light that scatters deeper into the brain to be obtained outside the scalp. However, this sounds implausible and localization would be an issue. Judgment: not likely. Confidence: 0.

(U//~~FOUO~~) A third option would be to optimize fNIRS for cortical imaging as a function of light wavelength and modulation. By studying the propagation characteristics of different frequency signals, an optimal band could be chosen to probe a specific area of the cortex.

(U//~~FOUO~~) Even though fNIRS is highly portable, EROS provides desirable characteristics with regard to spatial and temporal resolution, and it is robust against the artifacts that EEG has, due to its limited cortical extent and due to the relatively sparse literature using EROS as a dependent measure, we believe that fNRIS is not as ripe for technology surprise as is EEG. However, much information about cortical activity can be sensed using fNRIS, and if there is movement towards creating cheap, production level systems, then fNIRS could be just as feasible an area for tech surprise as EEG.

(U) Eye Movements

(U) How It Works

(U//~~FOUO~~) Since visual information must come in through the eye, measures of where the eye is looking can be used to derive information about what a person finds interesting, grabs his attention, or knows about the material. There are two primary ways in which eye movements are measured: with optical systems that bounce light off the retina and use the reflection to determine point of gaze, and electrical systems that use changes in the position of the charge across the eyeball to determine relative movements. Both systems involve a calibration period to develop a starting point (otherwise all measures are relative) and some method of accounting for head motion, whether they are optical or mechanical (a chin rest to limit head movement or possibly a gyroscope to account for it).

(U) Optical Imaging

(U//~~FOUO~~) Optical imaging has been used for decades and improvements are likely to be small changes. Most research systems involve a table top camera (or two) which emit and detect infrared light and account for head movement. These systems have the benefit of being non-intrusive, but are susceptible to large changes in head position or wide angles of vision (looking too far up or down, left or right). This can be accounted for by using a head-mounted system, which places a small camera just above or below the eye. A head mounted system accounts for head motion by moving with the head, but is more intrusive. Optical eye tracking is also susceptible to things that obstruct the light's path to the retina, such as blinking, some glasses and contact lenses, drooping eyelids, and long eyelashes.

(U//~~FOUO~~) More cameras: To better account for head motion, wide gazes, and other disruptions, a system with additional cameras could be used to provide additional information. Instead of one camera directly in front of the subject, two or three cameras could be placed at angles to the subject, perhaps also vertically in space. If the subject moved his head to one side one camera would primarily record useful information, but in general the cameras would be linked to integrate their measurements and obtain cleaner data. This may not help with blocking/reflecting issues like glasses or eyelids, however. Judgment: extremely doable. Confidence: 9

(U) Electrical Imaging

(U//~~FOUO~~) Since the eyeball has a differential charge from the lens to the retina, electrical recordings (via electrooculogram (EOG), which requires an electrode be placed near the eye) can be used to detect when the eye moves, and thus create relative points of gaze. While EOG recordings are commonly used to detect the fact that the eye has moved for issues of artifact correction in EEG, they are not typically used for actual eye tracking. The few sources that have investigated eye tracking via EOG raise issues of accuracy/resolution (typically not as good as optical tracking), limited range of gaze where the EOG and eye movements are linear, and the usual sources of noise in an electrical recording which requires a good amount of filtering/processing. Additionally, EOG would require placing electrodes on the subject, which would eliminate any potential discreet uses. Finally, eye tracking via EOG would need additional methods to determine head position as eye movements via EOG are relative, not tied to a particular gaze direction.

(U//~~FOUO~~) Accounting for motion: In optical systems, head movement is detected by a camera. Similarly, a camera or other optical system could be used to monitor head movements in electrical imaging. However, adding a camera largely defeats the purpose of using an electrical system; otherwise, tracking could be done with the camera. Instead, a gyroscope placed on the head could be used to account for head movements and provide a frame of reference for absolute eye position as opposed to relative eye movements. The wireless dry EEG system created by Emotiv contains gyroscopes; it could be in part for this purpose. Judgment: extremely doable. Confidence: 10

(U//~~FOUO~~) While eye tracking via various methods has proven to provide useful information about cognitive state, it does require an instrumented system and/or a cooperative subject. If these devices proliferate, or of other devices that can provide this information (for example, COTS EEG devices) proliferate, eye tracking could potentially be used in an unwarned manner. More research on the kinds of information that could be garnered (for example, knowledge of particular types of information) would need to be done to determine the actual threat posed if, indeed, the hardware required to do eye tracking does proliferate.

(U) Secondary Measurements

(U) How It Works

(U//~~FOUO~~) A number of methods exist for obtaining secondary information about cognitive states (as opposed to the methods discussed so far, excluding eye movements, which are assumed to measure cognitive states more or less directly). These methods include facial analysis (looking for microexpressions of emotion), thermal imaging (using changing patterns of heat in the face to determine stress or anxiety levels), and vital signs measurements (heart rate, breathing rate, and galvanic skin response via a polygraph machine or Radar Vital Signs Monitor, RSVM). All of these methods suffer from the fact that since they measure a secondary response, they must be interpreted in context and are more variable from person to person. If it were possible to know with 100 percent accuracy that someone was angry after being asked a question, for example, the reason for and meaning of that anger would still need to be explained. The experience of anger, as well as stress or anxiety, is highly individual and likely to differ across situations. Additionally, it is possible that 100 percent accurate knowledge is impossible: human raters usually only reach 90 percent accuracy when categorizing faces by their emotions, even though the ability to do so, and the expression of the six major emotions (joy, surprise, fear, disgust, anger, sadness) appear to be universal and most likely genetically determined. [redacted]

(b) (7) (E)
OGA

(U//~~FOUO~~) [redacted]

(U) Generally Useful Techniques

(U//~~FOUO~~) One major limitation to all of these techniques is how well their measures correspond to different brain states. This is the goal of most neuroscience research, as signals are collected during tasks thought to tap some process. However, this is not necessarily as true for secondary methods. Additionally, direct methods typically make assumptions, such as looking only at specific regions of the brain as opposed to coordinated network activity (univariate versus multivariate in fMRI) or interpreting the signal within a specific set of circumstances (ERP). Thus further research better characterizing the correspondence between cognitive states and imaging measures is necessary.

(U//~~FOUO~~) Another common research theme is in categorization. In many cases, the goal is to know if a person is in one of two states: has he seen this before or not, is he lying or not, is he paying attention or not, etc. A number of researchers have attempted to categorize imaging data, but few seem to get any better than 90 percent with any kind of categorization. This is due in part to the noise in the data, which may be reduced by some of the suggestions listed above, but also likely in part due to the fact that psychologists do not often talk with good signal processing or mathematical modeling people. Many researchers have attempted categorization with regression or other linear techniques which cannot possibly be optimal because the data are likely not linear. A better approach would be to use some form of nonlinear multivariate analysis, perhaps computer neural networks. An analysis technique like neural networks will require training to learn the proper weights and connections, but once this is done it is possible that categorization and analysis accuracy will improve. Some other mathematical technique analogous to regression may provide categorization with less training, or allow for categorization of new subjects based on prior data. In any event, better analysis algorithms will be necessary not only to ensure that some pattern of activity means what it is thought to mean but also to allow for real time analysis of data so that action can be taken immediately instead of hours or days later.

(U) Cortical Stimulation

(U) Goals

(U//~~FOUO~~) Once there is an ability to determine a person's cognitive state through neuroimaging or some secondary measure, there may be a desire to actively and directly alter that state. Thus, the goal of this section is to outline the current abilities of technologies that apply external stimulation to the brain, namely TMS and TDCS, and how they might be used to enhance positive attributes (memory, attention) as well as decrease negative attributes (fear, pain). Most of the enhancement topics have been discussed previously, so they will only be summarized here. Thus this review will focus on the technical aspects of TMS and TDCS and how likely they are to have uses influencing feelings of fear, pain, and fatigue.

(U) Transcranial Magnetic Stimulation (TMS)

(U//~~FOUO~~) TMS works by using a magnetic field to induce eddy currents in the brain, effectively "injecting" neural activity into a certain brain region. The TMS apparatus consists of a power source made up of a number of capacitors which store charge until stimulation is desired. At that point the charge is released very quickly, leading to a large current change through a coil of wire. This current change creates a high-intensity (typically 1.5 to 2 Tesla) magnetic field, which travels into the brain. The magnetic field finally induces an eddy current in underlying brain tissue, typically parallel to the skull, which is thought to create synaptic firing. The specific strength, localization, and depth of action of TMS depend on the characteristics of the coil (and, obviously, the amount of current put through the wire). Figure A-1 illustrates a TMS system.

UNCLASSIFIED//~~FOR OFFICIAL USE ONLY~~

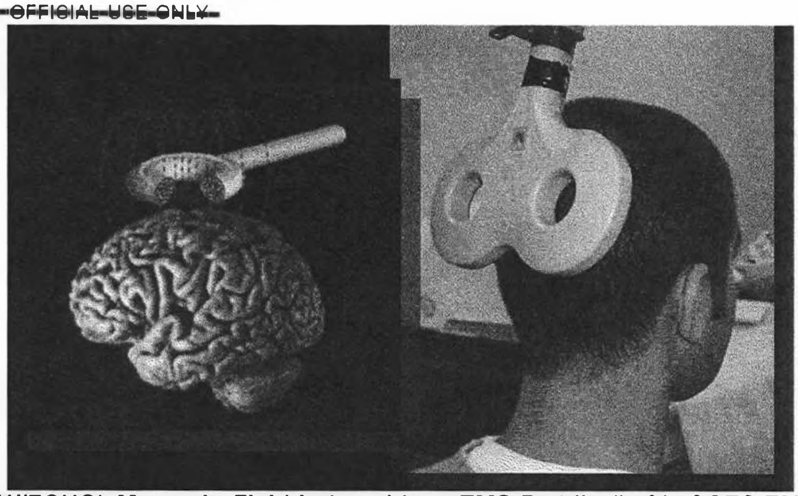

Figure A-1. (U//~~FOUO~~) Magnetic Field Induced by a TMS Paddle (Left); COTS TMS Device (Right)

(U//~~FOUO~~) The majority of neuroscience research coils are called figure-8 coils, because the TMS "wand" has two coils. Current is sent through each of them, leading to a stronger magnetic field directly below the region between the two coils and limited activations outside of this area. This configuration leads to magnetic field strengths of around 0.5 Tesla in the brain at a depth of around 2 centimeters below the scalp.

(U//~~FOUO~~) As such, TMS is limited to directly affecting shallow cortical brain regions, although there is evidence that other brain regions functionally connected to those that are targeted also change their activity (that is, there are downstream effects). Additionally, other coil configurations may enable deeper activation; [redacted][9] developed a design that reportedly has a much lower field drop-off with distance, allowing for stimulation of regions as deep as the nucleus accumbens without over-stimulating the closer cortical regions. While the field strength must be less at depth than it is at the cortical surface, their coil maintained a higher strength (~ 35 percent cortical strength) than the figure-8 coil (~8 percent) as far as 6 centimeters into a phantom brain. TMS does not inherently provide any anatomical information, but when used along with stereotaxic information, such as from an MRI, TMS can be used with accuracy of a few millimeters, or about on the level of fMRI spatial resolution. Again, many details of the induced activity will depend on the coil configuration.

(b) (6)
OGA

(b) (7) (E)
OGA

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) [redacted]

(b) (6)
(b) (7) (E)
OGA

(U//~~FOUO~~) [redacted]

(U) Transcranial Direct Current Stimulation (TDCS)

(U//~~FOUO~~) In contrast to TMS, TDCS works by directly applying current to the scalp which then travels into the brain. Electrodes are placed on relatively large (5 × 5 centimeter or larger) soaked sponges which are applied to the desired stimulation site, along with a reference elsewhere on the head. Direct current is then sent into the electrodes. A fair proportion of the current passes directly through the scalp and immediately out without entering the brain, but the rest appears to enter the brain directly below the electrodes. The current is thought to affect the resting potential of cell membranes. Stimulation can be anodal or cathodal; anodal stimulation raises the resting potential, making it easier for a neuron to fire due to other stimulation, while cathodal stimulation lowers the resting potential and thus inhibits neural activity.

(U//~~FOUO~~) The current is typically limited to 1 or 2 milliamps, although the important parameter is actually current density and thus electrode size is also important. It is also critical that the reference electrode be placed somewhere on the head because placing it on the body would lead to current passing through the brainstem and the body, which could lead to death. Limiting the current density is important for two reasons: since the electrodes are placed directly on the scalp, a high current could cause burning of the skin (note that this concern does not apply to TMS because the coil is encased in plastic, although the coil can become hot or even overheat to the point of failure); also, as with TMS, high levels of activity could lead to seizures if the electrodes are placed carelessly.

(U//~~FOUO~~) Electroshock therapy, which is not different from TDCS in principle, used current levels of 800-1000 milliamps and initially placed electrodes on either side of the head, which commonly leads to seizures (this can be avoided by placing both electrodes on one side of the head). Thus while the electrode size could be manipulated to allow for more specific targeting of stimulation, the current must also be adjusted to prevent adverse effects. However, TDCS is very simple from a technology standpoint; all that is needed is a power source (battery), a resistor to control the current, and electrodes. While TDCS uses direct current, it is typically applied in pulses similar to TMS, and so can be described as having a frequency as well.

(U) Enhancing Effects

(U//~~FOUO~~) As noted in the introduction to this section, most of the enhancing effects of TMS and TDCS have been discussed in previous reviews references needed.[11,12,13] As a brief summary, the positive effects of TDCS seem to be more consistent than those of TMS. This could be due to the fact that the anodal/cathodal distinction is much clearer than the frequency (or other less clear parameters) distinction in TMS. However, rTMS applied offline (not during a task) seems to have consistently positive effects; there could be a common mechanism underlying the two techniques. Both TMS and TDCS have been shown to have enhancing effects on mood, memory, working memory, and some procedural tasks; both have also been shown to have some disrupting effects. TMS and TDCS have frequently been applied to prefrontal cortex, however, which seems to be important for those processes, and so more focused research should be able to find positive effects in those tasks as well.

(U) Disrupting Negative Processes

(U//~~FOUO~~) While TMS and TDCS have been used to enhance desirable processes, it is less obvious that they could be used to disrupt undesirable processes. In this review, the processes considered are fear, pain, and fatigue. There is some positive evidence in this direction due to the known beneficial effects of TMS on depression; if TMS treatment can reduce a trait like depression, then it seems possible that it might affect other negative cognitive states.

(U//~~FOUO~~) Fear is perhaps the simplest of these three processes because it is typically localized to one brain region, the amygdala. While prefrontal cortex has also been implicated as a control region, it is somewhat unlikely that the emotional control regions are different from control regions that are used to guide other behaviors in the face of distraction or competing options. fMRI studies that use fearful stimuli, such as images of snakes and spiders or faces expressing fear, almost always find increased

(b) (7) (E)
OGA

activity in the amygdala compared to images of neutral stimuli or faces. Also, studies of patients with amygdala damage find that they have difficulty detecting or categorizing fearful faces and perhaps even further that they may not experience fear. The amygdala is a bilateral structure located deep in the brain, in the medial temporal lobes. As such, it is outside the range of effect of TMS and probably TDCS (unless it can be indirectly suppressed via connections to cortical regions). [redacted]

(b) (6)
(b) (7) (E)
OGA

(U//~~FOUO~~) [redacted]

(b) (6)
OGA

(U//~~FOUO~~) For example, [redacted][14] used "real-time" fMRI as biofeedback to teach subjects to raise or lower their subjective feeling of pain. Subjects were stimulated with a laser to a point where they rated the pain a 7 out of 10. The anterior cingulate has been suggested to be important to both pain sensation and pain regulation in previous studies, so it was used as a source of feedback in this study. Subjects were shown their ACC activity in real-time (actually slightly delayed due to the hemodynamic response, but otherwise unprocessed) and asked to either increase or decrease the signal using one of a few techniques (attend to or away from the pain, focus on the stimulated body part or a different body part, and so forth). After only a few 30-second blocks, subjects were able to do so, and this appropriately increased or decreased their subjective ratings of the pain.

(U//~~FOUO~~) The same result was found in subjects with chronic pain who were not stimulated with the laser. Control conditions, which failed to find significant pain modulation, included subjects with chronic pain who used heart rate, galvanic skin response, and breathing rate for biofeedback; normal subjects who were given the attention techniques but not biofeedback; normal subjects who were given one specific technique with twice as much training; normal subjects who were given feedback from a different brain region; and normal subjects who were shown a different subject's ACC activity. Thus the positive effect of the biofeedback seems to be due to controlling ACC activity. If TMS or TDCS could be used to affect activity in this region (perhaps using Roth's coil), then it seems plausible that pain sensation could be decreased. It is also possible that pain sensation could be decreased by decreasing somatosensory activity in general, although this may have the side effect of reducing all bodily sensations (pressure, temperature, body position, and touch in general).

(U//~~FOUO~~) Even more complex is the feeling of fatigue. Most neuroscience studies of fatigue involve making voluntary muscle contractions, such as squeezing a force gauge or contracting the biceps as hard as possible. As such, the conclusions likely differ from those that would be drawn from studies of long distance runners or other endurance tasks. However, fatigue has been tied to the sensorimotor cortex, supplementary motor cortex, prefrontal cortex, and the cingulate gyrus.

(U//~~FOUO~~) Previous reviews looking at articles more related to endurance also suggest that fatigue in the brain could be due to neurochemical modulations, such as reduced glucose and increased ammonia as a result of muscle activity. Fatigue obviously also occurs in the muscles themselves as resources like glucose are used up. Furthermore, there is an interaction between the brain and the body such that the brain can be urging full output when the muscles cannot provide it and alternatively the brain can feel tired and drive rested muscles at less than optimal levels.

(b) (6)
OGA

(U//~~FOUO~~) A review by [redacted][15] investigated the effects of applying TMS to motor cortex at different points during muscle fatiguing tasks (again, typically short-term exercises as opposed to endurance tasks). Motor cortex appears to become more excitable during and after muscle contractions, meaning that TMS has a larger effect (as measured by the motor evoked potential-MEP) and can be used to maintain cortical drive to the muscles. However, it is important to note that the MEP differs both across subjects and across muscle groups. For example, additional muscle fibers continue to be recruited during a biceps contraction up to 75 percent of maximum MEP, whereas most fibers in the hand are already recruited by 30 percent MEP. Thus it seems likely that TMS would be differentially effective depending on the targeted muscle group. Additionally, since most of these tasks use short bursts of voluntary contractions, it is unclear if TMS could be used to reduce fatigue during, say, a long march. It is also unknown if TMS would reduce the feeling of fatigue as opposed to the ability of motor cortex to drive muscle activity. Further research will need to be done to see if TDCS or TMS treatment after an endurance event can offset the feeling of fatigue and enable better performance on subsequent endurance/power tasks.

(U) Final Notes

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) Additionally, it is important to note that these "undesirable" processes do serve important purposes. Pain is a signal that something is wrong with the body; reducing or eliminating it may lead to further injury. Similarly, fear is a signal to be wary and cautious; a soldier without fear may make poor decisions and endanger others. Finally, as noted in the pain section, care must be taken to ensure that other associated functions are not disrupted along with the targeted process. Disruption of activity in the somatosensory cortex, for example, would likely reduce pain sensation but also pressure, temperature, body position, and touch sensation. This would leave a soldier open not only to further injury due to pushing through an existing injury, but also perhaps reduced performance due to a lack of bodily sensation (for example, perhaps a loss of balance or an increase in clumsiness caused by reduced feedback on limb positions or touch while holding things). To the extent that emotions may be processed in similar areas, disrupting the feeling of fear may also reduce a soldier's ability to feel happy or angry, which may be desirable depending on the situation. As with many of the other topics discussed in the reviews produced so far, more focused experimentation is necessary to ensure that the desired results are obtained without side effects.

(b) (7) (E)
OGA

(U) [redacted]

(U//~~FOUO~~) [redacted]

(b)(7)(E)
OGA

[redacted]

(U) [redacted]

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) [redacted]

(b)(6)
(b)(7)(E)
OGA

(b)(6)
(b)(7)(E)
OGA

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) [redacted]

(U) [redacted]

(U//~~FOUO~~) [redacted]

(b)(7)(E)
OGA

(b)(6)
(b)(7)(E)
OGA

[redacted]

(b)(7)(E)
OGA

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) [redacted]

(b)(6)
(b)(7)(E)
OGA

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) [redacted]

(b)(7)(E)
OGA

(b)(7)(E)
OGA

[redacted]

(U) [redacted]

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) [redacted]

(b)(6)
(b)(7)(E)
OGA

(U//~~FOUO~~) [redacted]

(U/~~FOUO~~) [redacted]

(U//~~FOUO~~) [redacted]

(b)(6)
(b)(7)(E)
OGA

(b)(7)(E)
OGA

[redacted]

(U//~~FOUO~~) [redacted]

(U) [redacted]

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) [redacted]

(b)(6)
(b)(7)(E)
OGA

(U//~~FOUO~~) [redacted]

[redacted]

(U//~~FOUO~~) [redacted]

[redacted]

(b)(6)
(b)(7)(E)
OGA

(U) [redacted]

(b)(7)(E)
OGA

(U//~~FOUO~~) [redacted]

[redacted]

(U//~~FOUO~~) [redacted]

[redacted]

(U//~~FOUO~~) [redacted]

[redacted]

(b) (7) (E)
OGA

(U//~~FOUO~~) [redacted]

(U) Pharmaceuticals

(U//~~FOUO~~) Here we consider a utility of pharmaceu ticals for fatigue management, improving vigilance and memory and learning enhancements.

(U) Fatigue

(U//~~FOUO~~) Attention is resource demanding. If resources aren't available due to fatigue, it is more likely that attentional lapses will occur. Fatigue can occur due to two main causes: sleep deprivation and exercise, either physical or cognitive.

(U//~~FOUO~~) As reviewed by [redacted][31] sleep deprivation is well-known to cause decrements in a variety of cognitive tasks. They review a number of drugs in use (or close to in use) both for aiding sleep and for aiding vigilance when sleep is not possible. The sleep drugs are temazepam (Restoril), zolpidem (Ambien), and zaleplon (Sonata), and each is effective at aiding sleep in one manner or another based on desired sleep length and if the problem is sleep onset or sleep maintenance.

(U//~~FOUO~~) The more relevant discussion to this review is about the drugs used to maintain alertness when sleep is not a plausible option (although sleep is probably the preferred method of recovering alertness for a variety of cognitive reasons). These drugs are caffeine, Modafinil, and amphetamine.

(b) (6)
OGA

(U//~~FOUO~~) Similar to their discussion of sleep drugs, [redacted] determine that each of these is successful at maintaining alertness during periods of sleep deprivation, with the main difference being duration of effect. With the proper dosage levels, each can be equally effective in maintaining performance. Amphetamines, however, are a strongly controlled drug class and can lead to issues of abuse with larger doses or if used without proper monitoring. Modafinil, being a more recent drug, has received a fair amount of recent research.

(U//~~FOUO~~) Supporting [redacted] claim, [redacted][32] found that Modafinil given at one of two dose levels maintained performance on a vigilance task as well as a large dose of caffeine during sleep deprivation. [redacted][33] had healthy, non-sleep deprived subjects perform both a selective and a divided auditory attention task and gave them either a placebo, Modafinil, droperidol, or both modafinil and droperidol. Modafinil is thought to act by increasing dopamine whereas droperidol is thought to decrease it. The authors found that while droperidol did indeed hurt performance as expected, Modafinil failed to increase it; the drug did, however, offset the effect of droperidol in subjects who received both.

(U//~~FOUO~~) [redacted][34] examined the benefits of Modafinil in healthy subjects as well, compared only to placebo. The authors found that Modafinil appeared to aid performance on spatial tasks and response speed (but not accuracy) on the rapid visual processing task (RVIP) task; however, upon performing the task a second time, subjects who were given the drug saw some impairment due to distractibility and impulsivity, perhaps due to a general increase in arou sal. Thus while it seems that Modafinil is suitable for use as a sleep deprivation deterrent, it probably is not especially effective in non-deprived subjects.

and it may be no better than caffeine. Given the cost and availability benefits of caffeine, it is probably preferred to Modafinil.

(U//~~FOUO~~) Using the RVIP for the simulated UAV task, [redacted][5] demonstrated that misses increase during performance on the task, consistent with other sustained attention tasks. There could be two reasons that sustained attention tasks are taxing over time: a decrease in arousal due to the gaps of inactivity between brief response events, or a decrease in attentional resources as they are used up over time. [redacted] suggested that both are in play, and that using a second task would actually increase arousal in such a way as to replenish the attentional resources, as opposed to providing more strain and reducing performance.

(U//~~FOUO~~) Thus, in a strictly behavioral paradigm, subjects performed the task at baseline as well as with an intervention: when a miss occurred, the subject would either hear an alarm tone or perform one of two difficulty levels of an auditory working memory task (hearing three or four digits and deciding if they could be contiguously ordered). Each intervention was equally successful in reducing misses throughout the task, and rates of misses over the course of the task were equivalent (a rising miss rate is indicative of vigilance decrements). The alarm, however, was subjectively more frustrating and less tolerated than the working memory tasks. Thus a simple behavioral task could be used in addition to the vigilance task to maintain attention. However, the distraction provided by the second task is unknown because targets did not occur during performance of the intervention.

(b) (6)
OGA

(U//~~FOUO~~) There are also studies that have looked at the effect of various drugs on maintaining performance after exercise. [redacted][36] suggested that some of the known exercise fatigue effects are in fact cognitive in addition to physical. Adenosine is a chemical that is produced during exercise, is found in the brain, and is known to increase during wakefulness and decrease during sleep; thus higher levels are associated with worse attention. The authors injected either an adenosine agonist (to increase levels) or caffeine (which is an adenosine antagonist) into rats and measured how long they maintained treadmill running. They found that rats that received the adenosine agonist were generally less active and stopped running earlier than those who received placebo injections, while caffeine increased running time and partially counteracted the effects of the agonist in rats who received both. While there is no cognitive task performed here, the understanding of the fatigue mechanism (at least partially due to adenosine levels) is important, and could be used to generate cognitive supporters/enhancers.

(U//~~FOUO~~) Similarly, [redacted][37] cited research showing that exercise increases serotonin, which results in fatigue. Tyrosine is an amino acid precursor to dopamine and a competitor with tryptophan, the precursor to serotonin. Carbohydrates are also tryptophan competitors and provide muscles with energy. The authors provided bike riders with one of a placebo drink, a drink with carbohydrates, a drink with tyrosine, or a drink with both carbohydrates and tyrosine while they performed an endurance time trial. Carbohydrates were found to help performance, perhaps due to some cognitive effects as they did decrease the level of tryptophan in the blood. Tyrosine had a subjective effect of reducing the feeling of fatigue, but had no effect on biking performance.

(U//~~FOUO~~) Yet another exercise effect was found by [redacted][38] Exercise releases ammonia in the muscles, which may then travel to the brain. Ammonia is known to interfere with brain metabolism in general, and with GABA in particular, because its precursor, glutamate, is used to remove ammonia from the brain. The authors verified that exercise (bike riding) indeed raised levels of ammonia in both the blood and the cerebral spinal fluid, and that higher levels of ammonia were associated with higher perceived levels of effort. These effects, however, were somewhat mitigated by providing glucose, which reduced levels of circulating ammonia.

(U//~~FOUO~~) Finally, [redacted][39] provided soldiers with either placebo or carbohydrate drinks during a day of exercise (a 12 mile march and two 3 mile runs interspersed with meals and a four-hour nap break). Throughout the day, the vigilance of the soldiers was tested by requiring a response to a beep from a watch device. The higher of two levels of carbohyd rates kept vigilance performance at baseline levels throughout the day while the lower level provided protection early in the day but faded to placebo levels towards the end of exercise. Carbohydrates also kept feelings of alertness at baseline levels.

(U//~~FOUO~~) Overall, cognitive fatigue occurs due to sleep deprivation, sustained task performance, and exercise. Feelings of fatigue can be tied to levels of dopamine, serotonin, adenosine, and ammonia in the brain, and probably other substances as well. A variety of drugs have been shown to be effective in counteracting the build-up of these chemicals, including caffeine, carbohydrates/glucose, amphetamines, and Modafinil. However, caffeine and carbohydrates are the baseline drugs in this case due to their wide use and availability, and it is unclear if any of the drugs reviewed are any more effective. Furthermore, it is unclear if using these drugs in a non-decremented state provides any benefit. [redacted] failed to find any benefit from Modafinil over baseline on either a selective or divided attention task, and [redacted]. found some benefit but only on the initial performance of the task. So while the efficacy of existing alertness drugs for boosting attention is unproven, better understanding of how they maintain vigilance levels during sleep deprivation or exercise may provide a window for understanding the neural processes involved and developing drugs that could boost performance in the future.

(b) (6)

(U//~~FOUO~~) In summary, attention is known to decrease most severely after periods of task performance (even as short as 10 minutes), sleep deprivation, and exercise, although it also likely fluctuates on a shorter time period. There are a number of substances known to counteract these effects, although the most effective may still be caffeine and carbohydrates, and it is unclear if current drugs can boost attention in normal, non-decremented conditions. However, given that most conditions involving sustained attention will, by definition, lead to decrements in attention (because they are sustained, and will result in increased lapses over time), these substances are useful for maintaining performance.

(U//~~FOUO~~) Side note: while we have no studies on Ritalin in this review, Ritalin affects dopamine levels. Since tyrosine and Modafinil are thought to as well, the same conclusions for those drugs may be applicable to Ritalin.

(U) Improving Vigilance

(U//~~FOUO~~) The neurotransmitter dopamine is present in a large portion of the brain including the prefrontal cortex, basal ganglia, and substantia nigra. The effects it has in each of these regions, however, are not consistent. Research into ADHD, for example, suggests that the disorder results from the combined effects of a lack of dopamine in areas of the prefrontal cortex (dorsolateral prefrontal and anterior cingulate) and an excess in the basal ganglia (striatum). The lack of dopamine in the prefrontal cortex is believed to impair working memory and result in the attentional deficits associated with ADHD. In the striatum however the excess can cause hyperactivity. The difficulty lies in the fact that both areas express the gene for dopamine receptor D4, meaning that any drug which specifically targets that receptor will work to alleviate attentional deficits while increasing hyperactivity. Drugs designed for the treatment of attentional deficits, such as Modafinil, have been shown to have similar effects on normal healthy adults, improving attention and working memory in certain areas but often increasing distractibility. Examples of these effects can be seen in [redacted] where subjects demonstrated increased performance on spatial and response speed tasks in some trials and increased distractibility in others. The problem, therefore, becomes how to deliver a drug that will increase levels of dopamine within the prefrontal cortex while avoiding the striatum.

(U//~~FOUO~~) Current drugs work by increasing the amount of dopamine present in the synapse. Increased dopamine means more binding to D4 receptors and an enhanced dopaminergic effect in the affected region. Drugs like methylphenidate accomplish this by blocking dopamine active transporters (DAT) which are membrane proteins called symporters designed to reuptake dopamine into the neuron, regulating its effects. Other drugs, like Modafinil, increase synaptic levels of dopamine via mechanisms that are not entirely known. However, as long as all of these drugs are targeted to neurons with dopamine receptor D4 (DRD4) they will carry both the positive and negative effects associated with increased dopamine in those areas. They are also dependent on continued presence of the drug and have no lasting effects once the drug is out of the system.

(b) (7) (E)
OGA

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) In summary the ideas for how to enhance attention via drugs are readily available today; however, they require large advances in both our basic understanding of the exact areas to be targeted and large developments in delivery mechanisms. The ideas, nevertheless, are plausible and something that should be considered as our understanding of nanomaterials and viruses increase.

(U) Memory and Learning

(U//~~FOUO~~) While the brain is typically described as an electrical machine operating on neural impulses, it is just as much a chemical machine operating on neurotransmitters. Declarative memory has been most tightly tied to a process called long term potentiation (LTP), which is a state of excitability created in a neuron by a convergence of inputs; LTP can last for hours to days and leaves the neuron more likely to fire again to later input. Consistent with the statement earlier that the hippocampus seems to operate at the theta frequency, [redacted][40] found that LTP could be induced in the dentate gyrus (a portion of the hippocampus) by applying stimulation at the positive phase of the theta rhythm. LTP seems to be related to alpha-amimon-3-hydroxy-5-meth yl-4-isoxazolepropi onic acid (AMPA) receptors, which then have a downstream effect on N-methyl-D-aspartate (NMDA) glutamate receptors[41]; unblocking these NMDA receptors induces LTP. [redacted][42] showed that NMDA receptors are critical to hippocampal function using knock-out mice; mice that had NMDA function turned off were impaired at learning a water maze, which is known to be strongly hippocampus-dependent. As opposed to these early indicators of LTP, later stages are supported by protein synthesis, notably the cyclic adenosine monophosphate (cAMP) response element-binding (CREB) transcription factor. This stage is taken to be related more to consolidation of a memory as opposed to encoding.[41] [redacted][43] noted that there are a number of CREB variants which may have different functions, and suggested that CREB may act as a type of "gain control" on memory, such that the levels of different CREB variants is connected to the amount and timing of training needed for learning. Thus drugs that can enhance the operation of either AMPA receptors or CREB are potential memory enhancers. Lynch mentions that some preclinical studies have indeed found memory benefits in normal older adults using drugs based on AMPA function. [redacted][44] and [redacted] note that CREB manipulations in animals have been successful in enhancing memory; animals with increased levels of CREB learn in a highly reduced number of trials. Thus drugs that can operate at one of these three levels (AMPA, NMDA, CREB) are likely to be successful at enhancing memory, whether by speeding learning or actually increasing the amount of learning that can occur. [redacted] notes that the limits on these drugs are going to be similar to those of all drugs aimed at the brain: issues include passing the blood-brain barrier, defeating pumps (such as P-glycoprotein) that remove substances from the brain and thus lower the functional dosage, and other general toxicological and pharmacokinetic effects. These effects must be taken very seriously, as unintended changes to hippocam pal function could lead to amnesia or seizures as well as other drug effects.

(b) (6)
OGA

(U//~~FOUO~~) Other drugs have been tested for their memory effects. [redacted][45] looked at the drug donepezil, which is an acetylcholine esterase inhibitor. This drug increases the levels of acetylcholine in the brain, which has been shown to effect at least of one of memory or attention, if not both. [redacted] administered a battery of tests on attention, executive function, and memory both before and after a month of receiving either donepezil or placebo. The authors found that the drug group showed significant increases on only the episodic memory tests; working memory, executive function, priming, and attention were all unaffected. [redacted][46] and [redacted][47] performed very similar experiments looking at the effects of Modafinil. Modafinil is a drug that counteracts sleepiness; it is similar to amphetamine, but works through a different mechanism and does not have the negative side effects. Both groups had subjects come into the lab, take one of two dosages of the drug, and perform a battery of tests soon after (Modafinil reaches peak blood concentration a couple hours after consumption). [redacted] found that neither dose of the drug affected any cognitive task, but did cause some changes in mood. In contrast, [redacted] found some mood effects but also a number of cognitive effects, including better performance on a go-no go task, the Tower of London (a complex cognitive skill task), working memory span, and memory matching tasks. The difference in the two studies to lead to such a difference in results is unknown; [redacted] had a larger group of subjects and started testing sooner after drug administration than [redacted] it's possible that the drug's cognitive effects start sooner and wear off faster than anticipated. These drugs, and others like off-label Ritalin use, are more likely to have wide effects and perhaps operate more at the level of attention or arousal than memory per se, although boosted attention should also lead to better memory in general.

(b) (6)
OGA

(U//~~FOUO~~) In summary, pharmacological interventions are likely to have cognitive enhancing effects. The main issues with drug use are specificity and side effects. A drug like Modafinil, or even caffeine, is likely to have a general "upper" effect, which will affect a number of systems. Drugs specifically aimed at a particular molecular action, like AMPA receptors or CREB function, are likely to have more specific effects. In terms of further research aimed at enhancing cognition, beyond the general concerns involved in finding a drug to operate in the brain, the critical issues are targeting specific brain regions or systems known to participate in a particular function (the hippocampus for memory, the prefrontal lobes for decision making, and so forth) and ensuring that side effects specific to those functions do not occur; for example, memory is limited for a reason and creating a drug that lead to photographic or near-photographic memory would in fact be undesirable if used in any kind of consistent or repetitive manner.

(U) Signal Detection and Processing

(U//~~FOUO~~) Signal detection theory (SDT) assumes that people distinguish signals from noise based on the amount of evidence for the signal being present on a given trial. In the memory domain, SDT is often used in experiments of yes/no recognition. That is, you would study a list of words and then be shown a test list consisting both old (studied) and new (unstudied) words. In this case, the "noise" distribution is made up of new words while the "signal" distribution is made up of old words. The act of studying the old words increases their "strength" in memory, moving them higher on the evidence scale (for example, the evidence for the oldness of a word can be viewed as the strength of the match between the presented test stimulus and a memory). It is of little use to consider individual trials, so instead distributions are created from types of stimuli, such as old and new. Items have different strengths, which leads to the distribution. For example, high frequency words are generally stronger than low frequency words, and some words are simply more well-known by a particular person. When words are studied the group increases on the whole, but some words are well-attended and others more poorly, which means they aren't all increased by the same amount of strength, creating a wider distribution in the old group than the new. This is illustrated below. The most common SDT model assumes normal distributions. The graph below shows five distributions. One is the noise distribution, taken to be the strengths of various unstudied words. By convention, the noise distribution is set to have a mean of 0 and a standard deviation of 1. These numbers are arbitrary because the important values are the differences between means and the ratios of standard deviations, so 0 and 1 are convenient. Since it's a normal distribution, half of the noise items are below 0 and half above, and all the other expected normal distribution rules apply. The other four distributions are possible old item distributions, labeled by mean and standard deviation. Two such distributions are at a standard deviation of 1 and two with higher SDs; the typical number in memory research is a ratio of around 0.8, so 1.2 would be typical (1/1.2).

UNCLASSIFIED//~~FOR OFFICIAL USE ONLY~~

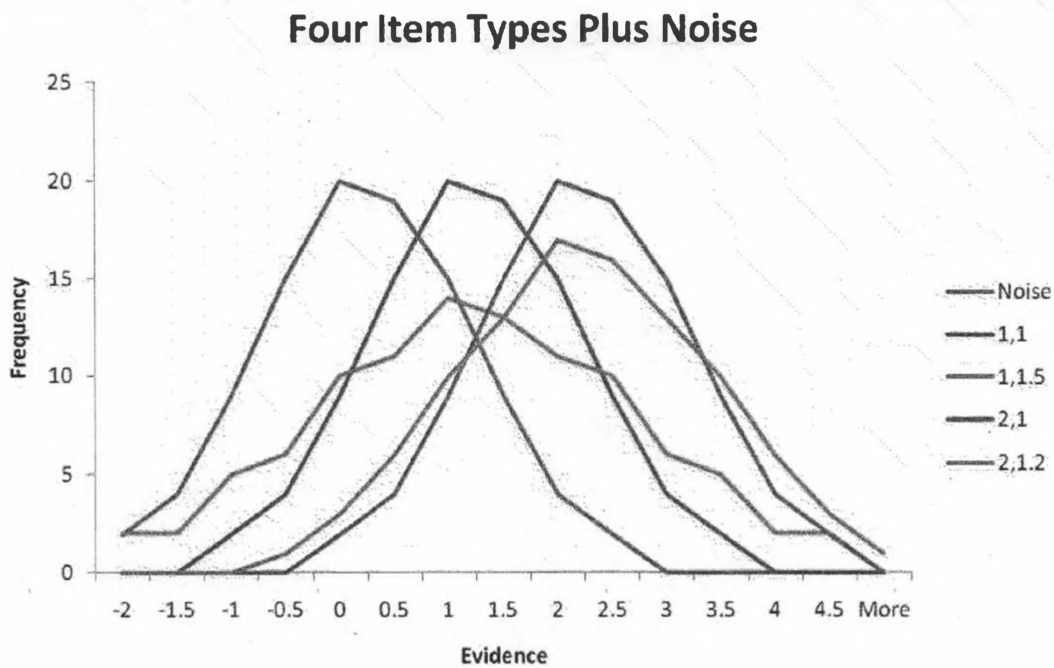

Figure A-2. (U//~~FOUO~~) Signal Detection and Processing

(U//~~FOUO~~) The way a person makes a yes/no decision is by setting a criterion. Criterion setting can be consciously controlled, but is generally a subconscious process. Criterion can be defined by the evidence level or by the likelihood ratio of old-new distributions; evidence level is easier. Let's say that you place your criterion at 1. That means that any time an item with strength 1 or greater comes up, you will say "old". If an item has strength less than 1, you will say "new". In the case of the 1,1 and 1,1.5 distributions, that means you will correctly say "old" to half of the items but incorrectly say "new" to half the items. This means your hit rate is 0.5 and your miss rate is 0.5. But, you will also say "old" to about 16 percent of new items (from the noise distribution) and "new" to 86 percent of those (the false alarm and correct rejection rates, respectively). So your overall accuracy would be (0.5+0.86)/2 = 0.73, assuming equal numbers of old and new items. One of the benefits of SDT is that it clearly defines the four types of responses; simply reporting a hit rate is common but unhelpful because the number of false positives (false alarms) is unknown. Actual memory is defined by having a higher hit rate than false alarm rate; if someone were just guessing, the rates would be equal.

(U//~~FOUO~~) If all you have is a yes/no test, the distributions above can only be estimated. The most common measure would be d' (d prime), which is the distance between the signal and noise distributions. More specifically, it is the distance in terms of standard deviations, and as such comes from the normal transform of the data. Since the distributions above were hand-created, the standard deviations are known: they are 1, 1, 2, and 2 in the cases above. Using raw data, it would be Z(HR) – Z(FAR), or normal transform of hit rate minus normal transform of false alarm rate. You can also derive the criterion from HR and FAR, and you can also estimate the area under the curve from a single HR/FAR pair, but it turns out this is a bad idea because it changes with criterion. Area under the curve has the benefit of being non-parametric, whereas d' assumes a normal distribution.

(U//~~FOUO~~) A better way to measure things is to collect data at a number of HR/FAR points. This can be done by varying payoffs. For example, offer $1 for every hit but take away 50 cents for every false alarm in one block and then reverse the payoffs in a second block. The first payoff structure will encourage "old" responses whereas the second will encourage "new"; this results in two HR/FAR pairs. With additional blocks and structures, more points can be plotted creating the received operating characteristic (ROC) curve. This curve describes an equal-accuracy line that gives the HR for any particular FAR. This is the true benefit of SDT: it separates accuracy from response bias or criterion (the willingness to say old or new). The more common way to create an ROC in memory research is to have subjects rate their confidence. High confidence responses will have lower FAR and misses, low confidence responses will have higher FAR and misses, and medium is in the middle. Each of the rated levels of confidence maps onto a different criterion. If a six-point scale is used, running from 1 = *high confidence new* to 6 = *high confidence old* response, then five criteria are used. The lowest one separates 1 from 2 responses (below is a 1, above is a 2); the next separates 2 from 3 (below down to the lowest criterion is a 2, above is a 3); the next 3 from 4, then 4 from 5, and finally 5 from 6 (below is a 5, everything higher is a 6). The underlying distributions can then be estimated from the data if a certain model (such as Gaussian SDT) is assumed. This allows for an estimate of accuracy (via a d' type measure or area under the curve), standard deviation of the old distribution, and criteria placement.

(U//~~FOUO~~) The ROC curve is simply a graph of HR/FAR pairs. With the known distributions above, this is easily done. Criteria are set and the HR and FAR rates are taken from the normal distribution as described above. Another example would have a criterion set at 0. This leads to 50 percent false alarms and 50% correct rejections from the new distribution, but 16 percent misses and 84 percent hits from the 1,1 distribution. Note, however, that the HR and miss rate change for the 1,1.5 distribution. This means that even though the distributions are equally far away from the noise distribution, the shape of the ROC changes. ROCs for each of the four old distributions above are plotted below.

UNCLASSIFIED//~~FOR OFFICIAL USE ONLY~~

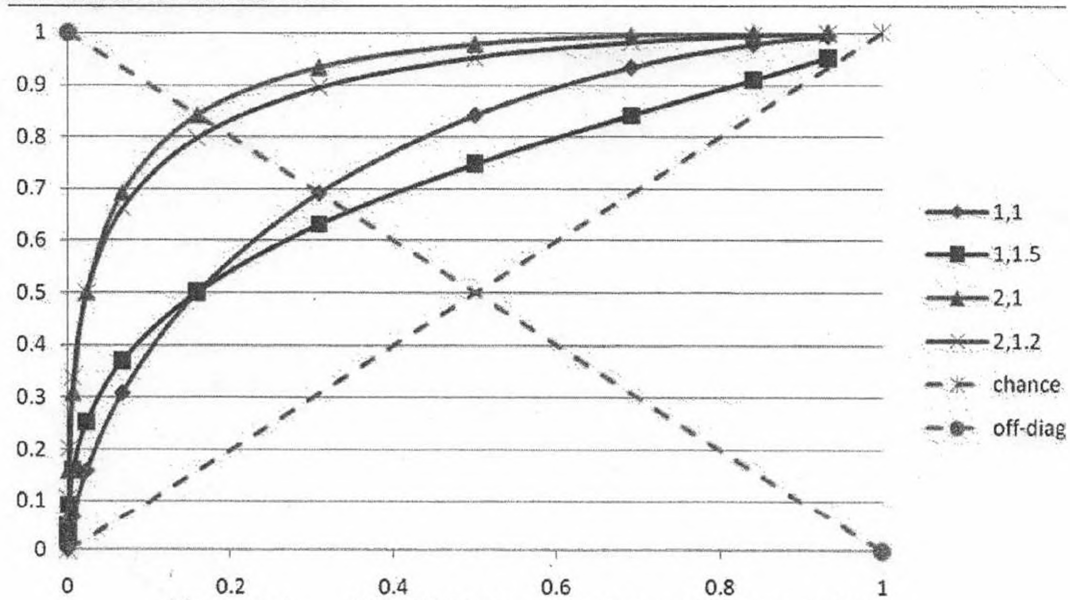

Figure A-3. (U//~~FOUO~~) Received Operating Characteristics

(U//~~FOUO~~) There are a few things to notice here. First, the two distributions with a mean of 2 are higher than those with mean 1. More specifically, they are further from the main diagonal. That diagonal is chance performance; any time HR and FAR are the same, regardless of their actual values, the subject is guessing. The closer to the top left corner, the better performance is. Secondly, the distributions with non-1 standard deviations are skewed to the left. This is easier to see with the 1,1.5 distribution compared to the 1,1. The 1,1 distribution is symmetric about the reverse diagonal (from 1,0 to 0,1) whereas the 1,1.5 distribution is skewed to the left. This results in it being above 1,1 at low FAR but below at higher FAR. It is also true for the mean 2 distributions; at the FAR=0.5 point 2,1 has a 2.5 percent higher HR but at FAR=0.006 2,1.2 is higher by 3 percent.

(U//~~FOUO~~) A number of measures of accuracy are available for ROC analysis. Two of the most popular are d' and A (area under the curve). The A values for the distributions above are 0.755, 0.708, 0.917, and 0.896. Thus it can be seen that the curves with higher standard deviations are viewed as equal to or worse than those with standard deviation 1 depending on the measure (note that d' assumes equal standard deviations, which is violated by two of the cases; an alternative measure using the normalized or zROC finds the same ordering).

(U//~~FOUO~~) However, many applications of discrimination require, or at least desire, low false alarm rates. This is broadly true in any case where the instance of a "hit" is very low. For example, say someone develops a test for cancer. Say that the test has a sensitivity (hit rate) of 99 percent and a specificity (correct rejection rate, or 1-FAR) of 99 percent. Then say that the test is given to 10,000 people, in which the incidence of cancer is 5 percent. The test would say that 590 of those people have cancer, but only 84 percent of those would. This isn't too bad overall, but notice the extremely high rates necessary to reach it. With the low (5 percent) incidence rate, a test with 97 percent specificity would quickly plummet to 63.5 percent hit accuracy. Instead of only 95 false positives, now there are 285. A corresponding drop from 99 percent to 97 percent hit rate means that instead of missing five cases of cancer, 15 are missed, but no extra false positives occur.

(U//~~FOUO~~) While it may be more desirable to have low miss rates with a cancer screening, there are cases where false positives are undesirable. In the case of lie detection, a false positive could cost someone his job, plus the expense of further investigation and possible litigation if the falsely accused sues. Courts also require a very low false alarm rate; they would much rather not catch those extra 10 people lost by a 2 percent HR drop as opposed to falsely accuse the extra 190 people caused by the 2 percent FAR increase. Thus for the purposes of lie detection, the desired performance measure is not just accuracy (HR + correct rejection/2), d', or area under the curve. Instead, performance should be measured as the HR at some standardized low FAR; say, 1 percent. For the curves above, HR at FAR=0.01 are 9.25 percent, 18.9 percent, 37.3 percent, and 39.3 percent. As was noted above, the non-1 standard deviation distributions have the advantage; the advantage increases when performance is lower (a curve with values 2,1.5 would have a HR of 41.4 percent, which is half the advantage over 2,1 compared to the 1,1.5 curve over 1,1). Thus it would be beneficial to create a test that not only discriminates well between liars and non-liars, but creates larger variability on some measure in the liars.

(U//~~FOUO~~) In summary, measures of accuracy are important in the realm of categorization, whether it be for memory, cancer screening, or deception detection. While SDT is an improvement over a report of overall accuracy, it is not optimal when judging performance on some tasks. It would be better to choose a false alarm rate, preferably a very low one, and report the hit rate at that criterion for the purposes of deception detection.

(U) Appendix B: Enabling Applications

(U//~~FOUO~~) There are a number of applications of emerging neurotechnologies that we focused on for the possibility of "off label" uses we will discuss in the final section of this forecast. These include cognitive state detection, cognitive state manipulation and enhancement, deception detection, and brain-machine interfaces. In this section, the state of the art for each of these is discussed, including both scientifically validated applications as well as capabilities that are claimed, but not scientifically validated.

(b) (7) (E)
OGA

(U) [redacted]

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) The overall theme for this application area is that you need to be able to detect an individual's cognitive state reliably in a real-time, operational environment. However, as each individual is different, and as existing methods for state detection (for example, EEG, fNIRS, fMRI) have significant technical difficulties from low signal to noise ratio, tendency towards artifacts, and sheer importability, there is significant work to be done in this area.

(U) Deception Detection

(U) Goals

(U//~~FOUO~~) The goal of the research area of deception (or lie) detection is to be able to tell with a high degree of certainty whether a certain piece of information is true or not. The most popular, if also most infamous, technology currently in use for that purpose is the polygraph. The polygraph measures physiological stress as a proxy for lying, but is known to be only somewhat better than chance (estimates range from 55 percent to perhaps 75 percent accurate, as it is used in the US, depending on conditions. Note that other countries use the polygraph under a different behavioral protocol and have thus reached higher levels of accuracy in detecting dishonesty.[48,49,50] Thus the polygraph has some amount of predictive ability, but nowhere near what is desired. Many other techniques have been and are currently being tested as replacements or addendums to the polygraph; they are discussed below.

(U) Common Elements

(U//~~FOUO~~) Before discussing different techniques, there are a number of points that apply to each and should be mentioned before getting into specifics.

(b) (7) (E)
OGA

(U) Research Goals

(U//~~FOUO~~) There are two primary avenues of research in deception detection. One is a theoretical understanding of lying. The goal here is to understand what lying is and how it is implemented in the brain, which would allow for the creation of a model and therefore allow for predictions of what lying "looks like" and how it might be detected. [redacted]

[redacted]

(U//~~FOUO~~) The other goal is the simple detection of lying, without regard for theory or reasons. This atheoretical approach may work faster than the theoretical approach, because of the disregard for basic research. However, it will be open to confounds, uncertainty across different conditions, and potential trickery (for example see vocal stress analysis [VSA], which "works" on a principle that may not exist) without a solid scientific theory to guide reasoning and predictions. Of course, if some technique provided evidence that is highly accurate, it would be foolish to not at least investigate its use even if a scientific basis was not immediately available.

(b) (7) (E)
OGA

(U) [redacted]

(U//~~FOUO~~) [redacted]

(U) The Measurement

(U//~~FOUO~~) It is critical to understand what each technology or technique is measuring to allow for at least some amount of theoretical knowledge of what exactly is being detected. For example, the polygraph is well-known to measure physiological variables like heart rate, sweating (via galvanic skin response), and breathing rate. These variables change with stress or anxiety, and are assumed to correlate with lying. However, since these are somewhat indirect measures associated with lying, and are also associated with stress that could be caused by a variety of sources, their efficacy is reduced.

(U//~~FOUO~~) [redacted] A marker that is commonly used in the literature is that of recognition; tests are constructed using materials that are sometimes random, and unlikely to be meaningful to anyone in particular, and sometimes related to the "crime" committed for the experiment (an item "stolen" from another room, for example). If a subject shows some sign of recognition of the crime-related item (which should not be the case for an innocent person), this is taken as evidence of guilt. There are a variety of ways of detecting recognition, as will be seen in the discussion of the various techniques. These markers (in general) may also differ

[redacted]

(U) Experiments

(U//~~FOUO~~) The vast majority of the research described here is typical of academic research, which is to say that it is done under low stress levels in controlled environments with college student subjects who are generally lacking in high levels of motivation. The application of this research, however, is likely to occur under high stress (for the investigator, the subject, or both) in potentially uncontrolled environments with people who are highly motivated to conceal what they know. [redacted] Additionally, almost all academic experiments are performed using individual, if not group, averages. That means that typically the best possible result is detecting that a person lied, but not necessarily about what piece of information. While tests are constructed to make a conclusion clear (for example, all the critical details tested refer to the same mock crime), it is unclear whether current paradigms are appropriate for deciding if a person lied about piece of information A but not piece B; that is to say, about individual trials or events.

(b) (7) (E)
OGA

(U//~~FOUO~~) Finally, as neuroscience becomes more popular and accepted by the general public, private companies are jumping into the research field more often. Being for-profit companies, their research (if it exists in some cases) is not published and their methods are often proprietary and may not be reproducible from available information such as websites and patents. This makes it difficult not only to verify their claims but also to evaluate the science in general. There is obviously a large demand for techniques to detect deception; if someone were to develop a good one, it would make perfect sense to keep that information private and attempt to make a lot of money off of it. The scientific community, however, largely views this as counterproductive and has a stake in dismissing, discrediting, or at least diminishing companies such as these for the very reason that their methods are secret. As such, it is possible that some particular method could be accurate or at least promising, yet be taken to task in the academic community. Groups interested in deception detection would be well served to not only pay attention to academic research but also to these private companies, especially looking for records of clients who could speak to not only the publicized success stories but also the failures.

(U) Methods for Deception Detection

(b) (6)
OGA

(U) Behavioral

(U//~~FOUO~~) There is some basic behavioral research into deception detection that has shown sensitivity to lying. One finding is that lies have, on average, a longer response time than truths,[53,54] although this may be paradigm-dependent (compare to [redacted][55]). However, the distributions overlap to such a degree that a single response time is unlikely to distinguish true responses from lies. One potential exception is in the case of speeded responses (note the error bars in [redacted]). A speeded response is one performed under some time deadline, such as being required to press a button in under a second. In this situation, responses tend to be "reaction times" rather than "response times", and it is more difficult to purposefully try to alter the response time (as a countermeasure). Accuracy may also be affected in this situation, allowing for some detection of a lie. It is important to note that [redacted]

[redacted] did not attempt to detect lies at the individual trial level, however. Thus in some situations at least, simple response characteristics may provide predictive information. (b) (6) OGA

(U//~~FOUO~~) Another behavioral method has to do with test construction. It largely is related to the concept of malingering, which is pretending to have some disease or exaggerating the symptoms. The concept can best be explained via amnesia. Imagine that someone is in a car accident and feigns memory loss to get a large insurance settlement. The person is tested with multiple choice questions about his life: what is the name of your wife, where were you born, and so forth. A malingerer will often studiously avoid the correct answers in order to show that he has no memory. However, someone with no knowledge of the person would still get some answers correct by chance (if each question has four answers, then approximately 25 percent correct). Thus someone attempting to hide knowledge will often perform worse than chance on a multiple choice test. Even someone aware of the concept will sometimes fail the test (if it is properly constructed) because the length of the test is unknown (thus making probability matching difficult) and a species-wide poor ability to estimate random responding. [redacted][56] demonstrated that such a test provides information that is statistically independent of, and adds predictive power to, that given by questioning with galvanic skin response. Namely, accuracy improved in one experiment from 73.5 percent to 78.5 percent and from 65 percent to 88 percent in a second where subjects lied about autobiographical details.

(U) Countermeasures

(U//~~FOUO~~) As mentioned above, it is possible for people to perform in such a way as to ruin any predictive power of response time (for example, simply by taking a long time to respond on every question). This can be countered to some degree by response deadline procedures, but this (as well as button pressing in general) requires a cooperative subject. Test construction can also be countered with better knowledge of randomness, but even statisticians are imperfect.

(U) The Future

(U//~~FOUO~~) There is some evidence that better behavioral procedures may boost the current accuracy of deception detection.[56] However, it is unknown how much of a boost should be expected in operational situations or how much of a boost is possible in general. Further research into test construction, especially with regards to malingering and question construction (avoiding questions or alternate answers that point to the correct answer or are easily dismissed; think of the SAT), should be beneficial to some deception detection situations. Behavioral testing does have the benefit of being portable (tests can be administered with paper and pencil or a laptop) and cheap (requiring only some software for stimulus presentation and response collection). Analysis is available in real time, but to the extent that response times must be averaged together or that accuracy is only usable over a number of trials, behavioral results cannot actually be used in real time; that is to say, the answer and response time to a particular question can be given immediately, but a single trial is insufficient to detect purposeful incorrect responding.

(U) Body Movement

(U//~~FOUO~~) There is some evidence that body movements, such as gesturing, posture, breathing, and so forth, offer some insight into lying. These are all likely related to the level of excitement or arousal that a person is feeling, although they may also be related to the level of cognitive effort being put into a discussion.[57] Typically, body movements decrease with lying because of the extra cognitive effort demanded by lying and because liars attempt to exert more control over themselves so as to not appear as liars.[52,57] However, it is known that the movement-based measures included in the polygraph (body shifting, breathing rate) are only somewhat predictive, and the evidence on other types of movement is mixed at best (for example, [redacted] found that only hand/finger movements, out of a number of variables, significantly differed between truth-tellers and liars). (b) (6) OGA

(U/~~FOUO~~) [redacted] found that asking subjects to tell their stories in reverse order did increase the difference between liars and innocent subjects on some physical dimensions, such as blinks and foot movements. They interpret these as signs of nervousness, however, so it is unclear how much they would contribute to deception detection, especially at the level of individual items. Accuracy can thus be expected to be in the range of the polygraph ([redacted] experiment includes a number of variables and will be discussed in a few sections of this document).

(b) (6)
OGA

(U) Countermeasures

(U/~~FOUO~~) As experience with the polygraph has shown, body movements are not especially predictive in the first place and furthermore can potentially be countered by people who know what body language is expected of innocent people. Thus body movement alone is unlikely to be of particular help in detecting deception. Furthermore, certain situations will make body monitoring untenable (such as a suspect in handcuffs, who cannot move his hands).

(U) The Future

(U/~~FOUO~~) While body movements are unlikely to be especially helpful, finding situations that would increase their accuracy (such as [redacted]) may be a reasonable goal. To the extent that body movements are helpful, the relatively recent technology called the Radar Vital Signs Monitor would assist in remote sensing (see *Educing Information* on page 67 in [redacted]). RVSM uses a high-frequency signal to detect body movements, which allows it to measure chest movements and eye blinks. The technology is portable but has limits under the use of radiation. The cost of an appropriate device is unknown, but a "flashlight" version suitable for breath detection is in the hundreds of dollars range.[58] Analysis may be available in real time (unknown).

(U) Eye Movement

(U//~~FOUO~~) Eye movements are known to change systematically with recognition.[59] Namely, fixations become longer and less frequent and the number of portions of a face or scene sampled also decreases. These effects don't require a specific task (that is, the subject need not be asked if they recognize the stimulus) and express themselves early in viewing and perhaps automatically. There is also evidence that blinks may be effective in distinguishing lying from truth-telling. Similar to response time, it is unclear if the distributions are separated enough to allow for item-level distinctions (this is not likely to be the case). Additionally, only one study has used eye movements (as opposed to blinks) to study deception, so further research is likely necessary. However, given paradigms used with other technologies, the fact that eye movements vary with recognition may prove to be very useful. The accuracy of eye movements in detecting deception is unknown.

(U) Countermeasures

(U//~~FOUO~~) Eye movements are only useful in terms of looking at some stimulus, such as a question or a picture on a computer screen; gaze direction (in terms of looking at an investigator or not) has no predictive power. So a subject could simply close his eyes or not look at what is being shown; eye movements also may not be useful in an interview setting. But, to the extent that all of these technologies may require a cooperative subject, this may not be an issue. The evidence pointing to the automaticity of some eye movement effects suggests that they may be particularly useful for evading countermeasures, but the lack of appropriate research means that countermeasures may simply have not been found.

(U) The Future

(U/~~FOUO~~) As mentioned, there is little research conducted on eye movements specifically regarding deception. But, the (at least somewhat) automatic recognition effects mentioned above may be a useful tool; they could be used as a marker of having seen some particular face or location. Further research is necessary to determine if there are countermeasures to eye movements (namely, exactly how automatic the effects are) as well as how predictive they may be at an item level. Eye trackers are portable, requiring only a special camera (one that sends out infrared light and can detect its reflection) and processing equipment, and are relatively cheap, having a cost of around $55K. Analysis in real time, at least for some measures (such as when a fixation occurs), is possible.

(U) Vocal Analysis

(U//~~FOUO~~) Voice analysis can be divided into two parts: characteristics of the voice itself, such as pitch, tone, and so forth, and the words used (for example, discourse or content analysis). Voice analysis (referring to the sound of the voice) has been used as a deception detection instrument since the 1960s starting with VSA. This method claims to measure micro-tremors in the vocal cords as an indicator of stress. This technique has not been found to have any predictive ability[60] [redacted] (b) (6) and the micro-tremors may in fact not exist. A more recent technique called LVA (for layered voice analysis) claims to analyze the spectrum of speech sounds for a variety of characteristics in a manner that allows for classification of statements as indicative of lying, arousal, attention, thinking, and even sexual attraction. However, the main producer of this method (Nemesysco) has been declared essentially phony by [redacted] and (successfully) threatened to sue to keep the results out of an academic journal. The technique has been tested by others[61] however and found to have no predictive ability in the hands of either the researchers or technicians from the company. Thus, while it is possible that some aspects of voice sound may have some predictive power, perhaps relating to stress, it is unlikely that any current technology will prove useful; VSA and LVA operate at chance.

(U//~~FOUO~~) Discourse analysis, on the other hand, has shown some ability to distinguish liars from truth tellers. Liars typically have trouble fabricating believable stories, so alibis may be inconsistent and in addition tend to lack sensory details (since the liar was not actually there). There are additional speech characteristics, such as slips and false starts that happen more often with lying.[57] These characteristics are only somewhat predictive, but may be useful. Since this is the last place that the [redacted] paper will be mentioned, the outcome of his study will be fully detailed: [redacted] found that telling a story in reverse chronological order lead to increased differences between liars and truth tellers in a variety of characteristics (eye blinks, speech patterns, body movements). Police officers with generally little interviewing experience then viewed tapes of the people who told the stories, some in order and some in reverse order. The officers who saw tapes of people telling their story in order were unable to distinguish the liars from the truth tellers whereas those who saw tapes with stories in reverse order performed above chance; however, accuracy was still only 58 percent. It is unclear what cues the officers may have been looking for, especially given their relatively low amount of experience; however, they rated the liars in the reverse condition as looking more nervous and appearing to have to think hard.

(U) Countermeasures

(U//~~FOUO~~) Virtually all of the research presented has been focused on college students or a similar population. Even when the subjects were given a cover story, there was little time to study the story or create one from scratch. Thus it is untested yet possible that someone with a prepared, well-rehearsed cover story would appear "natural" from a speech point of view.

(U) The Future

(U//~~FOUO~~) Future testing should be especially focused at the effect of practiced lies on speech. If it appears that some aspect of speech is still able to distinguish lies from the truth, then discourse analysis may be a tool used in addition with other techniques. The benefits are that analysis is likely to be cheap and portable, requiring only a recording device and a computer or trained coder, but a cost is that analysis of speech content (for example, words as opposed to pauses or pacing) will take time.

(U) Facial Scanning

(U//~~FOUO~~) Facial scanning is used primarily to detect emotion, which is then used as a marker for lying or deception. As such facial scanning is similar to the polygraph in that the investigator is looking for some mismatch between what is happening during a lie and what is happening while telling the truth; a person may smile but not with their eyes or display a flicker of anger. Most of the research in this area has been conducted by or based on work by [redacted] studies the movement of muscles in the face and has developed a system for categorizing the movements into different emotions or states, as well as combinations of those states. His research indicates that while the majority of people are very poor lie detectors, operating basically at chance, that he and his collaborators have been able to use recordings of people to detect liars at up to 80 percent accuracy.[61] This is based on a combination of body, vocal, and facial cues. There also appear to be a very small portion of the population that are consistently very accurate (in the 90 percent range) at detecting emotions. [redacted] has developed his facial system to the point that it can be taught in classes and even online.[62]

(b) (6)

(U) Countermeasures

(U//~~FOUO~~) To the extent that micro expressions (the very fleeting muscle movements that some people can detect and use to categorize emotions) are automatic, they may be transparent to countermeasure use. However, someone who was trained in [redacted] system may also have excellent control of their facial expressions and thus conceal their emotions. A more motivated subject might use Botox to reduce their ability to move the facial muscles.

(b) (6)

(U) The Future

(U//~~FOUO~~) Some facial recognition software already exists that is fairly good at detecting basic emotions in faces.[63] With either a computer system or a trained human, recognition should be fairly cheap, portable, and available in real time. However, to the extent that emotions are only a potential indicator of deception, their individual validity (as opposed to used in combination with other information) is unknown. Interestingly, [redacted][64] found that patients with left hemisphere lesions that impair their sentence comprehension were better able to detect liars from a videotape than normal controls or patients with right hemisphere lesions; this ability seemed to be due to facial processing. This finding has no immediate value (it would not be useful to give interrogators brain damage), but research into the basis of the result may provide information on what aspects of facial processing have become more developed in these patients.

(U//~~FOUO~~) A more recent advance has been to use thermal imaging of the face to detect deception. The theory is that the face will flush, particularly around the eyes, when a lie is told. A recent study[65] found that thermal imaging was able to distinguish liars from non-liars with 87 percent accuracy. Thermal imaging requires a special camera and analysis software, but should be comparatively cheap, portable, and offer fast (if not real time) analysis.

(U) Functional Magnetic Resonance Imaging

(U//~~FOUO~~) fMRI has been a tool of great interest recently in deception detection. It detects changes in blood oxygenation levels, which are interpreted as a marker of increased brain activity in the area of increased signal. fMRI is, by its nature, a comparative measure: activity in one state or during one task must be compared to another; there is no absolute measure of blood flow. As such, fMRI research has typically compared brain activity during trials of known lying to activity during trials of known truth telling. Studies typically find increased activity in regions of the prefrontal lobes and somewhat less consistently in the insula and anterior cingulate cortex.[51,54,55,66,67,68] The prefrontal regions are thought to be involved in cognitive control, executive control, and aspects of working memory and attention; the insula is thought to be involved in emotion and body state processing; and the ACC is thought to be involved in some aspect of error detection or processing. In general, lying is thought to involve using the same areas as telling the truth, but requires increased use of working memory, which is also affiliated with the prefrontal cortex.

(b) (6)

(U//~~FOUO~~) While some of this research has been strictly descriptive, others have attempted to use it specifically to detect liars. [redacted] used logistic regression on brain data in one set of subjects to create a classifier that categorized individual responses in another set of (four) subjects with 76.5 percent accuracy. In their paradigm, subjects chose one of two playing cards to lie about having, the other to tell the truth about having, and a variety of other cards displayed as distracters.

(U//~~FOUO~~) [redacted] (also with [redacted]) used a non-linear support vector machine (pattern classifier) on brain data in the same paradigm and achieved 88 percent accuracy both within and across subjects using a hold-out procedure.

(U//~~FOUO~~) Thus fairly high accuracy seems possible. However, the range of paradigms on which fMRI has been tested is limited. A variation on the Guilty Knowledge Test (GKT) has been the primary tool (this is also true of EEG). The few other studies that have looked at other paradigms or other kinds of lies find that there are brain activity differences there as well. For example, [redacted] had subjects describe a work and a vacation story that happened to them. They then developed and learned a false alternative scenario for one of those two categories. In the fMRI scanner, they answered short questions about the actual scenario, the false scenario (studied lies), and a false scenario in the unstudied category (spontaneous lies). The results indicated that spontaneous lies are more demanding than studied lies, although they activate largely the same brain regions.

(U//~~FOUO~~) [redacted] compared lying and truth telling about either autobiographical or non-autobiographical (semantic) facts and found that lying was only different from the truth in regards to autobiographical details. While it is a positive that these other studies find activity differences in largely the same regions as those that use the GKT, it is possible that other questioning scenarios may involve different brain regions or the same ones to a larger or small extent. Also, while the studies above show fairly good accuracy predicting across subjects, it should be noted that there are individual differences in brain activity during lying (for example, [redacted]).

(U) Countermeasures

(U//~~FOUO~~) Countermeasures to fMRI protocols are largely unstudied. Anecdotally, one might be able to fool the fMRI machine by "thinking hard" throughout the whole procedure (Mythbusters), which has some promise given that increased cognitive activity is the prime marker of lying. To the extent that predictive fMRI activity is indicative of markers used in other techniques (the insula or cerebellum might show higher activity related to increased physiological activity, such as measured by polygraphs), countermeasures to those techniques may also apply here. fMRI in general can be foiled by head movement and cannot be used on subjects with some types of metal in their body and is typically not used on people with claustrophobia.

(U) The Future

(U//~~FOUO~~) As mentioned above, only a few paradigms have been used in fMRI. Investigation of other paradigms will be important. For example, it is unknown how useful fMRI would be in an interview setting as opposed to a discrete answer (yes/no or single word answer) setting. Additionally, fMRI is often used in studies of memory. Applying other paradigms that focus more on recognition memory (such as mentioned in the eye movement section, and below under EEG) may open additional avenues of research. There are a variety of brain areas that demonstrate repetition priming, which is reduced activity to a stimulus that was recently seen or experienced. Priming is thought to be automatic, and thus may be particularly resistant to countermeasures.

(b) (6)

(U//~~FOUO~~) fMRI holds some promise in the arena of lie detection. Further development of algorithms (such as in [redacted] that separate lying from truth telling may be able to boost the predictive power of the technology, and complementary information from other technologies may bring accuracy [redacted] (b) (7) (E)

(U//~~FOUO~~) Even without field use, research using fMRI may still be useful in providing a model of lying. Once the model is established, other techniques could be used as field markers of brain activity. [redacted] showed, for example, that response time and accuracy correlated with certain brain regions' activity, and that subscales on a personality test correlated with other brain regions; [redacted][69] have shown that behavioral responses and galvanic skin response (GSR) correlated with activity in certain brain regions. It is possible that administering appropriate behavioral tests could stand in place of some fMRI tests. Other neuroimaging technologies are an obvious additional workaround, perhaps fNIRS in particular.

(U//~~FOUO~~) Additionally, it is worth noting that [redacted] are on the scientific boards of private lie detection companies Cephos and No Lie MRI, respectively. It is possible that better paradigms and/or analysis techniques have been developed than are reported in the scientific literature.

(b) (6)

(U) Functional Near-Infrared Spectroscopy

(U//~~FOUO~~) fNIRS is a relatively new technology that uses light to detect changes in the brain. It turns out that infrared light is particularly adept at penetrating the skin and bones, and that oxygenated and deoxygenated blood have different optical properties in this wavelength. Thus fNIRS is able to measure changes in blood flow and operates in a manner very similar to fMRI. However, since the measurement is based on light scatter and some light is absorbed by the brain, fNIRS can only provide information on the outer 2 centimeters or so of cortex and not any deeper structures. Only a couple of studies have investigated deception detection using fNIRS, and they largely replicate existing fMRI work. For example, [redacted][70] used a slight variation on the GKT task described previously and found increased prefrontal activity.

(U) Countermeasures

(U//~~FOUO~~) Countermeasures to fNIRS and fMRI are likely to be similar given they both measure blood oxygenation—at least in lying paradigms. However, fNIRS does not have the same requirements as fMRI regarding holding the head very still.

(U) The Future

(U//~~FOUO~~) While fNIRS is limited compared to fMRI, and thus may seem redundant at best, it does have the advantage of being more portable and cheaper. The [redacted] study investigated a limited area of the brain (essentially only the front of the head) using LEDs. Full-head fNIRS caps exist and their portability is an on-going area of research; in any event, they are likely to be as portable as EEG. fNIRS has the additional benefit of having greater temporal resolution than fMRI. While specific investigations of neural timing are missing in the deception detection literature, it is possible that timing could be a proxy for response time or some other distinguishing feature of lying. Despite its inability to measure activity below the top of the cortex, fNIRS has the potential to be the field equivalent of fMRI. Additionally, since the signal here is light as opposed to electrical or magnetic, fNIRS may be less disrupted by electromagnetic interference than fMRI or EEG. Thus, with improved algorithms and in combination with other measures, fNIRS may become a preferred method of deception detection.

(U) Electroencephalogram/ Event-Related Potential

(U//~~FOUO~~) Electroencephalogram (EEG) is the electrical signal produced by the brain. Event-related potential (ERP) is the EEG starting at some event (such as the onset of some visual stimulus) and typically averaged together over a large number of instances of that event (or category of an event, such as lies or previously studied words). ERP has limited spatial resolution, due to the inverse problem, but has very high temporal resolution. ERPs have been the other primary source of research into lie detection. This research focuses on a component of the ERP waveform called the P300, because it is a positive deflection that occurs around 300 milliseconds (or later) from stimulus onset. The P300 is sensitive to task-relevant, meaningful, relatively rare stimuli. The typical paradigm used in P300 research is called the oddball paradigm. One stimulus is declared to be "irrelevant"; no response is necessary if it is presented, but it is presented very often (say, 80 percent of trials). Another is declared to be the "target"; when this stimulus is presented, a response should be made, but it is presented rarely (the other 20 percent). The target, because of its task relevance and rarity, evokes a large P300 whereas the irrelevant stimulus does not; however, the absolute identity of the irrelevant and target stimuli is not important. A variation on this task has been used in lie detection, where a third category of stimuli exist called "probes." The probes are stimuli that should be relevant to a person who committed some crime, but would be irrelevant to anyone else.

(U//~~FOUO~~) The original research on this paradigm was conducted by [redacted] (1991).[71] They constructed crime scenarios involving a variety of pieces of information (a name, a location, and so forth). Each subject participated in one scenario (committed a crime) but was tested on two. Thus subjects would respond accurately to the irrelevant stimuli, targets, and probes for the scenario they were not involved in, but lie about probes for the scenario they were involved in. [redacted] reported that across two experiments (one involving 20 students in the mock crime scenario and the other involving four students who had committed different minor crimes), they were able to use the P300 to categorize a subject as guilty or innocent 87.5 percent of the time, and accuracy was perfect if a categorization was made. [redacted] has since expanded his work into a business called Brain Fingerprinting. [redacted][72] has strongly criticized [redacted] work. He has also shown[73] that the method is vulnerable to countermeasures; namely, performing some action to irrelevant stimuli (wiggling a toe, for example). This serves to make the irrelevant stimulus task-relevant. [redacted] found that using a countermeasure reduced lie detection accuracy from between 72 and 90 percent (the researchers used four different categorization techniques) to between 18 and 54 percent. Even worse, a second experiment showed that subjects trained on the countermeasure were harder to categorize even when told not to use the countermeasure on a later test. [redacted][74] developed a different paradigm for lie detection called the Complex Trial Protocol. In this paradigm, a probe or irrelevant stimulus is shown first and a simple reaction time response is collected (indicating that the stimulus was seen). A second or so later, the same stimulus would appear again in one of several colors, one of which was designated the target and responded to. This experiment was carried out over time similarly to the previous research with countermeasures to investigate the resistance of this paradigm to countermeasure use. [redacted] found that response time to the first stimulus (the "I saw it" response) was indicative of countermeasure use, and more importantly that the ERP waveform categorized lies with 100 percent accuracy without countermeasure use and 92 percent with. A second experiment with slightly different parameters reduced accuracy somewhat.

(b) (6)

(U//~~FOUO~~) [redacted][75] took a different approach by looking at ERP indicators of recognition. Specifically, there are known components associated with viewing previously studied (old) words that are smaller or not present when viewing unstudied (new) words. [redacted] had subjects perform an old/new test on words, responding either truthfully (saying "old" to old words and "new" to new), deceptively (the reverse), or randomly. [redacted] found that old and new words were distinguished in a variety of scalp regions, particularly over the parietal cortex (typical location of the parietal old/new effect, or late positive complex/LPC). While the researchers did not attempt to categorize responses using the waveform, it could have future value in lie detection.

(b) (6) OGA

(U) Countermeasures

(U//~~FOUO~~) As investigated by [redacted] in 2004, there are effective countermeasures to the GKT paradigm used with ERPs. However, subjects were told what countermeasures to use, so it is unknown how often or effectively countermeasures might be used in the field. The [redacted] paradigm is potentially less resistant to countermeasure use, as some aspects of remembering are automatic. However, it may be possible to develop a countermeasure, such as remembering something else while viewing a new or unknown stimulus. EEG in general, however, is sensitive to motion and other bodily electrical signals such as respiration, heartbeat, and eye movements/blinks. While these may not mask the P300 or LPC per se, they may make the signal too muddy to analyze at an individual trial level.

(U) The Future

(U//~~FOUO~~) ERPs, like fMRI, are very promising in the area of deception detection. While there is debate as the efficacy of existing methods (such as Brain Fingerprinting) and a range of reported accuracies, it is not unreasonable to tag the current state as better than the polygraph. If [redacted] 2008 technique proves to be generalizable, it may be of great interest (especially because it is publically accessible). Research into the area of recognition memory is also likely to be of great use. EEGs are relatively cheap, costing on the order of $50K to $100K and analysis can be relatively fast. Portability is an issue, as EEG recording is vulnerable to electromagnetic interference and so research is typically conducted in a shielded room. Recording is also usually done using a cap and conductive gel to boost the signal to noise ratio. Dry EEG caps that are more portable (and may even be wireless) exist (for example, the Emotiv Epoc headset), but noise is an issue. Thus the main barrier to EEG field use is technological; methods of boosting the signal to noise ratio and analysis speed would make EEG a very attractive choice.

(U) Other

(U//~~FOUO~~) Magnetoencephalography (MEG) is a technology that records the magnetic fields resulting from the electrical currents associated with neural activity. It has been shown in at least one case[76] to be able to detect lying at an individual level (during a "pay tax or smuggle" border crossing game) with accuracy in the mid-70 percent range. MEG is a nice technology in that it has almost the spatial resolution of fMRI coupled with almost the temporal resolution of EEG; it is sort of the best of both worlds. However, it is also the worst of both worlds in terms of being susceptible to electromagnetic interference [redacted] Work is currently being conducted on alternative methods for creating magnetometers sensitive enough to measure neural activity, but these technologies are still several years away from being available outside of specific laboratories, and they still require the use of a shielded room.

(b) (7) (E) OGA

(U//~~FOUO~~) In addition to neuroimaging, other evolving physiological measures that might hold some promise include genetic or hormone/blood testing. Genetic or hormone/blood testing will never be useful at the individual trial level; genetic testing for obvious reasons and blood testing because levels of biomarkers change too slowly. However, to the extent that they may serve as markers of other predictive variables, such as stress, personality, propensity for lying, and so forth, they may be useful as field tests for categorizing or profiling the type and state of the person to be questioned.

(U) Summary

(U//~~FOUO~~) In summary, there is no single technique or technology that is currently thought to distinguish the truth from lies with much more accuracy than the polygraph (topping out at perhaps 90 percent). Some have the potential to perform better, but appropriate testing is limited at best and it is unknown if a single technique is capable of the high accuracy desired. It is important to decide which research goal is more critical (or whether to pursue both): theoretical understanding of lying or simple detection. In either case, it is highly likely that multiple techniques will be necessary for characterizing what markers distinguish a lie from the truth under a variety of conditions. The use of converging methods is critical in theoretical research today, and multiple methods will provide additional predictive power in the applied detection of deception.

(U) Recommendations

(U//~~FOUO~~) Research in the areas of fMRI, EEG, fNIRS, behavioral testing, and perhaps facial scanning/eye movements is specifically recommended. fMRI and EEG are the two most developed technologies in terms of current research into deception detection and have peer-reviewed accuracy in some situations of around 90 percent, which is far better than the polygraph. Research in these areas should specifically look at generalization to other populations and paradigms, particularly interview situations with experienced liars. Research should also focus on paradigms that exploit known subconscious/automatic effects, such as priming, which should prove more resistant to countermeasures.

(U//~~FOUO~~) Once advances are made in these areas, parallel research should be conducted using fNIRS. EEG and fNIRS have the best probability of use in the field once technological hurdles have been cleared in portability, noise/interference issues, and analysis speed. Behavioral testing, both in terms of test construction and personality typing, will be useful for detecting deception as well as categorizing subjects. People are known to vary in a variety of important ways, and their ability to lie is one of them. As noted previously, it is very likely that some subjects will be nervous while other will not; some will have prepared stories and others will not. The particular markers that detect deception in each of these groups will differ from each other. Finally, facial scanning or eye movements are useful as potentially orthogonal measures that can add predictive power to models of lying. Each is portable and relatively cheap and may tap into somewhat automatic processes (particularly eye movements), and more importantly has shown predictive power on its own. While it is unknown if it is possible to detect deception with an arbitrarily high accuracy (say 99 percent), the more valid measures that are accessible, the better the chance there is of that goal being achieved.

(U) Improving Technologies for Deception Detection: What Would It Take to Make a Deception Detection Technology Better/More Portable?

(U) Functional Magnetic Resonance Imaging

(U//~~FOUO~~) Functional magnetic resonance imaging works by using a large magnet to align the spins of susceptible molecules in the body, then disrupting them with a radio-frequency signal and measuring the resulting signal as they relax back into alignment. Because oxygenated and deoxygenated blood have different magnetic properties, the comparative signals they produce can be used to track blood flow in the brain, which is taken as an indicator of neural activity in specific brain regions. fMRI can also measure other things so long as they have a magnetic component; for example, diffusion tensor imaging (DTI) measures water flow through white matter tracts via water content.

(U//~~FOUO~~) Currently, MRI magnets suitable for human use generate a field of 1.5 or 3 Teslas. Stronger magnets exist, but are not typically used in research due to cost and side effects (stronger fields cause reports of headaches and nausea). This requires a large electromagnet which is super-cooled by liquid helium. This makes fMRI technology expensive (on the order of millions of dollars) and non-portable; it also eliminates subjects with some kinds of metal in their body. Additionally, most MRI rooms are shielded to reduce outside signals. fMRI used for cognitive tasks also relies upon the averaging of a number of events of a similar type (that is, studied words that are later remembered) as a way to increase the signal by averaging out random noise.

(U//~~FOUO~~) fMRI is unlikely to become better in any field-applicable way. MRI machines can be constructed with permanent magnets or non-super-cooled electromagnets, which would increase portability, but they also produce weaker and less-consistent fields, which reduces the signal. A reduced signal can be compensated for by using longer relaxation times, but this trade-off is probably incompatible with the idea of field use (for example, even slower experiment timing is completely undesirable). [redacted] The major drawback to fMRI is the portability issue; this could only be overcome by developing some method of creating strong magnetic fields with subtle radio-frequency generators and receivers.

(b) (7) (E)
OGA

(U) Voice/Discourse Analysis

(U//~~FOUO~~) Voice analysis includes both sound characteristics (such as pitch) and content. There are many existing technologies for recording sounds, even from a distance away. The current drawback to voice analysis is that no existing method of analysis has been shown to have predictive power when it comes to deception detection. It is difficult to say if this is because none exists or because the major companies in the field have happened to skip over the basic research, lending an air of disreputability to the idea in general. Some level of legitimate, basic research should answer this question.

(U//~~FOUO~~) [redacted]

(U) Electroencephalograph

(U//~~FOUO~~) The EEG works by recording the small electrical field generated by groups of neurons firing. While this is a more direct neural signal than fMRI, the complete dynamics of what EEG records are still not well understood (whether it is pre-synaptic or post-synaptic firing, what these might mean different from each other or different from the potential change within a cell, and so forth). Due to the physics difficulty in deducing electrical sources from the field measured at the scalp, the spatial resolution of EEG is fairly poor. However, the temporal resolution is very good, working at the millisecond level. The EEG signal itself can be analyzed for frequency information, such as the power at alpha, theta, etc, or the signal can be "cut up" and locked to events (such as displaying a stimulus) and averaged together to create an ERP. An ERP is then analyzed for components, or deflections in the waveform that correlate with a putative cognitive function. Once established, a component can then be used to make assumptions about the presence or degree of some kind of processing (for example, mnemonic), although this analysis is typically strongly tied to the paradigm that generates it.

(U//~~FOUO~~) The current typical EEG system can be described by the number of channels, or recording electrodes, that is uses. Research systems use anywhere from 16 to 256 channels. The signal and spatial resolution increase with more channels, although there is a limit due both to limited space on the head for electrodes and the correlated activity of electrodes near each other. Denser recording systems increase the cost; a good 128 channel system costs around $100K. Current limits to EEG include the ability to amplify and record the electrical signal and analyze it. Because of the small magnitude of the signal (on the order of microvolts), brain activity is typically overwhelmed by outside sources and physical electrical signals such as eye movements, respiration, and heartbeat. This requires a good recording, usually aided by scraping the scalp to remove dead cells and using conductive gel to lower the physical resistance, as well as the use of signal amplifiers. In addition, any ERP waveform is then generated by averaging together a number of individual waveforms, typically at least 15-20, to average out the random electrical noise and draw out the brain signal.

(U//~~FOUO~~) EEG has a strong potential to become more portable and field-usable. The technical drawbacks are recording in a non-obtrusive way that doesn't harm the signal to noise ratio. For example, scalp scraping and conductive gel may not be desirable in all circumstances, and in fact some "dry" EEG systems do exist. However, the cost is reduced signal quality. This could potentially be overcome with better amplifiers located at the recording channel, but these are currently very expensive (on the order of hundreds of dollars per channel). The signal issue is especially important for ERPs, similar to as discussed with fMRI; lower signal means more trials needed to average together. An alternative is to record the EEG from inside the skull; the skull and scalp smear the electrical signal, so moving inside of them reduces the noise. However, this is obviously invasive. Localization even then is still not perfect, but it is also untenable to place electrode arrays throughout the brain. EEG analysis can be conducted in real time or near-real time, but ERP analysis is slower due to averaging; this time loss can be reduced if fewer trials are needed to generate the average. Thus the most desirable outcome for EEG would be to find a way to record dry EEG from a relatively small number of channels and transmit the information wirelessly to an analysis device. If a solid or powdered conductive agent (as opposed to a gel) were developed, this could aid in reducing signal loss. Alternatively, if some relatively non-invasive manner of recording from within the skull were developed, perhaps targeted nanorobots, the signal would be much greater and better spatial resolution could be achieved.

(U) Facial Analysis

(U//~~FOUO~~) Facial analysis works by correlating specific muscle movements with known emotional states. For example, fear is denoted by a widening of the eyes and other stereotyped movements. With training, people can learn to recognize even split-second muscle movements and thus detect fleeting or concealed expressions of fear, anger, surprise, and so forth. These emotions can then be used to decide if a person is lying or telling the truth with some degree of accuracy.

(b) (6)
OGA

(U//~~FOUO~~) Current facial analysis is almost exclusively based on the work of [redacted]. While Ekman's research has found that the majority of people are very poor at telling when someone is lying or not, he has also shown that people can be trained to become better. There is also work with computer recognition software which indicates that computers can be trained to recognize these microexpressions as well.

(U//~~FOUO~~) The current limitations on the use of facial analysis are the training of either people or computers to become accurate microexpression detectors. More critical, however, is research correlating how these expressions can be used to detect lying. People vary in their response to telling a lie, and people differ in how practiced and willing to lie they are. Thus people will very likely differ in their emotional response, and thus facial expressions, while lying. Furthermore, emotions will likely never be direct consequences of lying (for example, the emotion will be a secondary feature, similar to the physiological responses measured by a polygraph). Research needs to be conducted both on these individual differences as well as the correspondence between lying and microexpressions in general to [redacted]

(U) Eye Movements

(U//~~FOUO~~) Eye movement technology works by shining an infrared light on the eye. This light is then reflected by the pupil and the retina back to the camera, which uses the two to calculate the direction of gaze. The principle measures from eye tracking are then fixations (where the eye lingers as it takes in information) and saccades (the jumps made between fixations during which very little information is acquired). It can also measure pupil dilation and blinks. Eye tracking is used heavily in linguistics for studies of reading, but also somewhat in studies of memory. As with fMRI and ERP, trials are averaged together. This is not due to "signal" issues, as it is easy to plot out where the eyes have traveled on a single display, but in order to get averages of the characteristics (fixation duration, number of fixations) on one type of stimulus compared to some other type (famous versus never-seen faces, for example). The eye characteristics will depend on a number of factors, including salient points of the display and the task being performed by the viewer, and so averaging is necessary to differentiate classes of stimuli.

(b) (7) (E)
OGA

(U//~~FOUO~~) Current eye tracking systems differ in a number of ways. They can be placed on a table in front of the subject or head mounted on a cap or headband. They may consist of a single camera or multiple cameras. There are even versions that are essentially contact lenses that work by detecting their own movement, and thus the eye (assuming that the lens does not slip). Eye movements can also be measured, although perhaps more noisily, but recording electrical activity around the eye (electrooculogram: EOG). Thus there are a number of potential trade-offs. A head-mounted camera may better account for head movements, but may be uncomfortable. A table-mounted camera may be more convenient but have issues with head movement and distance (since the camera deals with angles, distance is important).

(U//~~FOUO~~) Current limitations in eye tracking are related to the particular system used. In general, eye movements would be useful for lie detection in their use as measures of recognition memory, which are thought to be mostly automatic. There is anecdotal evidence that gaze changes with cognitive processing in general (for example, looking up to the left denotes accessing visual memory). [redacted] are available instantaneously, measures that require averaging obviously will not be. Research that allows for the development of algorithms may lead to fast processing that detects recognition with only a few repetitions. An alternative is to develop a discrete contact lens-type eye tracker, although this may be unlikely given the need to also transmit data to some analysis computer (the lens may always be noticeable). Additionally, absolute measurement of gaze is contingent on a registration procedure where the subject is asked to look at a series of predetermined locations in order to standardize the eye tracking; otherwise, eye movements are known only in a relative sense. If absolute measures are needed (such as knowing that the subject is looking at point A on a display and not point B), then a more covert manner of registration should be developed.

(U) Positron Emission Topography

(U//~~FOUO~~) PET works in a manner similar to fMRI. A radioactive tracer is injected into the subject, and the emitted particles are then detected by receivers placed around the head (this is typically a torus-shaped system, like an MRI machine). The direction and number of particles is then used to determine their spatial source. When the tracer is radioactive water or a compound that binds to glucose, it can be used as a measure of blood flow and thus acts just like fMRI. However, PET is limited in that the particle counts are accumulated over time, and so temporal resolution is even worse than fMRI, and spatial resolution isn't quite as good. For these reasons, as well as the invasiveness of injecting a tracer, the radioactivity, and the non-portability of the apparatus, PET is a very poor choice as a lie detector. However, it is mentioned here because if the receivers could be made portable, and appropriate radioactive tracers were available, PET could be somewhat useful. Better tracers would have shorter half-lives to allow for faster imaging, although this is limited by the time it takes for the tracers to reach the brain and then the regions receiving differential use (for example, the more active areas of the brain). Shorter half-lives also requires that the tracer be created closer in time to use; thus, most research PET machines also have a facility capable of creating the tracers on-site. Were these obstacles overcome, however, PET could be useful in addition to fMRI through the creation of tracers that bind with neurotransmitters: it might be useful, for example, to visualize dopamine use throughout the brain. fMRI cannot do that.

(U) Body Movements

(U//~~FOUO~~) Body movements are gestures, fidgeting, and so forth. They can be measured by people trained to code video tape. Some body movements have been correlated with lying, such as hand movements. However, much like with facial analysis, the question is if the movements are somehow related to lying or if they are a downstream measure that correlated with anxiety or some other trait.

(U//~~FOUO~~) Body movements currently can be measured by trained coders or to some degree with remote sensors like radar. Analysis is comparative, as with eye movements (for example, movements can be detected immediately, but are likely only to mean something at an aggregate level compared to some other subset of movements).

(U//~~FOUO~~) The main limitation on body movement analysis is the level of research. It is unknown how well body movements alone correlate with lying, what they mean on their own (if a particular movement means anxiety or fear), and how they may differ across people.

(U) Physiology

(U//~~FOUO~~) Physiology is meant to include bodily measures such as breathing rate, heart rate, galvanic skin response, and so forth. These measures are included in the polygraph, and as such are known to have some predictive power. However, they are secondary measures, being directly related to stress or anxiety and thus only somewhat to lying. Also, the polygraph is known to be beatable. However, under some circumstances the polygraph or associated measurements may be useful.

(U//~~FOUO~~) Current polygraph technology relies on direct contact, such as pulse meters and electrodes for measuring GSR. Improvement on this design lies primarily on making remote sensors, such as the Radar Vital Signs Monitor (RSVM). This machine uses high frequency signals acting as radar to detect changes in chest movement as a proxy of breathing and heart rate, and can even detect blinks.

(U//~~FOUO~~) To the extent that physiological measures are desirable, the best advances would be along the line of remote sensing. The RVSM is a good step in this direction, but is limited by its radiation effect. A system that allowed for more extended use would be desirable.

(U) Functional Near-Infrared Spectroscopy

(U//~~FOUO~~) fNIRS works by shining near-infrared light through the scalp into the brain. The tissue of the brain then interacts with the light in predictable ways and the phase changes can be detected when the light scatters back to detectors on the scalp. Much as with fMRI, oxygenated and deoxygenated blood and tissue can be detected because of their different optical properties, and so neural activity can be imputed from their relative levels. However, due to absorption, fNIRS can only be used to image the outer layer of cortex in the brain.

(U//~~FOUO~~) fNIRS is functionally similar to fMRI in terms of the information it provides. However, the temporal resolution can be much better, and the system can be more portable. The system itself is much more similar to EEG in that sensors are placed on the scalp; however, not as many channels are possible because both light emitters and detectors are necessary, and a certain amount of distance is needed between them to allow for computation of the path the light took.

(U//~~FOUO~~) fNIRS is similar to EEG in its future applications. Primary advances will be in making a portable system that can image the whole head; portable systems that image only a portion of scalp are currently available. Some manner of scanning deeper into the head would be of great interest; perhaps this could be accomplished by using stronger light sources or by somehow placing them inside the skull (although this would obviously be highly invasive). The ability of fNIRS to detect other substances, such as neurotransmitters, would also be desirable as mentioned under PET. Since fNIRS works at a much better timescale, if it could be used to detect useful neurotransmitters, PET would be completely obsolete. Anything that boosts the fNIRS signal would also be useful for averaging purposes as discussed with earlier technologies.

(U//~~FOUO~~) Another way to utilize light for detecting cortical activity is called event-related optical signal (EROS). Due to swelling of active cortical areas light scatters differently from those resulting in altered intensity (relative to inactive areas). This difference in scattered light intensity affords a map of cortical activity. The EROS technique has the same limitations as fNIRS (shallow probing area, two-dimensional image).

(b) (7) (E)

(U//~~FOUO~~) [redacted]

(U) Cognitive State Management

(U) Goals

(U//~~FOUO~~) The goal of the research area of enhancing vigilance is to find ways to improve attention in normal subjects as well as counteract fatigue and attentional lapses in subjects who have undergone long periods of activity, such as those under the effects of sleep deprivation. This document is focused on possible methods of increasing baseline levels of attention as opposed to sustaining attention against the decrements caused by fatigue. There are two sections, technological and drug interventions.

(U) Technological Interventions

(U//~~FOUO~~) As discussed in the vigilance summary section, attention is thought to work at the level of the brain via top-down signals from a control network in the frontal and parietal regions sent to the sensory processors of interest. An example would be that if the important feature to look for is something red, then the fronto-parietal network would enhance V4, the portion of visual cortex thought to be responsible for color processing. This enhancement is thought to have two forms; a base line increase in activity before any stimulus is presented, and then increased response to a matching stimulus compared to that same stimulus if it is not to be attended. This is particularly true if there is also a stimulus to ignore near the stimulus to be attended, as inhibitory activity also occurs (attention raises the activity to a single stimulus as well, but not as much). However, it is possible that the baseline increase is related only to the expected spatial location of the stimulus and not the content of that stimulus; there is evidence that baseline shifts in activity fail to predict response to the stimulus itself. The exact mechanism that creates these enhancements is unclear, but it is thought to be primarily due to cholinergic projections from the fronto-parietal regions to the upstream regions. It may also be reflected in synchronization between different processing regions; synchrony of neural firing between brain regions may function as a tag or mechanism for noting which distinct features (color, motion, shapes, and so forth) go together in a single object or percept.

(U//~~FOUO~~) Given that the target of attention is basic sensory processing regions, it seems that the best method for boosting attention would then be to attempt to stimulate these regions, to make them more receptive to incoming stimuli. However, the evidence showing that baseline activity increases are unrelated to stimulus-processing activity makes this an unattractive choice. A better option may be to simply increase the functioning of the fronto-parietal attention regions in the first place and assume that they will target the proper sensory regions as they would normally. The two mechanical methods currently available to do this are TMS and TDCS.

(U//~~FOUO~~) TMS has been applied to the parietal lobes in a few studies and has demonstrated effects on spatial attention. Because the hemispheres have mutual connections, researchers have been able to apply TMS to one parietal hemisphere and affect the deployment of attention to the other side. The parameters of stimulation are important; low frequency repetitive TMS (rTMS) applied to the right hemisphere would aid attention to the right visual field by releasing the left hemisphere from inhibition (there is a cross-over in spatial attention such that the left hemisphere directs attention to the right visual field) whereas high frequency rTMS would aid attention to the left visual field. But, applying TMS to parietal regions has been shown to affect earlier sensory regions while that feature was present (medial temporal lobe activity was different if motion was present along with TMS versus without TMS). The main issue with using TMS would be to determine for sure a) what brain location should be stimulated for the particular task to be accomplished, b) what parameters aid performance, as opposed to inhibiting function (which seems to happen often with TMS), and c) if the enhancement lasts beyond the application of TMS and for how long.

(U//~~FOUO~~) TDCS, on the other hand, seems to have much more consistent and beneficiary effects on a variety of tasks. Anodal TDCS almost always seems to benefit performance on tasks via excitatory stimulation whereas cathodal TDCS does the reverse. TDCS also has specific effects on different tasks depending on electrode placement; putting the electrodes over the temporal cortex improves word recognition memory but not attentional cueing, and placing them over motor cortex improves motor skills. Since the frontal and parietal regions involved in attention are relatively near the cortical surface, they should be amenable to TDCS interventions. The main issue with TDCS would be to determine a) the proper electrode placement to aid performance, and b) if the enhancement lasts beyond the period of treatment. Research on a motor skill learning task showed that TDCS aided learning, but performance dropped off at the same rate afterwards as it did in subjects who did not receive TDCS. This may not be an issue, since the better training results in higher performance even after the drop-off, but it is important to note that the changes may not be permanent and retraining may be necessary; it is also unclear how long changes may last for cognitive as opposed to motor abilities.

(U//~~FOUO~~) In summary, the best chance for vigilance enhancement seems to be in stimulating the attentional network itself as opposed to the target of that network. This is due to research on the neural mechanisms of attention as well as the fact that some sensory processing regions are too deep in the brain to be affected by current stimulation technologies. TDCS is currently more promising than TMS, although with research either one should prove effective. Research into improving attention with either method should be possible immediately; improvements in the process will come from determining the optimal stimulation and training parameters. If the enhancement effect is such that it would be desirable to make a portable system (due to short-lived effects or some other reason), TDCS would likely be preferred because all that is really needed are electrodes and a power source. A more invasive option would be an implanted device that could be activated remotely, although this would be inflexible because the brain placement would be permanent. While TMS and TDCS are the major brain stimulation methods currently in use, other methods could be imagined ; nanorobots could be developed that can directly stimulate specific neurons or groups of neurons, for example. The hurdles with a technology like this would be specificity (the technology may work at a more precise level than the theories are created at; fMRI and EEG measures neural populations as opposed to individual or small groups of neurons) and targeting/activation (getting the robot to the proper location and turning it on and off).

(U) Drug Therapy

(U//~~FOUO~~) The neurotransmitter dopamine is present in a large portion of the brain including the prefrontal cortex, basal ganglia, and substantia nigra. The effects it has in each of these regions, however, are not consistent. Research into attention-deficit hyperactivity disorder (ADHD), for example, suggests that the disorder results from the combined effects of a lack of dopamine in areas of the prefrontal cortex (dorsolateral prefrontal and anterior cingulate) and an excess in the basal ganglia (striatum). The lack of dopamine in the prefrontal cortex is believed to impair working memory and result in the attentional deficits associated with ADHD. In the striatum however the excess can cause hyperactivity. The difficulty lies in the fact that both areas express the gene for dopamine receptor D4, meaning that any drug which specifically targets that receptor will work to alleviate attentional deficits while increasing hyperactivity. Drugs designed for the treatment of attentional deficits, such as Modafinil, have been shown to have similar effects on normal healthy adults, improving attention and working memory in certain areas but often increasing distractibility. Examples of these effects can be seen in [redacted][34] where subjects demonstrated increased performance on spatial and response speed tasks in some trials and increased distractibility in others. The problem, therefore, becomes how to deliver a drug that will increase levels of dopamine within the prefrontal cortex while avoiding the striatum. (b) (6) OGA

(U//~~FOUO~~) Current drugs work by increasing the amount of dopamine present in the synapse. Increased dopamine means more binding to D4 receptors and an enhanced dopaminergic effect in the affected region. Drugs like methylphenidate accomplish this by blocking dopamine active transporters (DAT), which are symporters designed to reuptake dopamine into the neuron, regulating its effects. Other drugs, like Modafinil, increase synaptic levels of dopamine via mechanisms that are not entirely known. However, as long as all of these drugs are targeted to neurons with DRD4 they will carry both the positive and negative effects associated with increased dopamine in those areas. They are also dependent on continued presence of the drug and have no lasting effects once the drug is out of the system.

(b) (7) (E)
OGA

(U//~~FOUO~~) Most drugs reduce unwanted side effects in the exact way that attention deficit drugs have attempted to do, through receptor subtype specificity. [redacted]

(U//~~FOUO~~) In order to affect a more lasting result it is necessary to change the structure of the neuron itself. In order for this to be accomplished a method would have to be developed to induce expression of DRD4 gene, located on chromosome 11. Unfortunately, there are no second messenger pathways that result in increased expression of the gene and coding for the DRD4 protein, making standard pharmaceutical approaches inapplicable. Many researchers believe that the "long" variant, one of several polymorphisms in the DRD4 gene, predisposes individuals to ADHD and other attentional deficits. Other polymorphisms may be in part responsible for individuals with superior baseline attention, coding for [redacted]

(U//~~FOUO~~) In summary the ideas for how to enhance attention via drugs are readily available today; however, they require large advances in both our basic understanding of the exact areas to be targeted and large developments in delivery mechanisms. The ideas, nevertheless, are plausible and something that should be considered as our understanding of nanomaterials and viruses increase.

(U) Summary

(U//~~FOUO~~) We currently have access to enough knowledge on the workings of the prefrontal cortex to attempt to modify its action, at least in some domains. These modifications could be aimed at improving attention at least for limited periods of time; it is unknown if any current methods short of gene therapy or possibly years of practice can change stable cognitive abilities like attention in normal subjects. Drug therapy is an attractive option with some existing options such as dopamine or Ritalin, but the drugs have too broad an action to be applicable right now. Developments in drug targeting are key to improvements [redacted] plausible mechanisms for altering attention and are available now, but research is needed to ensure that behavior on attentional tasks is altered during stimulation, to find the stimulation parameters that create beneficial change, and the long-term effects (both positive and negative) of such a treatment. Studying the long-term effects will be important for any mechanism of brain change, especially drugs. It is known, for example, that Parkinson's patients who are treated with L-dopa develop a tolerance and eventually the side-effects outweigh the benefits of treatment. Research will be needed to ensure that even infrequent but repeated treatment via drug or brain stimulation does not lead to tolerance, dependence, or undesirable changes in other aspects of the subjects' identity.

(U) Cortical Enhancements/Disruptions

(U) Improving Vigilance

(U//~~FOUO~~) As discussed in the vigilance summary section, attention is thought to work at the level of the brain via top-down signals from a control network in the frontal and parietal regions sent to the sensory processors of interest. An example would be that if the important feature to look for is something red, then the fronto-parietal network would enhance V4, the portion of visual cortex thought to be responsible for color processing. This enhancement is thought to have two forms; a base line increase in activity before any stimulus is presented, and then increased response to a matching stimulus compared to that same stimulus if it is not to be attended. This is particularly true if there is also a stimulus to ignore near the stimulus to be attended, as inhibitory activity also occurs (attention raises the activity to a single stimulus as well, but not as much). However, it is possible that the baseline increase is related only to the expected spatial location of the stimulus and not the content of that stimulus; there is evidence that baseline shifts in activity fail to predict response to the stimulus itself. The exact mechanism that creates these enhancements is unclear, but it is thought to be primarily due to cholinergic projections from the fronto-parietal regions to the earlier regions. It may also be reflected in synchronization between different processing regions; synchrony between brain regions may function as a tag or mechanism for noting which distinct features (color, motion, shapes, and so forth) go together in a single object or percept.

(U//~~FOUO~~) Given that the target of attention is basic sensory processing regions, it seems that the best method for boosting attention would then be to attempt to stimulate these regions, to make them more receptive to incoming stimuli. However, the evidence showing that baseline activity increases are unrelated to stimulus-processing activity makes this an unattractive choice. A better option may be to simply increase the functioning of the fronto-parietal attention regions in the first place and assume that they will target the proper sensory regions as they would normally. The two mechanical methods currently available to do this are TMS and TDCS, as discussed in the Cognitive State Management section above.

(U//~~FOUO~~) In summary, the best chance for vigilance enhancement seems to be in stimulating the attentional network itself as opposed to the target of that network. This is due to research on the neural mechanisms of attention as well as the fact that some sensory processing regions are too deep in the brain to be affected by current stimulation technologies. TDCS is currently more promising than TMS, although with research either one should prove effective. Research into improving attention with either method should be possible immediately; improvements in the process will come from determining the optimal stimulation and training parameters. If the enhancement effect is such that it would be desirable to make a portable system (due to short-lived effects or some other reason), TDCS would likely be preferred because all that is really needed are electrodes and a power source. A more invasive option would be an implanted device that could be activated remotely, although this would be inflexible because the brain placement would be permanent. While TMS and TDCS are the major brain stimulation methods currently in use, other methods could be imagined ; nanorobots could be developed that can directly stimulate specific neurons or groups of neurons, for example. The hurdles with a technology like this would be specificity (the technology may work at a more precise level than the theories are created at; fMRI and EEG measures neural populations as opposed to individual or small groups of neurons) and targeting/activation (getting the robot to the proper location and turning it on and off).

(U) Memory and Learning

(U//~~FOUO~~) As with vigilance, the most promising neurotechnology for enhancing memory and learning involves stimulating different parts of the brain with either TMS or TDCS.

(U//~~FOUO~~) [redacted][77] had subjects learn a list of word pairs and then either stay awake for 90 minutes or sleep for the same period, which allowed for one REM-SWS (rapid eye movement—slow wave sleep) cycle. Subjects also either received transcranial direct current stimulation (TDCS) or "placebo" where the electrodes were applied but no current was sent into the brain. Subjects who slept showed a small improvement in memory, but those who received TDCS showed a larger improvement. Subjects who stayed awake, even if they received TDCS, showed no improvement or a slight decrease in memory. Subjects who received TDCS while sleeping also showed decreases in theta (4-8 Hz), alpha 1 (8-10 Hz), and beta 1 (15-20 Hz) EEG activity compared to those who simply slept. These EEG rhythms will be important to later discussion of neurofeedback. This study used stimulation set to 15 seconds on, 15 seconds off. In a similar study, [redacted][78] showed that TDCS set to the slow wave oscillations associated with SWS (frequency of .75 Hertz, although actual sleep can also extend into the delta 1-4 Hz range) lead to a larger increase in cued recall (remembering one word from a pair given the cue of the other member of the pair) than just sleeping or than applying theta stimulation. Importantly, each of these studies included a procedural memory test and showed no change in performance between presenting an odor cue[79] or applying TDCS [redacted] or not; thus the effect is specific to declarative memory.

(b) (6)
OGA

(U//~~FOUO~~) Another method of stimulating the brain is transcranial magnetic stimulation (TMS). TMS works by creating a strong magnetic field focused at a particular brain structure, which then induces a current in the brain. This current is essentially the externally-generated firing of a population of neurons. Thus, TMS can create neural activity in a non-invasive manner. However, it is unknown exactly what parameters of TMS will create what response, and if it differs across brain regions. [redacted][80] showed that rTMS (repetitive TMS) can disrupt function. They applied rTMS at 10 Hz while subjects performed a change blindness study (a display is presented, then a quick mask, followed by either the same display or the same display with a change in it; it has been found that the mask causes changes to often go undetected even if they are relatively large and obvious without the mask). Applying rTMS to the right parietal cortex caused poorer performance than left parietal or no TMS; it was thought that the TMS disrupted attention or working memory processes critical to conscious perception of the change. On the other hand, [redacted][81] found that rTMS at 10 Hz applied to the superior parietal lobule during the delay period of a location or object working memory task lead to increased accuracy on the spatial task (parts of the parietal lobes are thought to be important to spatial working memory in particular). Furthermore, [redacted] showed that there was a range of effects of rTMS: for some subjects, the rTMS reduced power in the alpha range and these subjects showed increased accuracy. For other subjects, however, the rTMS increased power in the alpha range and decreased performance. [redacted][82] found that 7 Hz rTMS applied to the left prefrontal cortex during encoding of words lead to increased later recognition of those words, but found that semantic judgments made at encoding were slowed. They thus concluded that the rTMS actually disrupted encoding, but caused subjects to reprocess the words in such a way that it benefited later memory. Finally, in an offline paradigm, [redacted][83] applied rTMS at 5 Hz to the prefrontal cortex of older adults before performing a name-face learning task (the paradigm is called "offline" because stimulation is applied outside of performance of the task as opposed to during). The authors found that rTMS increased associative memory and increased activity in frontal and visual areas compared to a baseline learning block and subjects who received sham rTMS.

(b) (6)
OGA

(U//~~FOUO~~) In summary then, it appears that TMS can both help and hurt performance on declarative tasks. In some cases, as suggested by [redacted] it can help later performance by hurting current performance. It is possible that TMS works by causing the brain, at least locally, to adopt activity at the frequency being applied to it [redacted] found affects in the alpha band while applying rTMS in the alpha frequency), but this is unknown. In general, TMS is plagued by unknowns: what frequency is appropriate for a given task or brain region, when and how long should it be applied, how strong the stimulation should be, how long the affects last, will it help or hinder processing, and so forth. What is known is that rTMS can indeed aid declarative memory in some cases; further research will be needed to determine how best to do so on a consistent basis, however.

(U//~~FOUO~~) TDCS, on the other hand, seems potentially more likely to provide benefits immediately. Research on sleep has already shown that performance can be enhanced by stimulating the brain at a certain sleep stage. Comparing the two, TDCS may be the more useful technology; it does not require specific brain localization to be effective and does not require the generation of a strong magnetic field (although TMS is still comparably portable). However, with more research, TMS may be preferred because its localized effect may result in more circumscribed cognitive effects.

(U//~~FOUO~~) Brain stimulation will be limited in its application to procedural memory (memory for how to do things that is not usually encoded primarily linguistically—for example, how to ride a bike), mostly because of the variety in procedural skills; the brain regions involved in motor skill learning will not be the same as those involved in cognitive skill learning, and they likely will not respond to the same kind of stimulation. However, there is evidence that stimulation can improve procedural learning. [redacted][84] applied TDCS to motor cortex of subjects learning a motor adaptation task (exerting different levels of pressure between thumb and index finger) and found that subjects who received the TDCS improved faster on the task than subjects who did not receive TDCS. Interestingly, this improvement seemed to be due to a decreased loss of training between sessions; improvements during the session were equal in the two groups. However, once training ended, both groups declined in performance at equal rates. However, due to the better performance of the TDCS group during training, their performance was still better than at the beginning of training even three months later.

(b) (6)
OGA

(U//~~FOUO~~) In summary, brain stimulation likely holds promise for speeding procedural learning, if not boosting it above normal levels. All of the critiques from the declarative memory section hold here, however; the characteristics (frequency, strength, location, and so forth) of the stimulation are likely to be unknown at this point. With more research however, the [redacted] study indicates great promise for brain stimulation as a method for improving procedural memory.

(b) (7) (E)
OGA

(U) Analogy Making

(U//~~FOUO~~) [redacted]

(b) (6)
OGA

(U//~~FOUO~~) [redacted] used rTMS during analogy making. Subjects saw a pair of displays either simultaneously or sequentially. One was a source picture depicting colored, patterned shapes. The second was a target picture. In the analogy condition, subjects judged if the target picture kept the relations from the source picture; changes could be either spatial or object based. In a "literal" condition, subjects decided if the target picture was the same as the source picture; again, changes could be spatial or object-based. While performing the tasks, subjects received rTMS at 5 Hz during the first, middle, and last 10 seconds of test blocks. The locations could be over a left PFC region identified as more active in a similar task during PET, the right PFC homologue, or over left motor cortex. There was no effect on response times from right PFC rTMS; left motor cortex rTMS improved response times in the sequential literal and both analogy conditions; and left PFC rTMS improved response times in both analogy conditions. The authors conclude that the left PFC is important to analogy making because the effects were limited to analogy making in that case. However, the motor cortex effects are concerning.

(U//~~FOUO~~) [redacted][87] gave rTMS to severely depressed patients along with a battery of neuropsychological tests both before and after treatment. Treatment consisted of 20 5-second trains of 10 Hz rTMS to the left PFC per session for 10 to 20 sessions over a few weeks. Subjects were given the same tests both before and after treatment, and scores of individual tests were combined into composites reflecting general cognitive abilities (memory, attention, working memory/executive function, and motor speed). Subjects who responded to the treatment by lowering their depression scores performed worse (overall; not in terms of improvement) on attention but better on motor speed than those who did respond to the treatment. In general, subjects improved on the composite working memory, memory, and motor speed tasks. However, this study is marred by the change in mood of subjects over the course of the experiment (mood is known to have cognitive effects), the lack of different versions of the tests, and the lack of a control group. It is very possible, and in fact likely, that the improvement in scores is due to a practice effect.

(b) (6)

(U//~~FOUO~~) [redacted][88], however, performed a similar experiment with a sham TMS group. Again using depressed patients, subjects were given a battery of tests before and after treatment. Subjects received five sessions of 20 Hz rTMS to the left middle frontal gyrus in 20 2-second trains or sham stimulation for five minutes. Interestingly, both groups improved equally in mood over the course of the experiment. Neither group improved on any test except for the rTMS group on Trails B, a test of switching ability. However, the authors used conservative non-parametric tests, and so some small effect may be present. Across these two studies then it appears as though rTMS has no effect on cognitive processing, although the primary goal was to assess the ability of rTMS to reduce depression without hurting cognition. It is possible, however, that rTMS to the frontal lobes does increase some executive function (in the Trails B test). Frontal executive functions are related to analogy making, so this could be a promising result. Similarly, [redacted][89] reviewed a number of TMS studies across a variety of paradigms and concluded that the data suggested that TMS could be used to increase normal capabilities. [redacted] only cites th[redacted] study in regards to analogy making, but cites studies in other domains that have found enhancement in normal subjects as evidence that TMS would be viable with further research.

(b) (6)

(U//~~FOUO~~) [redacted][90] used TDCS on subjects while they performed a probabilistic classification task. This task is called the weather prediction task; subjects see simple symbols on the screen in some combination (square and triangle, triangle, square and circle and diamond, and so forth) and guess if they mean "rain" or "sunshine." The subject then receives feedback, and over time learns to associate the symbols with either rain or shine. The symbols are typically probabilistic, however; the square might indicate sun 85 percent of the time, for example. This reduces the ability of subjects to make explicit connections between symbols and the weather outcome. Subjects were given TDCS for 10 minutes; after five minutes the task started and the final block was completed after stimulation was turned off. Stimulation was administered over either the left PFC or the occipital lobe and was either anodal or cathodal (or sham). Subjects who received sham stimulation to either location learned over the course of trials, improving from chance (50 percent) to 60 percent in 50 trials. Subjects who received occipital stimulation showed no difference from sham stimulation. Subjects who received anodal stimulation over left PFC showed improved learning in the fourth block of trials, while those who received cathodal stimulation showed numerically but non-significantly lower performance throughout the experiment. Thus it appears that TDCS can improve implicit learning on a task that involves generalization and classification learning. It should also be noted that 50 trials of this task is not many, and that greater effects might have been found with further training. Importantly, however, is the conclusion that TDCS improved a process that relies on the PFC.

(U//~~FOUO~~) In summary, studies where rTMS has been applied to the prefrontal cortex have been equivocal on their effects on executive function, but that could be due in part to the fact that they were not intending to enhance performance per se, and stimulation may have been focused on regions other than those that support analogy making. In contrast, a behavioral intervention with older adults and a TDCS study improved performance on prefrontal-related tasks, suggesting that analogy making could indeed be enhanced in normal adults. Combined with other data on other tasks showing the beneficial effects of TMS and TDCS, it would be worthwhile to put more direct effort into studies aimed at improving analogical reasoning.

(U) Disrupting Negative Processes

(U//~~FOUO~~) While TMS and TDCS have been used to enhance desirable processes, it is less obvious if they could be used to disrupt undesirable processes. In this review, those processes will be fear, pain, and fatigue. There is some positive evidence in this direction due to the known beneficial effects of TMS on depression; if TMS treatment can reduce a trait like depression, then it seems likely that it might affect other negative traits.

(U//~~FOUO~~) Fear is perhaps the simplest of these three processes because it is typically localized to one brain region, the amygdala. While prefrontal cortex has also been implicated as a control region, it is somewhat unlikely that the emotional control regions are different from control regions that are used to guide other behaviors in the face of distraction or competing options. fMRI studies that use fearful stimuli, such as images of snakes and spiders or faces expressing fear, almost always find increased activity in the amygdala compared to images of neutral stimuli or faces. Also, studies of patients with amygdala damage find that they have difficulty detecting or categorizing fearful faces and perhaps even further that they may not experience fear. The amygdala is a bilateral structure located deep in the brain, in the medial temporal lobes. As such, it is outside the range of effect of TMS and probably TDCS (unless it can be indirectly suppressed via connections to cortical regions). A more likely avenue of effect is to enhance the prefrontal control regions that allow for completion of a task in the face of fear.

(U//~~FOUO~~) In contrast to fear, pain is a more complicated process. Painful or irritating stimuli evoke activity in a variety of brain regions, including the amygdala, nucleus accumbens, putamen, anterior cingulate, insula, orbitofrontal cortex, and somatosensory cortex. It is likely that some of these regions are not directly related to pain sensation directly but instead to control mechanisms or emotions associated with experiencing and dealing with pain. Of these regions, the only ones likely to be affected directly by TMS or TDCS are the orbitofrontal and somatosensory cortices (although note that [redacted] 2002 coil was designed with the intent of activating the nucleus accumbens). There is promising evidence of modulating pain sensation if activation could reach the anterior cingulate, however. [redacted] used real-time fMRI as biofeedback to teach subjects to raise or lower their subjective feeling of pain. Subjects were stimulated with a laser to a point where they rated the pain a 7 out of 10. The anterior cingulate has been suggested to be important to both pain sensation and pain regulation in previous studies, so it was used as a source of feedback. Subjects were shown their ACC activity in real-time (actually slightly delayed due to the hemodynamic response, but otherwise unprocessed) and asked to either increase or decrease the signal using one of a few techniques (attend to or away from the pain, focus on the stimulated body part or a different body part, and so forth). After only a few 30-second blocks, subjects were able to do so, and this appropriately increased or decreased their subjective ratings of the pain. The same result was found in subjects with chronic pain who were not stimulated with the laser. Control conditions, which failed to find significant pain modulation, included subjects with chronic pain who used heart rate, galvanic skin response, and breathing rate for biofeedback; normal subjects who were given the attention techniques but not biofeedback; normal subjects who were given one specific technique with twice as much training; normal subjects who were given feedback from a different brain region; and normal subjects who were shown a different subject's ACC activity. Thus the positive effect of the

(b) (6)

biofeedback seems to be due to controlling ACC activity. If TMS or TDCS could be used to affect activity in this region (perhaps using [redacted] coil), then it seems plausible that pain sensation could be decreased. It is also possible that pain sensation could be decreased by decreasing somatosensory activity in general, although this may have the side effect of reducing all bodily sensations (pressure, temperature, body position, and touch in general).

(U//~~FOUO~~) Even more complex is the feeling of fatigue. Most neuroscience studies of fatigue involve making voluntary muscle contractions, such as squeezing a force gauge or contracting the biceps as hard as possible. As such, the conclusions likely differ from those that would be drawn from studies of long distance runners or other endurance tasks. However, fatigue has been tied to the sensorimotor cortex, supplementary motor cortex, prefrontal cortex, and the cingulate gyrus. Previous reviews looking at articles more related to endurance also suggest that fatigue in the brain could be due to neurochemical modulations, such as reduced glucose and increased ammonia as a result of muscle activity. Fatigue obviously also occurs in the muscles themselves as resources like glucose are used up. Furthermore, there is an interaction between the brain and the body such that the brain can be urging full output when the muscles cannot provide it and alternatively the brain can feel tired and drive rested muscles at less than optimal levels.

(b) (6)

(U//~~FOUO~~) A review by [redacted] investigated the effects of applying TMS to motor cortex at different points during muscle fatiguing tasks (again, typically short-term exercises as opposed to endurance tasks). Motor cortex appears to become more excitable during and after muscle contractions, meaning that TMS has a larger effect (as measured by the motor evoked potential) and can be used to maintain cortical drive to the muscles. However, it is important to note that the motor evoked potential (MEP) differs both across subjects and across muscle groups. For example, additional muscle fibers continue to be recruited during a biceps contraction up to 75 percent of maximum MEP, whereas most fibers in the hand are already recruited by 30 percent MEP. Thus it seems likely that TMS would be differentially effective depending on the targeted muscle group. Additionally, since most of these tasks use short bursts of voluntary contractions, it is unclear if TMS could be used to reduce fatigue during, say, a long march. It is also unknown if TMS would reduce the feeling of fatigue as opposed to the ability of motor cortex to drive muscle activity. Further research will need to be done to see if TDCS or TMS treatment after an endurance event can offset the feeling of fatigue and enable better performance on subsequent endurance/power tasks.

(U) Final Notes

(U//~~FOUO~~) In summary, TMS and TDCS appear to be plausible techniques for reducing the effects of pain and fatigue, and perhaps fear. Most of the evidence is circumstantial, however, as direct tests of these theories were not found in the literature search. So long as cortical regions such as prefrontal, motor, and somatosensory cortex are the basis of some process, they should be amenable to modulation via TMS or TDCS. Prefrontal cortex will likely be a key region as it is broadly involved in control processes and attention and thus may be a target for the enhancement of "drive" or "willpower." However, as noted in the fatigue section, it is unknown if modulation will affect these processes themselves or the perception of them or both; for example, muscles may be driven to perform better when fatigued, but it is possible that a subject will feel just as tired if not more so as his effort continues.

(U//~~FOUO~~) Additionally, it is important to note that these "undesirable" processes do serve important purposes. Pain is a signal that something is wrong with the body; reducing or eliminating it may lead to further injury. Similarly, fear is a signal to be wary and cautious; a soldier without fear may make poor decisions and endanger others. Finally, as noted in the pain section, care must be taken to ensure that other associated functions are not disrupted along with the targeted process. Disruption of activity in the somatosensory cortex, for example, would likely reduce pain sensation but also pressure, temperature, body position, and touch sensation. This would leave a soldier open not only to further injury due to pushing through an existing injury, but also perhaps reduced performance due to a lack of bodily sensation (for example, perhaps a loss of balance or an increase in clum siness caused by reduced feedback on limb positions or touch while holding things). To the extent that emotions may be processed in similar areas, disrupting the feeling of fear may also reduce a soldier's ability to feel happy or angry, which may be desirable depending on the situation. As with many of the other topics discussed in this document, more focused experimentation is necessary to ensure that the desired results are obtained without side effects.

(U) Brain-Machine Interfaces

(U//~~FOUO~~) Here we explore the world of brain-machine interfaces (BMI, also known as. neuroerogonomics) where human brains drive machines. Both invasive and non-invasive applications are considered.

(U) Goals

(U//~~FOUO~~) The goal of the research area of neuroergonom ics is to develop methods for using the brain to control devices instead of pressing a button or even reading muscle activity from the body. This is a highly complicated field because it requires knowledge and integration of ps ychology, biology, and engineering. Neuroergonomics is a subfield of human factors, understanding how people interact with their environment, and an example of BMI. Most of the academic research, especially that involving invasive surgery, is related to restoring motor and/or communication ability to people with movement disabilities like spinal cord injuries or ALS. More recent research, especially involving noninvasive techniques like EEG, has focused on improving or extending normal human performance using BMI. While the latter is more in line with the ultimate neuroergonomics goal, the former is important for understanding some of the basic principles involved and setting a baseline comparison for noninvasive BMI. Any form of BMI will involve at least three parts: a recording device to get input from the brain, a translating device to read the brain activity and transform it into something useful, and then the output device to perform some desired action. Output devices include anything from a cursor on a computer screen to a robot. A typical fourth component is feedb ack to the user (which also distinguishes open loop from closed loop systems), but this is not critical to actual control of a machine and so will only be discussed briefly.

(U) Invasive Brain-Machine Interfaces

(U//~~FOUO~~) Invasive BMI are those that involve placing some manner of recording device, usually an electrode array for measuring electrical activity, into the body. Given the goal of using brain signals, this means opening the skull and placing the electrodes on or in the brain tissue. As such, research on invasive BMI has been limited to animals and small trials on humans with disabilities for whom drastic surgery is reasonable. Beyond the surgery itself, invasive BMI carries a number of risks. All current systems use a socket in the skull to pass wires from the electrodes to an external device for analysis, meaning infection is a constant concern. The electrodes themselves have not been reported to last more than around a year, meaning that additional surgery is needed to repair or replace the electrodes. Perhaps related to the previous issue, there is a chance of scarring; while biomaterials are advancing to create increasingly innocuous electrodes, there is some chance of physical damage to brain tissue from the electrodes themselves or from an immune response. And while animals have had electrodes implanted for recording for a few years, the long term effects of an implant in the brain is unknown. Finally, electrodes can only record from the area they are implanted, meaning that invasive BMI can only be used for certain applications related to their location (for example, motor control from motor cortex, visual analysis from occipital regions).

(U//~~FOUO~~) While all of these issues make invasive BMI sound completely unattractive, and the long term goal is to have more versatile noninvasive BMI, existing research is promising and provides information on where BMI technology is now and could be in the future. In attempting to decode neural firing from motor cortex it has become apparent that single neurons are not in control of any single movement. Any one neuron has a preferred direction of motion that it will fire maximally for, but it also responds to movements in a different direction. It is actually the overall firing level in a group of neurons that determines a movement, and in fact implanted BMI record from tens to hundreds of neurons.[91] Perhaps the most important implication of this finding is that multiple types and directions of movement can be read out from the same region (although perhaps not optimally if the movements are very different). Finally, invasive BMI have provided one very important proof of concept: that brain activity can be used to control movements without involving any muscle activity. Because of their use in paralyzed patients and their limited spatial area of recording, it can be concluded that arm movements or muscle activity from the head are not needed or contributing to control of the external device. As we will see in the noninvasive BMI section, this is an important issue for deciding what exactly is being recorded and what could be used as a control signal. It is also the case that normal subjects, and animals, learn to stop moving their body when controlling an external device because it is unnecessary. This raises the interesting question of if it would be possible to control an external device as well as one's own body simultaneously.

(U//~~FOUO~~) The remainder of this section will focus on two different invasive BMI experiments. The first taught a monkey to control a robotic arm.[92] The task goal in this case was to reach out to one of a few locations, use a pincher-like grabber to retrieve a piece of food, and then bring it back and eat it. The experiment was originally performed with a different monkey who did not have control of the grabber, and then extended to a second monkey who provides the current data. The monkeys had a few days of initial training with the arm via joystick control so that they could learn about the apparatus; the authors presume that this would be unnecessary for a human. Transitioning from the joystick to cortical control was too difficult, however; it is unknown if this would be true for a human. Thus a training program was created wherein the monkey would initially watch the arm move itself (controlled by the experimenter) for a few trials and then control would gradually be given to the monkey by combining its input with experimenter control (called the assisted task; this also served as the calibration task at the beginning of each testing day to create the algorithm parameters and lasted for 16-48 trials).

(U//~~FOUO~~) The first monkey learned first to move the arm in one dimension, straight out from its mouth and back, then in the vertical and horizontal plans, and finally in three dimensions. These three training periods took a month. Once the monkey attained full autonomous control, there was an additional month of training wherein the monkey learned to do the full movement including having food placed in the grabber and placing it in the mouth; finer grabber placement (for example, target accuracy) was also trained. The experimental monkey additionally learned to control the grabber and was given control over opening and closing it. This monkey's training also lasted approximately two months, although it learned what the first monkey did in approximately a month (that is, most of the training was focused on integrating grabber control). The delay in the system was approximately 150 milliseconds, which is comparable to that of natural body command.

(U//~~FOUO~~) The experimental monkey, across two days of testing, achieved a 61 percent success rate, where success is defined as reaching out, grabbing the food, and successfully placing it in the mouth. It is important to note that "chance" success on this task is probably close to 0, and certainly not 50 percent, so 61 percent is fairly high. The movement was broken up into component parts to better analyze the data. The monkey was almost always successful at reaching to the proper target (98 percent) and maintained high accuracy with each successive component task (~90 percent), but failures at any given point led to the overall accuracy rate. The most difficult part of the task actually seemed to be getting the food off the grabber and into the mouth (~83 percent); this could be due to the food chosen or the material of the grabber (the overall accuracy prior to this stage was 76 percent). In any event, this was a large jump in success over the lab's previous work, in which a monkey was only at 90 percent accuracy at the first portion of the movement.

(U//~~FOUO~~) The second experiment (associated with BrainGate) taught a human tetraplegic patient to control a computer cursor.[93] After training on the cursor, the patient was also able to manipulate some simplified computer programs, his television, and very simple movements of a robotic arm (opening and closing a hand, and moving a one-dimensional arm). Similar to the monkey study above, the system was calibrated each day before testing. The calibration involved watching a technician move a cursor on a computer screen to random targets for four minutes to generate an initial filter algorithm (actually a linear regression of neural firing onto target cursor position), followed by four minutes of the patient tracking the technician's cursor with his own. The subject was told to imagine moving an actual mouse or computer cursor on the screen. After the calibration, the subject performed tasks like moving the neural cursor from the center of the screen to one of several peripheral targets, which the subject was able to do in 2.5 seconds with an accuracy that increased across six days from 72 to 90 percent. Normal subjects took around 1 second for the same task. The patient also had some difficulty on a similar task that involved moving to targets without hitting "obstacles" on the screen. Overall, the patient demonstrated that cursor control in two dimensions could be achieved with some accuracy, and that control could be co-opted to allow for control of other external devices like a television or simple mechanical devices. It is likely that better accuracy could be achieved through use of a more complex algorithm; this level of control using a linear regression is fairly impressive.

(U//~~FOUO~~) The general lessons learned from invasive BMI experiments are that neural activity can be used to control external devices in real-time, that they can be used for a number of external devices (perhaps without the need for retraining, although that will depend on compatibility issues), and that even relatively simple algorithms can be used to categorize small populations of neurons' activity by the intended/imagined movement. The issues raised above are still important ones, but it seems that BMI research is well on the way to restoring function to patients with movement disabilities. The next section on noninvasive BMI will look at research into allowing normal subjects to control external devices.

(U) Noninvasive Brain-Machine Interfaces

(U//~~FOUO~~) In contrast to invasive BMI, noninvasive BMI methods measure neural activity from outside the skull, typically using EEG but also occasionally fNIRS. The advantages to noninvasive BMI are largely the converse of invasive; no surgery is required, fixing equipment is relatively easy, there are no long-term health issues, and activity can be recorded from anywhere on the scalp allowing for flexible applications. The major downside to noninvasive BMI is that the techniques currently available suffer from signal to noise issues at the level of single trials and activity can only be recorded from relatively shallow brain regions. In the case of EEG, electrical activity in the brain sums together from a variety of neural sources, which reduces the ability to focus in on a particular brain region. Furthermore, that activity is then "smeared" by the skull and scalp, introducing more noise into the system and further reducing spatial resolution.

(U//~~FOUO~~) fNIRS can read one of a few signals; one uses blood flow and so has better spatial resolution, but can only measure activity in the outer centimeter or so of cortex and lacks temporal resolution due to the sluggishness of the hemodynamic response. A fast NIRS signal can also be measured that corresponds more to neural firing, like EEG, but suffers from noise issues like EEG. Thus the major issue with EEG, which is the primary method for noninvasive BMI, is separating signal from noise. This is accomplished somewhat by using a larger number of electrodes (research apparatuses use as many as 128 channels) but primarily by averaging together a number of trials of the same type to create an ERP. The main hurdle for

EEG then is in trying to minimize contributions of noise (both bodily and external) while finding algorithms that can effectively pull out the signal on a single trial without using too large a time window in order to allow for real-time analysis.

(U//~~FOUO~~) [redacted][94] evaluated fNIRS (the slow hemodynamic measure) in an offline manner to see if it could be used to support a BMI. Sensors were placed over the motor cortex and recorded while subjects tapped either their left or right fingers or imagined tapping their fingers. The data were analyzed offline using both a support vector machine method and a Markov model. Supporting previous EEG data, it was found that actual tapping produced a larger signal than imagined tapping, but the left and right conditions were still separable. The two methods were used to classify left or right hand tapping: the support vector machine had an accuracy of 87.5 percent for actual tapping and 73 percent for imagined, while the Markov model had an accuracy of 93 percent and 89 percent (accuracy is across subjects). While it should be trivial to separate left from right, it is important to note that the subjects received no training or feedback, which would likely raise accuracy. The Markov model was thought to do better because of its probabilistic nature, which would account for variability in tapping initiation and speed, especially in the imagine condition. The authors suggest that great increases in analysis would be necessary to use the fast fNIRS signal, but don't test it directly. It is also unclear how well their method would transition to real-time use, considering the slowness of the hemodynamic response, but they suggest that using a moving window for analysis would be effective.

(b) (6)
OGA

(U//~~FOUO~~) A German group involved in creating an EEG BMI (the Berlin brain-computer interface, or BBCI, project) has published results using machine learning without subject training.[95,96] Their goal is to create a real-time, flexible, fast information transfer system. Like the previous study, subjects in the [redacted] study performed finger tapping, but at three different speeds. The EEG was analyzed to look at the readiness potential (a negative signal that increases in magnitude before an actual movement occurs) and the beta activity over motor cortex (subjects wore a 128 channel EEG cap, however, so many electrodes were available). Using classification based on the period before an actual movement, accuracy decreased with tapping speed (going from 95 percent using the readiness potential at the slowest speed to 73 percent using synchrony at the highest speed) but the information transfer rate still increased with increasing speed.

(U//~~FOUO~~) In a second experiment, the same methods were used to attempt to distinguish between which of two fingers were used, as a proof of concept towards "mental typewriting". Accuracy was not as high as in the previous case, and in half of the cases categorization was at chance even though there were differences in the signal. But, the ability to statistically separate the signals, even if they weren't classified accurately, indicates that classification should be possible.

(U//~~FOUO~~) In a third experiment, subjects underwent a brief calibration test to determine which two movements out of left hand, right hand, and left foot could best be categorized for a subject and then underwent a two-option cursor task. Importantly, one subject could not be calib rated accurately, and so the BMI experiment was impossible. Accuracy in this feedback experiment (subjects saw if they moved the cursor to the correct location) averaged 90 percen t and the information transfer rate was slower than that in the initial tapping experiment.

(b) (6)
OGA

(U//~~FOUO~~) Overall, this study showed that relatively high classification rates are possible using imagined movements even with little training or calibration. However, classification still was not perfect for the relatively easy task of distinguishing left from right hand activity. The more recent paper [redacted] 2008) points out that calibration takes as little as 7 to 20 minutes and that some subjects can operate the cursor without any calibration. However, around a third of subjects (which they claim is a universal number) fail to achieve use of the BMI at all. Their machine learning technique allows for the algorithm to adapt to the subject during calibration, although it still picks out universally relevant features (activity from electrodes near motor cortex, frequencies in the 12-Hz range).

(U//~~FOUO~~) In this experiment, subjects learned to use a hex-o-spell program. An array of six hexagrams was presented arranged around a central arrow; each hexagram contained five letters or characters. Imagined use of the right hand caused the arrow to move clockwise and imagined left hand activity stopped the arrow, causing it to point at the desired hexagram. Imagined foot activity caused the arrow to grow until it reached the hexagram, which was considered a choice. The process was then repeated with the characters in the chosen hexagram such that a single letter could be chosen. The screen returned to the full alphabet hexagram after a choice to repeat the procedure. Two subjects used the program all day at an electronics trade show, which is a fairly ecologically valid atmosphere. Accuracy rates aren't reported, but corrections could be made by choosing a "delete" character, and typing was performed at a speed of around five characters per minute, which is very high for a BMI system. Overall, the BBCI system works fairly well but cannot be used by around a third of subjects and for the most part only makes binary choices. However, the accuracy given the relatively fast calibration time is pretty good.

(U//~~FOUO~~) The final experimental paper involves the creation of a wireless EEG system for monitoring driver attentiveness.[97] The system is very useful from an ecologic point of view: four dry EEG electrodes are placed in a baseball cap to lie on the subject's forehead (along with one by the ear acting as ground). Signals are recorded and amplified at the electrodes and then transmitted to the wireless transmission unit, which is also in the cap. The wireless signal can be either Bluetooth or RF, allowing for different effective distances. The signal is transmitted to a processing unit that analyzes the input and can then send back a signal to the wireless device, which can activate a warning signal if necessary. Subjects performed a calibration session to create a model in which the EEG data is subjected to principal components analysis and then fed into a linear regression; this model was then used for real-time testing.

(U//~~FOUO~~) The task was for subjects to drive down the center of the lane of a driving simulator while the car is randomly drifted away from center. Distance from center was taken as a measure of drowsiness. The system was able to analyze 2 seconds of EEG data and send back a signal in just under 2 seconds, meaning that a driver would be warned 2 seconds after the system determined that he was drowsy in the prior 2 seconds. The system had an average training accuracy of 92.5 percent but only 74.5 percent accuracy during online experimenting. Overall, their system is fairly impressive. The accuracy of the system isn't especially good, but their algorithm and behavioral measure are pretty crude. As an engineering feat, the baseball cap EEG seems effective.

(U//~~FOUO~~) There are also a number of companies involved in creating BMI. It is more difficult to evaluate these systems since they rarely, if ever, have any direct published data on how it works or how well it works. However, some of their claims can be evaluated relative to what is known from the academic literature. Honda's BMI more than likely relies on motor signals to control a robot called Asimo since the user is shown a picture of which hand to move and directs Asimo to do so without overtly moving. The Honda system, however, uses both EEG and NIRS. The system appears to have a large number of electrodes as well as a helmet section containing the NIRS equipment. Honda's BMI does not appear to incorporate machine learning algorithms, instead requiring extensive user training.

(U//~~FOUO~~) Another system, developed by Toyota and Riken, can control a wheelchair via thought. This system appears to have at least three mental commands (forward, left, and right) but it is unclear what signal it uses as an input. It utilizes a method called blind source separation (BSS) to isolate relevant signals from noise and process them in real time. [redacted][98], employees of Riken, used the same BSS method to separate signal from noise in their BMI study. The study required users to focus on checkerboard images flashing at different frequencies in front of, behind, or to either side of a car

(b) (6)
OGA

(b) (6)
OGA

presented on a computer screen in order to direct the vehicle along the road. A calibration run was done in which the user focused on each image while EEG recordings were taken to classify what signals indicated each direction. In this study the BSS method had a roughly four second delay, however, the Toyota wheelchair reports a delay of only milliseconds. This could be because it does not use the same signal as the [redacted] study. A recent update showed that recording from the occipital lobes was more effective and faster than recording from the frontal lobes, and accuracy could reach as high as 96 percent.

(U//~~FOUO~~) EmSense (emsense.com) uses a slim headband to record EEG and is specifically aimed at marketing applications. It claims to measure brainwaves, breathing, eye-blinks, head movement, body temperature, and changes in heart rate. This is likely true, as everything on that list after brainwaves is typically removed from the EEG signal as noise. EmSense claims that these measures can be used to create measures of levels of arousal, engagement, and positive or negative emotion. These measures are combined with a separate eye tracking system in some cases to analyze the response to movies, commercials, packaging, video games, store displays, and so forth.

(U//~~FOUO~~) Given that there are only one or two sensors (the number is unclear), the EmSense system likely works on power analysis of some frequency bands in the ongoing EEG, and perhaps uses the "noise" in the system to measure blink rate, heart rate, and so forth. It seems unlikely that their algorithms for detecting engagement or emotion are better than those developed in the academic community, although they have a much larger body of data than most experiments.

(U//~~FOUO~~) A competitor called NeuroFocus (neurofocus.com) also does marketing research based on EEG. They claim to measure attention, engagement, memory retention, novelty, persuasion, and emotional engagement. A video showing their system says that they use a full-head cap (although perhaps with dry EEG) of 68 or 128 channels along with eye tracking.

(U//~~FOUO~~) A reduced system similar to EmSense's is provided by OCZ (ocztechnology.com) for use in video games. There is a single headband with three electrodes placed over the forehead. Their system reads out responses and maps them onto PC commands such that their system can replace a keyboard and mouse. A description of their system admits that it uses signals from the eyes and muscles as well as the brain, although it also claims that with practice a user can talk while using the system, implying that at least some functions are independent of and resistant to some motor signals. The brain signals it uses are claimed to be the alpha, beta, and gamma frequency waves.

(U//~~FOUO~~) A Swedish company called Interactive Productline has created an EEG-based game called Mindball (mindball.se). Their system also appears to use three forehead sensors, and the game is to move a ball towards a goal (either alone or against an opponent) by relaxing; the game is supposed to teach relaxation through neurofeedback. To that extent, at least, the product is fairly altruistic. It is not perfectly clear, but the system seems to work by registering alpha and theta activity and teaching the subject to modify their levels in the same way that neurofeedback is used to treat ADHD, depression, and some other conditions.

(U//~~FOUO~~) The most stripped-down system is provided by NeuroSky (neurosky.com), which uses a single EEG electrode. This system measures activity in various frequency bands but seems to categorize them into either a relaxed or attending state. NeuroSky is not available directly to consumers but instead is sold to other companies, and the games listed on the website seem to verify this two-option system.

(U//~~FOUO~~) The final company to mention also has the strongest claims. Emotiv (emotiv.com) has a 14-electrode system (and gyroscopes for head movement) that is supposed to measure facial expressions, emotional states, and cognitive thoughts in real time. Like the OCZ system, it can be linked to existing games to bypass a keyboard and mouse, but there are Emotiv games as well. An article about the company states that 30 states can be distinguished by the system, but it is unclear how many of them belong to which of the three categories (facial, emotional, and cognitive). Online demos indicate that the system is very good at recognizing facial expressions, being able to mimic winks, eyebrow movements, smiles and frowns. It also claims that the system can replace a mouse by tracking eye movements and letting some other state determine button clicks.

(U//~~FOUO~~) However, many of the demos show users using obvious motor movements while using the system. An Emotiv-controlled wheelchair works by smiling or winking with one eye or the other to move forward and turn. Users playing a video game move their arm to push an animated block or lift a stone. It is unclear how many of the supposed 30 states are not related to facial movement and how many are truly "cognitive" as opposed to the "emotional" relaxed/excited states measured by simpler systems like NeuroSky. Given that the device operates in real-time, it likely uses frequency bands like the other systems described. One nice feature of the system, if true, would be its ability to measure eye movements. This would eliminate the need to have an additional eye tracker for any related applications.

(U//~~FOUO~~) In sum, it appears that while EEG (and fNIRS) systems can be used for real-time BMI, it is uncertain how advanced they are. The academic studies discussed seem to max out at around 90-95 percent correct categorization of no more than four, and typically two, responses. These systems have been based on activity from motor and visual cortex, which are the two most mapped and stereotyped systems in the brain. Both systems, for example, are mapped onto space in a very straightforward fashion; the left motor cortex controls the right side of the body and vice versa, and the same is true for visual cortex mapping onto the visual field (with an additional reversal of top and bottom). So while it seems like it should be straightforward to determine if someone is actually moving their right hand, let alone imagining moving it, this has been a challenge so far. The claims made by companies should thus be taken with a grain of salt. They appear to mostly rely on analysis of frequency bands, thus achieving a kind of neurofeedback system (this is exactly the goal of the Mindball game). Power analysis of EEG data is not difficult; instead the issue lies in determining if the system is measuring what it claims to measure.

(U//~~FOUO~~) The previous sections for this report have laid out the problems that other researchers have had in using frequency bands to measure attention and other cognitive processes and the problems that some users have in actually using the feedback effectively. A simple system like NeuroSky or Mindball's seem plausible, in which change is effected gradually towards one of two options. The stronger claims made by marketing companies seem somewhat implausible, but marketing research typically approaches more of an experimental status and thus may have some truth to it. The systems that claim larger numbers of options, like OCZ and Emotiv, seem the least plausible. OCZ clearly states that they use eye and motor movements in their algorithms, which make their keyboard mapping claim more likely, although the number of functions that the system is supposed to allow is unclear. Emotiv is the most mysterious of the companies; the facial recognition seems very good and the demos are impressive but seem to rely on muscle movement as opposed to brain signals per se. It is important to note that depending on the requirements, this could be perfectly acceptable. If, however, a BMI needs to determine actual cognitive states in order to categorize some signal (distinguishing simple boredom from sleepiness, for example, or selecting one choice out of 10 when eye and muscle movements cannot be made), it is unclear that any current system is near that capability.

(U//~~FOUO~~) The issue of training is also an important one; while some systems claim to work out of the box, available research indicates that training always increases the effectiveness of the system. Optimal performance will never occur unless the individual differences in user brain morphology and strategy are accounted for.

(U) Summary

(U//~~FOUO~~) In summary, BMI research has made great improvements even in the last 10 years. However, performance does not seem to be quite at the level needed to allow for use in the field: even simple binary discriminations are not perfectly accurate, and performance on an objectively complex but relatively simple movement like grabbing a piece of food and placing it in the mouth was only accomplished 60 percent of the time. As mentioned in the introduction, there are four components to a BMI, and each will be described below with a discussion of what would improve the efficacy of a future BMI system.

(U) Signal measurement

(U//~~FOUO~~) Getting input from the brain is the first step in a BMI system. The main issue in this stage is the signal to noise ratio; invasive systems have an edge here because they are closer to the neural generators of interest and thus reduce noise from other brain regions as well as external sources like the scalp or the rest of the body. However, invasive systems are undesirable in general. Noninvasive systems are preferable but need to address the issue of noise. Portable systems have even further to go, as they must forego extensive preparation (like scalp scraping) and EEG gel which are used in research contexts to reduce noise.

(U//~~FOUO~~) Reducing noise is critical because the signal must be clear enough that different states generate signals distinctive enough from each other to allow categorization. However, signals based on the power of different frequency bands is probably somewhat resistant to noise, since heart rate and electrical noise happen at known frequencies and can be removed and other sources like eye blinks do not occur with any specific frequency. Combining fNIRS with EEG may be helpful as a way to help clear out signals due to eye or muscle movements; the fast fNIRS signal could be used as an additional source of signal to complement or boost that measured by EEG, or the slow signal could be used to constrain the activity.

(b) (6)
OGA

(U//~~FOUO~~) Additional sensors will also improve power, although this is countered by the desire to have less obtrusive systems. Ecological systems, like the baseball cap developed by [redacted] are good but must include some method for keeping the electrodes from moving too much; a relatively tight headband like that used by some of the companies discussed above is probably a good idea. Creating a new method for recording brain activity in a portable, noninvasive manner is unlikely, so improvements in this area will come from creating more sensitive EEG electrodes; perhaps they could emit a small amount of conductive material to boost local accuracy without covering the head in gel. Another alternative would be to create electrodes that extend into the scalp without causing pain; however, it is unlikely that this would reduce the noise sufficiently to be worth it, and it is also possible that the electrodes would become more susceptible to scalp muscle activity.

(U//~~FOUO~~) A different issue is that of what signal to measure. Many of the systems described work (or seem to work) on facial expressions, muscle movements or imagined muscle movements, and perhaps eye gaze. These options are not available in some situations, however. While subjects learn to imagine muscle movements without actually making them, it is unclear if one can imagine some movement while actually performing another one. Eye movements are also unhelpful, since anyone in a dual-task situation is almost certainly using their eyes for something else. Facial expressions may also be undesirable in some situations: some are involuntary and could lead to unintended actions via BMI, and others like blinks could be disruptive toward other tasks. The conditions under which the BMI will be used are thus paramount, as a non-disruptive signal must be found that the user can generate on command while performing some other non-BMI task.

(b) (7) (E)
OGA

(U) [redacted]

(U//~~FOUO~~) [redacted]

(U//~~FOUO~~) The systems created by various businesses, like NeuroSky, Emotiv, and EmSense, are portable, wireless, and real-time. Without knowing what kind of processing they do, however, it is hard to tell if they are using algorithms of the necessary accuracy and complexity. To the extent that they seem to accurately measure something, they may be of current use. If Emotiv can truly replace a keyboard and mouse (for example) that would almost certainly be of use in some situations.

(U//~~FOUO~~) [redacted]

(U) Output

(U//~~FOUO~~) [redacted]

(U) Feedback

(U//~~FOUO~~) Feedback is what creates a closed loop system; the user can see (or otherwise sense) the outcome of his mental activity and adjust appropriately. For some applications, like computer control, feedback is trivial; the user will receive the same feedback that would occur from non-BMI use. Other applications, however, would benefit from further thought. Control of robotic arms or other limbs would benefit from a sense of balance or strain that cannot be afforded by vision. Control of remote devices, such as in teleoperation, will require transmitting information back to the user. The state of this information will be important. Presenting an image on a screen, for example, removes some visual cues and may reduce performance compared to actually being there. The feedback modality should be chosen with the task in mind; a visual cue is unhelpful if it is supposed to alert a sleepy driver to the fact that his eyes have closed. An auditory cue may not be effective if it is used in a noisy environment, or one that is quiet enough that the sound is startling. The feedback should also not interfere with BMI operation; a visual cue is unhelpful if the system works via eye movements or blinks, and a tactile signal may interfere with a motor controlled system. A human factors approach to feedback will greatly improve the efficacy of any system.

(b) (7) (E)
OGA

(U//~~FOUO~~) [redacted]

(U) Acronyms and Abbreviations

ADHD	attention-deficit hyperactivity disorder
AMPA	alpha-amimon-3-hydroxy-5-methyl- 4-isoxazolepropionic acid
BBB	blood-brain barrier
BMI	brain-machine interfaces
BSS	blind source separation
cAMP	cyclic adenosine monophosphate
CNS	central nervous system
CuO	copper oxide
COTS	commercial-off-the-shelf
CREB	[cAMP] Response element-binding
cRGD	cathelicidin resistance gene
CuZn-Fe_2O_4 copper-zinc	ferrite
DAT	dopamine active transporters
DRD4	dopamine receptor D4
dsRNA	double-stranded RNA
DTI	diffusion tensor imaging
e_3O_4, Fe_2O_3 iron	xides
EEG	electroencephalography
EOG	electrooculogram
EROS	event-related optical signal
ERP	event-related potential
fMRI	functional magnetic resonance imaging
fNIRS	functional near-infrared spectroscopy
GKT	Guilty Knowledge Test
GNP	gold nanoplasmonic particles
GSR	galvanic skin response
LTP	long term potentiation
LVA	layered voice analysis
MEG	magnetoencephalography
MEP	motor evoked potential

mRNA	messenger RNA
MWCNT	multi-walled carbon nanotubes
NIA	Neurotechnology for Imagery Analysts
NIR	near infrared
NMDA	N-methyl-D-aspartate
PbSe	lead selenide
PEG	polyethylene glycol
PET	positron emission tomography
PMC	PubMed Central
RES	reticuloendothelial system
RF radio	frequency
ROC	received operating characteristic
ROS	reactive oxygen species
RSVM	radar vital signas monitor
RSVP	rapid serial visual presentation
rTMS	repetitive transcranial magnetic stimulation
RVIP	rapid visual processing task
SDT	signal detection theory
SETI	search for Extraterrestrial Intelligence
SLN	solid lipid nanoparticle
SQUID	superconducting quantum interference device
SWCNT	single-wall carbon nanotubes
TDCS	transcranial direct current stimulation
TiO_2 titania	
TMS	transcranial magnetic stimulation
VSA	vocal stress analysis
ZnO	zinc oxide

(b)(6)
(b)(7)(E)
OGA

(U) References

1. [redacted]
2. [redacted]
3. [redacted]
4. [redacted]
5. [redacted]
6. [redacted]
7. [redacted]
8. [redacted]
9. [redacted]
10. [redacted]
11. [redacted]
12. [redacted]
13. [redacted]
14. [redacted]
15. [redacted]

(b) (6)
(b) (7) (E)
OGA

[redacted]

(b) (6)
(b) (7) (E)
OGA

30. [redacted]
31. [redacted]
32. [redacted]
33. [redacted]
34. [redacted]
35. [redacted]
36. [redacted]
37. [redacted]
38. [redacted]
39. [redacted]
40. [redacted]
41. [redacted]

(b) (6)
(b) (7) (E)
OGA

42. [redacted]

43. [redacted]

44. [redacted]

45. [redacted]

46. [redacted]

47. [redacted]

48. [redacted]

49. [redacted]

50. [redacted]

51. [redacted]

52. [redacted]

53. [redacted]

54. [redacted]

55. [redacted]

56. [redacted]

(b) (6)
(b) (7) (E)
OGA

57. [redacted]
58. [redacted]
59. [redacted]
60. [redacted]
61. [redacted]
62. [redacted]
63. [redacted]
64. [redacted]
65. [redacted]
66. [redacted]
67. [redacted]
68. [redacted]
69. [redacted]
70. [redacted]

(b)(6)
(b)(7)(E)
OGA

71.

72.

73.

74.

75.

76.

77.

78.

79.

80.

81.

82.

83.

84.

(b)(6)
(b)(7)(E)
OGA

85.

86.

87.

88.

89.

90.

91.

92.

93.

94.

95.

96.

97.

98.

(b) (6)
(b) (7) (E)
OGA

[redacted]