

The Prospects of Whole Brain Emulation within the next Half-Century

Daniel Eth

DETH@STANFORD.EDU

*Department of Materials Science and
Engineering
Stanford University
Stanford, CA, 94305, USA*

Juan-Carlos Foust

FOUSTJC@GMAIL.COM

*Department of Computer Science
Stanford University
Stanford, CA 94305, USA*

Brandon Whale

BWHALE@STANFORD.EDU

*Department of Science, Technology, and
Society
Stanford University
Stanford, CA, 94305, USA*

Editors: Randal Koene, Diana Deca

Abstract

Whole Brain Emulation (WBE), the theoretical technology of modeling a human brain in its entirety on a computer—thoughts, feelings, memories, and skills intact—is a staple of science fiction. Recently, proponents of WBE have suggested that it will be realized in the next few decades. In this paper, we investigate the plausibility of WBE being developed in the next 50 years (by 2063). We identify four essential requisite technologies: scanning the brain, translating the scan into a model, running the model on a computer, and simulating an environment and body. Additionally, we consider the cultural and social effects of WBE. We find the two most uncertain factors for WBE's future to be the development of advanced miniscule probes that can amass neural data *in vivo* and the degree to which the culture surrounding WBE becomes cooperative or competitive. We identify four plausible scenarios from these uncertainties and suggest the most likely scenario to be one in which WBE is realized, and the technology is used for moderately cooperative ends.

Keywords: Brain Activity Mapping, Scalable Neuroscience, Mind Uploading, Technological Forecasts, Substrate Independent Minds, Whole Brain Emulation, 2063



This work is licensed under the Creative Commons Attribution 3.0 License.

1. Background

1.1 Mind Uploading, Whole Brain Emulation, and Substrate-Independent Minds

Mind uploading (or simply *uploading*) is the theoretical process by which a particular mind—the set of thoughts, feelings, memories, and other cognitive faculties intrinsic to an individual—is transferred from an implementation on one substrate to an implementation on another, such as from a biological brain to a computer. (Koene, 2012a) The objective of mind uploading is to create an entity that exhibits all relevant functions of the mind when implemented in a different substrate. This corresponding new implementation of the mind is referred to as a *substrate-independent mind* (SIM). Here, we consider the approach to uploading known as *Whole Brain Emulation* (WBE), by which the neurophysiological and neuroanatomical details of the individual’s entire brain are emulated with such precision that the emulation (the SIM) behaves in a manner typical of the person. While other approaches to uploading, such as by modeling a mind based on external observations, may be possible, these other approaches are not presently well-researched and are outside the scope of this article.

Mind uploading could be performed for a variety of reasons. Productive benefits could be achieved by uploading the brain of a productive individual, especially when one considers the ease of copying a computer program and running it multiple times simultaneously. For instance, copying a SIM based on Elon Musk a thousand times could lead to many worthwhile technological ventures. It has even been proposed that realization of WBE might lead the growth rate of the economy to increase by one or more orders of magnitude, since uploading and copying SIMs would support large, continual growth in the labor force. (Hanson, 2008) Another proposed benefit of mind uploading is legacy, by having a similar, yet distinct entity persist after one’s death. Significantly more profoundly, if uploading the patterns of the brain copies the conscious self in addition to the thought processes of the person, mind uploading would allow for a literal immortality of the sorts promised by certain religions. Commonly accepted perceptions of self allow for breaks in consciousness, such as during sleep, without termination of the self. Therefore, if these commonly accepted perceptions are accurate, continual experience throughout the mind uploading process would not necessarily be required for continued existence from the biological person to the SIM. Another benefit of uploading is once a SIM is created, tinkering with the underlying computer program would allow for influencing the behavior of the SIM. A SIM could therefore see the benefits of cognitive or emotional enhancement (for instance, by increasing the level of simulated dopamine in the brain). Furthermore, once the human brain is understood in sufficient detail, the SIM should be able to “download” specific skills or memories, or fuse with other SIMs to form a hive-mind machine. While in the future philosophical implications of WBE will be more than simply musings, for now WBE remains a fiction. Similar to other technologies, that fact has not stopped the media from discussing WBE, albeit usually in a science-fiction context.

1.2 In the Media

While direct references to the concept of mind uploading have been a part of our pop culture for half a century, the more general idea of giving a mind to a piece of technology has been around much longer. The 1818 novel *Frankenstein* touched on the idea of using science to create a creature with a mind. (Shelley, 1818) Since at least the 1950s, the media has portrayed WBE more directly, as in *The Tunnel under the World*. (Pohl, 1954) Portrayals of mind uploading have often been quite dystopic, as in *The Matrix*. (*The Matrix*, 1999) Mind uploading has been represented in popular culture as a form of cloning, overcoming death, androids taking over the

World, and more recently integrating into another civilization. (*The Sixth Day*, 2000; *Freejack*, 1992; *The Creation of Humanoids*, 1962; *Avatar*, 2009)

Some recent references to WBE have consisted of specific predictions about the realization of the technology. Ray Kurzweil's *The Singularity is Near* predicts that WBE will come to exist in the 2030s. (Kurzweil, 2005) The 2045 Initiative, founded in 2011, has as its expressed mission to enable WBE by creating a "hologram-like avatar" for the mind by the year 2045. (2045 Initiative, 2013) It is likely that as technology improves and we approach WBE, there will be a shift in the portrayal of mind uploading in the media. So far the use of mind uploading in the media has mostly been a metaphor for mankind's relationship with technology or for man seeking the unattainable. If WBE appears viable, however, stories will likely focus more on its significance for the world. Additionally, as mind uploading appears more feasible, its frequency of appearance in the media will likely increase. This has already been the case, as the frequency of the phrase "mind transfer" has risen drastically over the decades. (Figure 1)

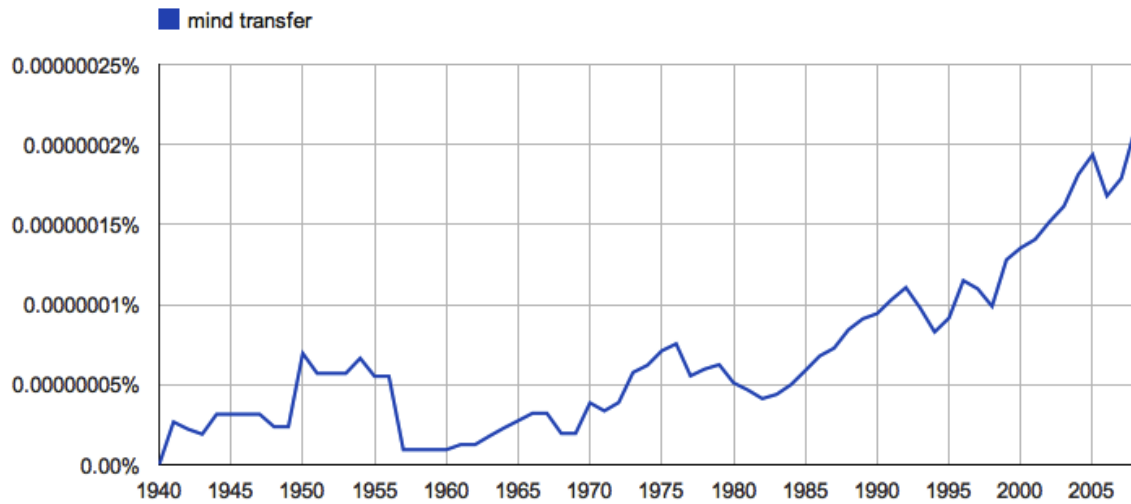


Figure 1: The normalized frequency of the phrase "mind transfer" in books published in American English from 1940 to 2008. (Jean-Baptiste, 2010) The increase in frequency indicates a growing public interest in mind uploading and related concepts since 1940, and we expect this trend to continue as mind uploading becomes closer to reality.

2. Requisite Technologies

Once WBE is developed, the process of mind uploading requires a number of primary steps, which may be completed sequentially or simultaneously. Here, we consider a sequential ordering of the steps for simplicity, as each step can be investigated independently in a sequential procedure. First, the brain will be scanned for necessary information. Then, this information must be translated into a model of the brain (using a more general map of the human brain) that can function as a computer program. After that, this model or program (the basis for the SIM) will be run on a computer, once provided a suitable body and environment. All of these steps are elaborated in the following sections. The uncertain driving technological forces (which will be expanded upon later) are **in bold**.

2.1 Scanning the Brain

Much of the uncertainty regarding the necessary capabilities of brain scanning for WBE is the existence of a *scale separation* between the smaller and larger functional units of the brain. Scale separation—a disconnect between macro and micro properties—enables one to model the “larger parts” of the system with their own unique set of rules without modeling the “smallest parts” of the system. (Sandberg, 2013) One example from elsewhere in science is that between classical mechanics and quantum mechanics. Quantum mechanics describes the behavior of particles according to highly complex rules (such as the Schrödinger equation). When a very large number of these quantum particles act in aggregate, however, they may be described by classical mechanics, which consists of much simpler rules (such as Newton’s Laws). Despite the fact that these macro classical interactions are fundamentally composed of the micro quantum interactions, the classical interactions can be modeled without any knowledge of the quantum interactions.

A similar scale separation potentially exists in the brain. For instance, rules governing neurons or ion channels may be developed which can operate without any knowledge of the atoms in these neurons or ion channels (and thus would require less fine scanning technologies than if knowledge of all atoms were required). A scale separation at the synaptic level or higher likely exists in the brain, for the following reason: the environment in which the brain operates is “noisy,” where external perturbations (such as changes in motion) and internal perturbations (such as changes in temperature) don’t tend to disturb the operation of the brain except in extreme cases. (Sandberg, 2013) Therefore, it is probably the case that the specific interactions on the scale of this noise are nonessential to the behavior of the brain and can be averaged over, since if they were essential, the information they were meant to transmit would be destroyed by the noise, and the brain would not operate properly. The amplitude of this noise is on the level of specific concentrations of metabolites and neurotransmitters, so a scale separation likely exists at a level higher than this. That said, a scale separation probably does not exist at or above the level of spiking neurons, as some neurons communicate without even spiking. (Lytton, 2002) We will therefore use these two levels as the lower and upper bounds for a scale separation in the brain.

In addition to the level of scale separation, the required capabilities of brain scanning will depend on what information is necessary to include in our models. For instance, certain components in the brain are undoubtedly for housekeeping purposes, and it might not be necessary to model them for mind uploading. Depending on the level at which scale separation occurs in the brain, scanning a brain could be achieved through **advanced miniscule probes that amass neural data (AMPAND)**—microscale or nanoscale probes that enter the brain and scalably record data *in vivo* from inside or between neurons in real time—or more **traditional brain scanning** (such as MRI or SEM). While the possibility exists that traditional brain scanning techniques may prove adequate for performing mind uploading once WBE is realized and a general model of the brain exists, they do not appear capable of realizing WBE. This situation is because they will not be proficient in gathering the neuroscientific knowledge necessary for producing this general model of the brain, a point we expand upon in the next section. Therefore, the development of AMPAND appears necessary for WBE.

2.2 Translating into a Model

Once a particular brain is scanned, the next step is to turn this scan into a runnable computer program. This step would use a general model offering a broad framework for the human brain, such that data from a brain scan (addressed in Section 2.1 “Scanning the Brain”) would serve as inputs in generating the parameters to model this particular brain. This general model would not necessitate a full understanding of the reasoning behind the strategies used at each level of

function in the brain, but instead could be based on the functionally relevant behavior of the components used by the brain for computing and the connectivity between these components. (Koene, 2012b) One could imagine this process of translating the brain into a model as: a particular brain is scanned, cell and synaptic attributes are determined from this scan, and then the general model is used to construct a person-specific program with corresponding simulated cells and synapses.

While a scale separation likely exists in the brain at the synaptic level or higher, in creating a general model for the brain, we will likely need resolution somewhat finer than this. To see the mechanisms of operation of the synapses, we will likely need to observe the entities with which synapses perform their roles, including metabolites, neurotransmitters, and ions. Since these substances are the chemicals with which synapses operate, if we can observe them in action, we have hope for determining under what conditions synapses perform their various activities. Observing these specific molecules is likely necessary, since these individual molecular types appear quite important for proper brain function; genetic defects that eliminate even one neurotransmitter (such as dopamine) can have drastic effects on a person's behavior.

Unfortunately, the **traditional brain scanning** technologies (alluded to previously) do not appear capable of accomplishing this task. Traditional brain scanning can largely be broken down into two categories: destructive (such as SEM) and nondestructive (such as MRI). While the destructive scanning technologies have extraordinary spatial resolution, they are clearly not fit for the job, as they destroy the brain in the process of scanning it and therefore can only be used to determine the structure of the brain, not the activity. Activity might be obtainable from structure once we have a general model of the brain, but in order to create this general model, we will need to determine activity to discern how the molecules interact *in vivo*.

The nondestructive brain scanning techniques for scanning the whole brain, such as MRI or MEG (Magnetoencephalography) are not likely to acquire the necessary capabilities within the next 50 years. The spatial resolution of MRI for brain scanning purposes has seen roughly exponential improvement for a number of decades. (Figure 2) Other similar nondestructive techniques, like MEG, have seen a similar trend and currently stand with about the same resolution, so the following line of logic will apply for them, too. First, these technologies will probably run up against some hard limits eventually, as noise will become a bigger issue as resolution improves, and the skull makes it hard for these generally noninvasive techniques to peer deeply into the brain with precision. Therefore, this exponential growth will probably plateau at some point. Second, other specifications besides resolution, such as contrast, will likely limit their abilities. In fulfilling the task of gathering enough neuroscientific data for a general model of the brain, an MRI would not only need to visualize a structure, but also classify it from myriad possibilities. Third, even if these technologies do continue their exponential improvements through 2063, they still won't have the spatial resolution necessary for creating the model by that date. If the trend for MRI continues, the resolution in 2063 will be between 1 and 10 microns. (Kurzweil, 2012) This resolution is an order of magnitude too large to see synapses (presumably necessary to observe for gathering knowledge for a general model), and two orders of magnitude too large to see microtubules and synaptic vesicles (probably necessary to observe for gathering this knowledge).

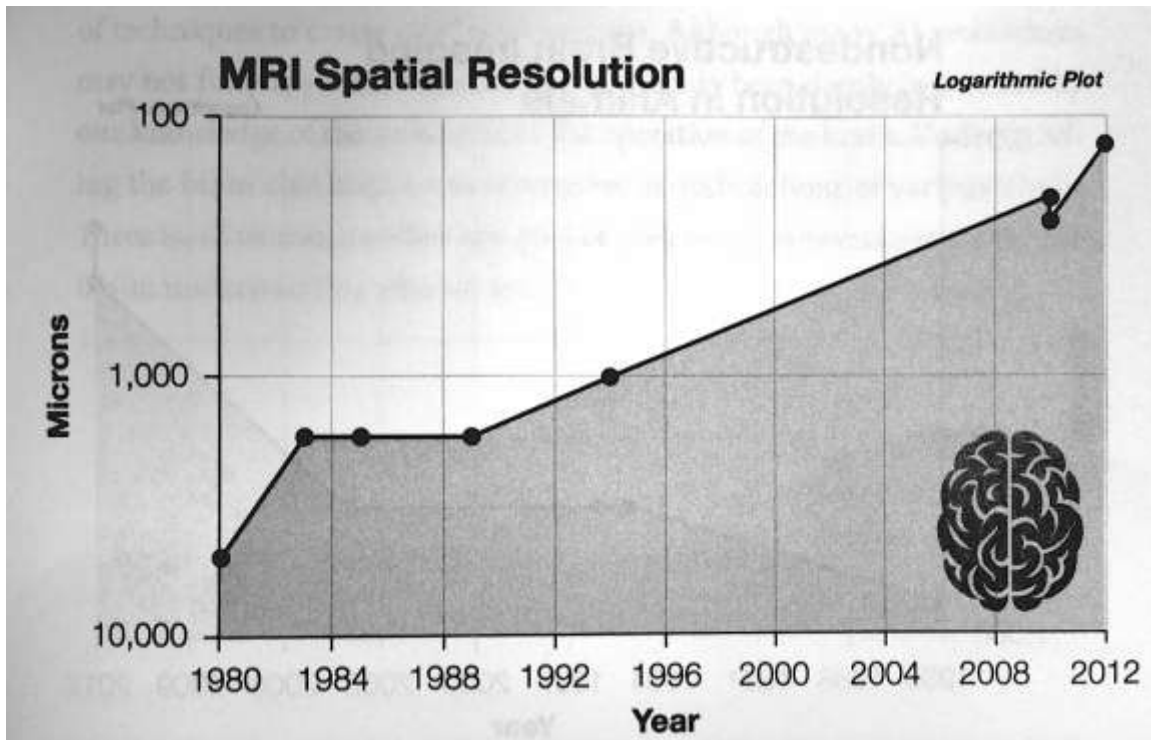


Figure 2: Spatial resolution of Magnetic Resonance Imaging (MRI) for use in brain scanning, in microns. Since the y-axis is logarithmic and the curve is more or less linear, the resolution has improved more or less exponentially. (Kurzweil, 2012) Unfortunately, even if this trend continues for the decades to come (which we find unlikely due to potential hard limits in the resolution of MRI, such as imposed by the skull), MRI still will not have the spatial resolution needed to see pieces of neuroanatomy relevant to WBE (including synapses) by 2063.

On the cellular scale, techniques have been developed to monitor single neurons with great precision. The patch clamp technique, for instance, allows for the study of the electrical properties of ion channels or entire neurons, *in vivo*. Techniques such as the patch clamp have been used to create models of individual neurons and small groups of neurons, but not large circuits of neurons. Ultimately, in creating a general model for the brain, it will be necessary to model the activity of large neuronal circuits, which can contain millions of neurons. (Alivisatos et al., 2012) Current techniques that operate on the cellular scale fall short in that they generally only monitor neurons for at most a few signals, instead of the many types of signals expressed by neurons. In order to simultaneously monitor so many neurons for so many chemicals, **AMPAND** (addressed in Section 2.1 “Scanning the Brain”) must operate in parallel. Furthermore, these AMPAND should be able to stimulate specific neurons in particular ways, so that cause-and-effect can be elucidated. (Alivisatos et al., 2013) Despite the fact that the traditional brain scanning techniques are inadequate in gathering all the necessary neuroscientific data, they still may provide much insight. For instance, a system might be created for using AMPAND in combination with traditional brain scanning techniques to gather the necessary data. One development that has the potential to hasten the accumulation of neuroscientific data is the **automation of neuroscience**. If highly trained humans are needed to perform the physical work of gathering data, the process will likely be slow and error prone. If, on the other hand, robots can automate the process, it may be streamlined.

Regardless of the level of automation, mapping the brain will be expensive, so **neuroscience funding** is also critical. Furthermore, once we have the necessary data, extracting insights will present a major challenge. Determining important pieces, cause and effect, and rules of operation will be necessary. This amount of data will be impossible for a human to sift through. To make sense of this data, **data mining and machine learning** must be used.

2.3 Running on a Computer

Significantly improved computational hardware capabilities are a prerequisite technology for WBE, as the requirements to run a SIM dwarf our current computational resources. Such levels of hardware appear to be attainable if and only if exponential increases in computational abilities continue (which we argue below will be the case). It is most likely that the first SIMs will be run on supercomputers. Assuming continuation of **exponential gains in processing power** and **exponential gains in computer memory**, we can predict when supercomputers will have the capabilities to emulate the human brain, based on certain assumptions about where a scale separation occurs. If all that is needed for mind uploading is a model based on rules of spiking neurons and synaptic states without actual ion concentrations (an optimistic situation), then 8 PB (8×10^{15} bytes) and 1 EFLOP/S (10^{18} FLOP/S) are needed, which should be available to a \$1M supercomputer by 2019 for both the memory and the processing power. If instead the SIM needs to be simulated all the way down to neurotransmitter concentrations in the compartments (a pessimistic situation), then 1 EB (10^{18} bytes) and 10 YFLOP/S (10^{25} FLOP/S) are needed, which should be available to a \$1M supercomputer around 2029 for the memory and 2044 for the processing power. (Sandberg and Bostrom, 2008) Special-purpose hardware, or a significantly larger financial investment, could allow for the realization of these requirements even sooner. PCs have consistently been about 20 years behind supercomputers in terms of hardware. While some algorithmic improvements will occur in the intervening years, they will be small compared to the exponential increases in hardware, so WBE might be available to PCs 18 years after supercomputers, or between 2037 and 2062. It is worth noting that if computational power turns out to be the bottleneck for WBE, then it will indeed take around 18 years for WBE to go from supercomputers to PCs. If, on the other hand, some other requirement (such as gathering the neuroscientific data) turns out to be the bottleneck, then we should very quickly go from SIMs on supercomputers to SIMs on PCs, as computing power will continue to improve past the required level before WBE is realized.

2.4 Environment and Body

The main options for environments for a SIM are a **virtual reality environment** and a **robotic avatar**. A virtual reality environment would consist of the SIM acting through a virtual body in a virtual world. Sensory data from the virtual world would then be fed back to the SIM. While such feedback would be difficult with current neuroscientific understanding, the understanding required for WBE (described above) is probably much higher than simply for feedback, so we do not think this will be a bottleneck for WBE.

In addition to virtual reality, a SIM could inhabit a robotic avatar. This arrangement would consist of a robotic body (likely an android but not necessarily) controlled by the SIM. The signals from the SIM to move the limbs would be transmitted to move the corresponding limb of the robotic body. Additionally, the robot would use sensors (such as cameras) to observe its environment. This information would then be fed to the SIM so that it would observe the surroundings of the robotic body as if it were the robot. The program for the SIM could be stored inside the robotic body, or it could be stored elsewhere and control the robotic body remotely.

3. Driving Forces

In addition to the aforementioned driving technological forces (**in bold**), a number of other driving forces, such as social and cultural ones, could shape the possibilities for WBE. We consider both the technological driving forces and a number of the social/ cultural ones below. For each of these identified driving forces, we assign a value of -10 to 10 for both uncertainty and impact, based on our considerations. For uncertainty, -10 indicates we are totally certain about the development or absence of an event, while 10 suggests we are totally uncertain about the development (such that we have no idea if it will or will not happen). For impact, -10 means the force does not impact WBE at all, and 10 means it affects it in every way possible (for instance, if it is totally necessary for WBE).

3.1 Advanced Miniscule Probes that Amass Neural Data

These AMPAND could take one of many forms; they might use low-power CMOS circuitry, or whatever circuitry has replaced CMOS by 2063 (resembling a recent proposal for neural dust by Seo et al., 2013), or they might be based on DNA (resembling a proposal for molecular ticker tapes by Kording, 2011), or they might be unlike anything proposed to date, so long as they can scalably report all functionally relevant data from within the brain. As mentioned in Section 2.2 “Translating into a Model,” whatever form AMPAND assume, their requirements appear to be prerequisites for WBE. No other technology is poised to give us the necessary information in making our models. Scanning the brain in real-time, *in vivo*, has proven quite cumbersome so far, and every current technique for measuring neural activity would require orders of magnitude improvements in key parameters to simultaneously record from all neurons in a human brain with millisecond resolution, just for neuronal spikes. (Marblestone et al., 2013) We also don’t know at what level a scale separation exists in the brain, so we don’t know what capabilities these AMPAND would need. Adding to the uncertainty, we’re not sure what information they would need to gather; many neurotransmitters haven’t even been discovered yet. On the bright side, the US’s recently announced Brain Research through Advancing Innovative Neurotechnologies Initiative (BRAIN Initiative, discussed in Section 3.4 “Neuroscience Funding”), a large-scale, multi-billion dollar effort to map the human brain in real-time, has called for the development of new tools to record and control the activity of most or every neuron in a brain circuit. The project claims that by year 15 of the project, one million neurons could be observed simultaneously. (Alivisatos et al., 2013) The hardest part of this task might not be recording the data, but getting the data out of the brain once it is recorded. (Dean et al., 2013) A number of ideas have been proposed for retrieving this data, including wireless transmission and recording the data in DNA that would be sequenced at a later time. (Alivisatos et al., 2012) Furthermore, market forces for nanomedicine, for example for fighting cancer, will likely lead to advances that progress us along the path to AMPAND. Already, scientists at Harvard have created nanorobots out of DNA that *in vitro* release antibodies in the presence of leukemia, killing these cells, while not releasing the antibodies in the presence of other cells. (Douglas et al., 2012) By combating diseases using nanorobots, we will develop techniques for observing and manipulating biology on the cellular scale, which is just what we need for AMPAND. Overall, we are slightly optimistic that AMPAND will be developed within the next 50 years. That said, we are not confident in our prediction, and AMPAND provide a highly uncertain, very important condition for WBE.

Uncertainty: 7

Impact: 9

3.2 Traditional Brain Scanning

As mentioned in Section 2.2 “Translating into a Model,” these technologies do not offer the potential to gather all necessary data by themselves. They could, however, be moderately useful in piecing together some of the puzzle, and in advancing the field of neuroscience before AMPAND can be developed. They have seen relatively smooth growth for the past few decades (Figure 2), and the same market forces that have been increasing their capabilities (medical, especially imaging for neurological diseases) exist today and are likely to exist for a while. Additionally, as funding for neuroscience will probably increase (see Section “3.4 Neuroscience Funding”), these technologies will face larger forces to improve. Eventually, however, they will probably reach some more fundamental limits regarding noise, at which point they will likely plateau. Overall, these traditional brain scanning methods will continue to have some degree of impact with a considerably more predictable future.

Uncertainty: -4

Impact: -2

3.3 Automation of Neuroscience

Since Data Mining and Machine Learning, in addition to AMPAND, each have their own category in this paper (Sections 3.5 and 3.1) and we didn’t want to “double count” these topics, the uncertainty in this section relates to automation of neuroscience not inherently encapsulated within those other two sections. In this area of automation of neuroscience, much progress has already been made. In 2012, researchers at MIT and Georgia Tech developed a method for automatically finding and recording information from neurons in a live brain. (Trafton, 2012) This method previously required an expert to perform, and the robot currently performs it more accurately and faster than humans. Other techniques, such as automatically performing SEM on a brain and automatically performing patch clamp on neurons *in vivo*, have also recently been developed. (Center for Brain Science, 2010; Kodandaramaiah et al., 2012) This area is another one in which increased funding for neuroscience will be beneficial. Additionally, many mechanical processes are being automated in the work place now. Automated neuroscience will gain from improvements in automation technology developed elsewhere. Improvements in automation of neuroscience seem relatively certain. This automation mostly applies to automation of traditional brain scanning methods. Since these methods aren’t as impactful, neither is this automation.

Uncertainty: -5

Impact: -4

3.4 Neuroscience Funding

Funding for neuroscience must remain high for enough knowledge to be generated. Currently, the National Institute for Neurological Disorders and Stroke spends over \$1.6B per year studying the brain (Figure 3). We believe that this funding will probably increase as the American population ages and cognitive diseases such as Alzheimer’s and Parkinson’s become larger problems relative to other diseases (since the proportion of these dementias to other diseases increases with age). More directly related to WBE, the EU recently announced that it has decided to spend \$1.6B over ten years simulating the human brain through the Human Brain Project. (Dvorsky, 2013) The United States recently announced its own project to study the human brain (the BRAIN Initiative), which has been suggested to cost around \$3B. (Markoff, 2013)

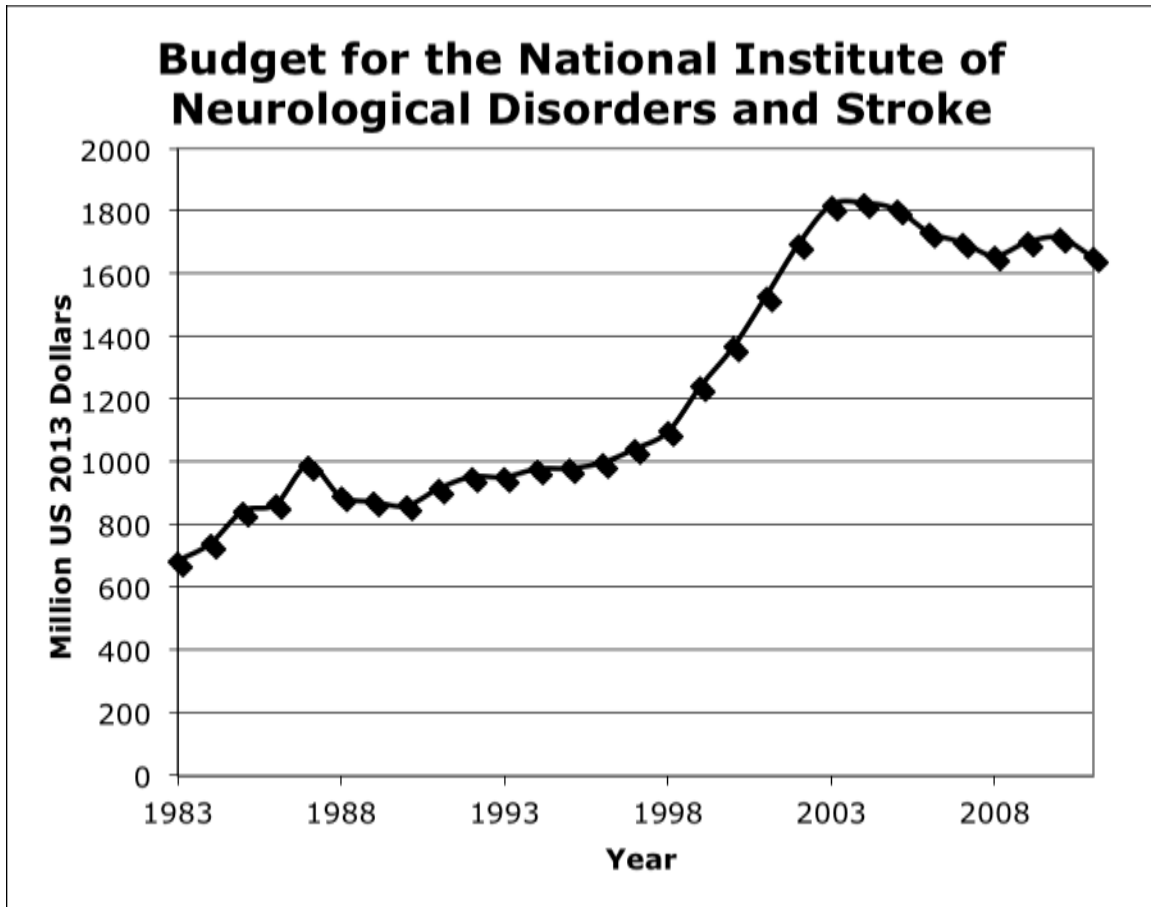


Figure 3: Annual budget of the National Institute of Neurological Disorders and Stroke (NINDS) in constant 2013 dollars. (NIH, 2012a, 2012b; Coinnews Media Group, 2013) The NINDS spends more money annually on neuroscience than any other institute in the world (except for the NIH, which contains the NINDS). We believe this funding to be essential for learning enough about the brain to achieve WBE, and we think that the aging American population will put pressure on the US government to increase this funding in hopes of combating cognitive diseases that often afflict the elderly.

The announcement of these projects might come to resemble JFK's call to put a man on the moon or a similar point in technology development. The Human Genome Project generated \$800 billion for a \$3.8 billion investment. (Tripp and Grueber, 2011) If these projects follow a similar path, future investments in mapping the brain could be even bigger. Countries might realize the economic benefits of these projects, from drug discoveries to AI improvements to WBE. If an "arms race" of brain mapping takes off, funding for neuroscience could skyrocket.

International competitiveness aside, funding for neuroscience outside of America will likely increase, and the impact of this research could also increase. As the developing world continues to "catch up" to the developed world, their total wealth, and especially investments, should increase. Some of these investments will be used towards neuroscience. Additionally, as globalization increases, advances in less developed countries are more likely to be seen by researchers in the developed countries and included in their models.

The level of funding for neuroscience is quite important for WBE, though not as important as some of the stronger technical requirements. There is uncertainty with regard to the funding for large-scale brain mapping projects, but relative certainty for more basic neuroscientific research.

Uncertainty: 2

Impact: 6

3.5 Data Mining and Machine Learning

Computer algorithms for finding the important knowledge within the vast amounts of data are needed for generating a model of the human brain. The machine learning and data mining algorithms we already have developed are seeing much success (from bioinformatics to stock trading) and it is uncertain how much more advanced they must become for WBE purposes. This is one area that is sure to continue to see much market incentive to increase, as better data mining and machine learning promises to be a gold mine (though there is a question of how transferable improvements in algorithms will be to different fields). Additionally, many of these algorithms will function much better with improved hardware capabilities, and so as hardware continues its exponential increases (discussed in Section 3.6 “Exponential Gains in Processing Power” and Section 3.7 “Exponential Gains in Computer Memory”), these algorithms will improve in capabilities, too. Overall, this is an area that is high impact but somewhat likely.

Uncertainty: -3

Impact: 8

3.6 Exponential Gains in Processing Power

A look at the history of computation shows an exponential increase in capabilities holding much further back than the 1965 formulation of Moore’s Law. Calculations per second (measured in either FLOPS or MIPS, shown in FLOPS in the graph below) per constant dollars have seen more than exponential growth for over 110 years. (Figure 4)

Amazingly, not only has there been smooth growth for over 110 years (through war, the Great Depression, the boom and the bust of the internet), but this growth extends through five paradigms of computing: electromechanical, relays, vacuum tubes, transistors, and integrated circuits. This trend of at least exponential growth in processing power will likely continue for decades – long enough for WBE to be computationally feasible. By *computationally feasible*, we mean that a computer would have enough power to simulate interactions throughout an entire human brain down to the neurotransmitter level (though that might not be required) at real-time equivalent speeds, given estimates about how frequently neurotransmitters are transmitted per synapse, estimates for how computationally heavy simulating these interactions with rules will be, and scaling up to the level of the brain. (Sandberg and Bostrom, 2008)

A number of people have speculated that exponential gains in processing power will soon slow considerably, but there are reasons to be confident that these alarms are false. Shrinking of silicon transistor parts will likely reach a limit, but there appear to be many other materials and architectures for computation poised to take its place (and only one needs to succeed), such as: 3D chips, optical computing, carbon nanotube transistors, molecular computing, and DNA computing. (Corrado et al., 2011; PhysOrg, 2012; Anthony, 2012; Carell, 2011) Another oft-raised concern regarding continual increases in hardware is the energy requirements, but the energy efficiency of computers has seen its own exponential increase (with the energy required to perform a set of calculations halving every 1.5 years for 60 years), so this problem also appears surmountable. (Koomey, 2012)

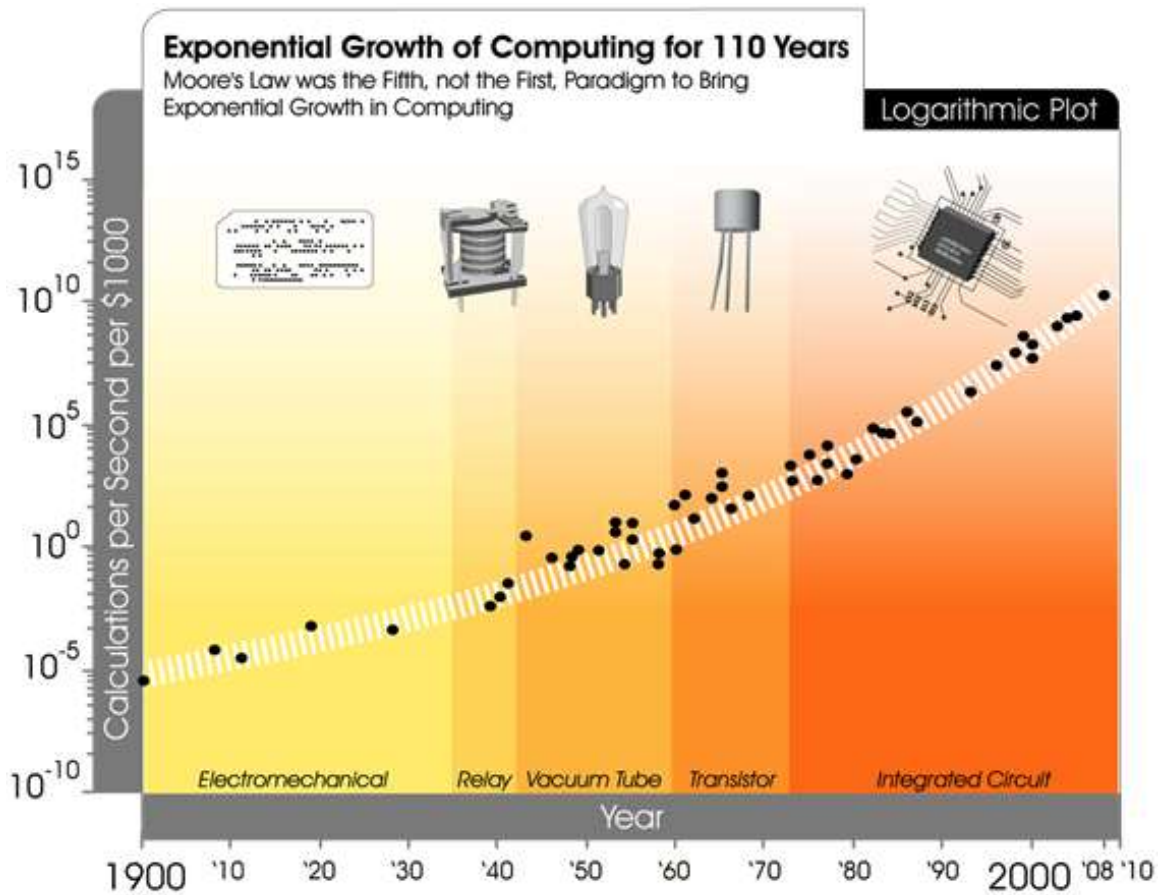


Figure 4: Calculations per second (measured in floating-point operations per second) per constant \$1000 of hardware. Note that the y-axis is logarithmic, so the increasing slope indicates more than exponential growth. (Kurzweil, 2010) This more than exponential growth spans back over 110 years, and we believe it will continue through 2063, by which point the hardware requirements for WBE will be available on computers costing no more than what personal computers cost today.

While most technologies follow a sigmoidal (“S shaped”) curve of growth, where exponential gains are quickly followed by a plateau, economic reasons have allowed computational power to remain in exponential growth. One reason is the larger trend of exponential increases in processing power is made up of smaller sigmoidal curves of narrower technologies. Another reason is that as computational power increases, new applications continuously open up, so there is always a strong drive to improve the technology (for example, iPhones weren’t available 20 years ago). Additionally, we use our most advanced computers to design the next round of computers, so there is a positive feedback loop for the technology. These claims of a continuation of exponential increase in hardware have been echoed by CTO of Intel, Justin Rattner, who thinks there is “no end in sight for Moore’s Law.” (Brodkin, 2009) One piece of good news about this exponential increase is that even if our analysis of the computational

requirements for WBE are wrong (for instance, if our analysis neglects to include an essential piece of the puzzle), the time to reach the actual computational requirements for WBE will not differ significantly from our estimates unless we are drastically wrong. For instance, if our predictions turned out to be based on only one-sixteenth of what is necessary to simulate, then computationally, we are only four doublings (about six years) away in our estimates.

Uncertainty: -8

Impact: 9

3.7 Exponential Gains in Computer Memory

The same logic as above also applies to computer memory (though the trend with memory hasn't been quite as smooth). The major difference is improvements in memory aren't quite as necessary, since we are much closer to the memory requirements for WBE (at whatever scale we need to emulate), and processing power is far more of a bottleneck. If the exponential gains in computer memory stop, the impact to WBE will be quite small.

Uncertainty: -7

Impact: -4

3.8 Virtual Reality Environment

Virtual worlds today are likely good enough for a mediocre environment for SIMs. Second Life might not be as aesthetically pleasing as reality, but it is possible to function in such an environment (consider the fact that there are many people today who spend the majority of their waking hours playing video games, "living" in virtual worlds). There are also huge market forces to improve the graphics of video games, as the video game market is over \$75B annually. (Duke University, 2013) Additionally, the increases in computational power (addressed in Section 3.6 "Exponential Gains in Processing Power") should provide for greatly enhanced graphics. The virtual world requirement for WBE should be available long before our 2063 deadline and should not be the bottleneck, as the simulation of the brain itself should be much more computationally heavy. That said, the bandwidth and complexity of signals required for a virtual reality environment indistinguishable from reality would be immense and should not be understated, requiring high fidelity input through many systems, notably visual, auditory, olfactory, gustatory, and somatosensory (including slight breezes against body hair, temperature variations, feedback from musculature, strain and body stress, sensations of gravity, inertia, and more). The only alternative to a virtual environment would be a robotic avatar (or no body and environment at all), but as discussed below, virtual environments will generally be superior to robotic avatars in many regards. Therefore, virtual reality environments will be quite important.

Uncertainty: -6

Impact: 6

3.9 Robotic Avatar

Robotics is much less advanced than virtual reality, and it remains to be seen if it will be at an adequate level for WBE by 2063. There have been many recent gains in robotics, however, such as a demonstration of bipedal walking. (Boston Dynamics, 2009) At the very least, crude robotic avatars will be available, as one could construct a simple robot today that is controlled by human thought (though not with human-like motion). Sensory input (picked up from sensors on the robot) would be fed to the SIM in a manner similar to how the input for a virtual environment would be fed to the SIM (see Section 3.8 "Virtual Reality Environment"). Robotic avatars are

therefore almost strictly less likely to exist than are virtual reality environments, as creating robotic avatars requires solving the hardest part of creating suitable virtual reality environments, in addition to solving other complicated robotics problems. Even if robotic avatars reach an apt level, they will be inferior to virtual reality environments for a number of reasons. First, virtual reality environments will enable SIMs to perform actions that are impossible in the real world, such as flying. Second, robotics will be much more costly, as a robotic avatar will require power, raw materials, maintenance, and space, while a virtual reality environment will principally require server time. The main benefit of a robotic avatar would be the ability to influence the physical world (for instance, with physical labor). If robotic avatars have not reached a sufficient level by the time WBE emerges (or if they are deemed costly), then there is the possibility for a hybrid technology, in which the robotic avatar is used for real-world actions and sensations, and a virtual reality environment plugs the experiential gaps that arise from lapses in what the robotic avatar offers. While the gains in robotics appear uncertain, they also don't seem as necessary compared to many other driving forces mentioned here, as virtual worlds will exist.

Uncertainty: 7

Impact: -5

3.10 Cultural Movements

Cultural movements surrounding WBE could have a large impact on its use. Transhumanists and hedonists will very likely be in favor of WBE, as they will feel it promises immortality and an enhanced mode of living. Certain religious groups and naturalists will likely see it as a threat to divinity, humanity, or nature. Modern movements against technology usually have been ineffective, and when they have been effective (such as the ban on embryonic stem cell research), researchers usually find workarounds, the research goes elsewhere, and the ban is not permanent. While these bans can slow down research since the overall level of funding for the research can decrease and many of the experts will not follow the research to less resistant locations, outright halts to research have not occurred recently. WBE also will cause people to question the nature of identity and consciousness. Those who determine that identity is not as rigid as they previously thought will have less of a motive to act in their own self-interest as compared to that of the whole. The immortality of SIMs who are backed up will allow them to act much riskier. On the flip side, those who have not yet had their brains uploaded but want to will act with much more caution to avoid an unnecessary early death. Movements for new economic systems will arise as SIMs become cheaper than workers and create unemployment among biological humans. WBE also presents a challenge that will unify people working toward the same goal. The future popularities to the above ideas and movements are unknowable. We believe that the cultural implications of WBE (whether the technology exists or is simply in development) are quite uncertain and will have a large impact on the use of the technology.

Uncertainty: 8

Impact: 8

3.11 Cognitive Enhancement

Human enhancements through means not related to WBE are likely to exist long before WBE. Already many people use weak forms of enhancement, such as caffeine. Further cognitive gains could affect WBE for a number of reasons. Cognitive enhancers could improve the capabilities of scientists working on WBE, thus making them more productive. Additionally, improvements to cognition will lend credibility to the idea that limitations to the human condition are escapable and thus increase WBE's appeal. Separately, other enhancements might be a gateway technology, where people who originally were against WBE for moral or safety reasons might change their view after becoming enhanced. While the future of human enhancements are quite uncertain, their impact on WBE is likely less than the other technological and social forces mentioned here, which more directly relate to the technology of WBE and surrounding issues.

Uncertainty: 9

Impact: -7

3.12 Suspicions

WBE will introduce many suspicions that currently don't exist, and exacerbate some that do. Biological humans and SIMs might not trust each other since the biological humans could worry that the SIMs are smarter and more powerful, and the SIMs might fear lack of compassion from the biological humans. Citizens might not trust governments with WBE technology, since they might fear SIMs will give governments too much power. Similarly, governments might distrust people with the technology, since they will fear the technology will potentially upset their power. Governments likely won't trust each other with the technology, since it could be used as a weapon (although the two biggest projects to map the human brain proposed to date—the Human Brain Project and the BRAIN Initiative—are both being conducted in the open and involving collaborators from multiple countries). People won't trust corporations with the technology, since they will fear unemployment due to replacement from a SIM that can be copied easily. SIMs may be wary of hackers that could read or control their minds. These suspicions could be exacerbated or eased based on laws surrounding uploading and SIMs. If no changes are made in the laws, SIMs would probably not have any rights and would be considered property. Technologies do not fall into place fully-formed, however, and a gradual progression to the existence of SIMs, for instance first through humans augmented with neuroprosthetics, will give society time to debate the ramifications of WBE and enact laws accordingly. While the suspicions surrounding the technology will exist, it is uncertain to what degree, and whether these suspicions will be stronger or weaker than cultural movements for unity.

Uncertainty: 6

Impact: 8

4. Critical Uncertainties

The above driving forces can be plotted for the impact they would have on WBE and their uncertainties, based on the values we assigned. (Figure 5)

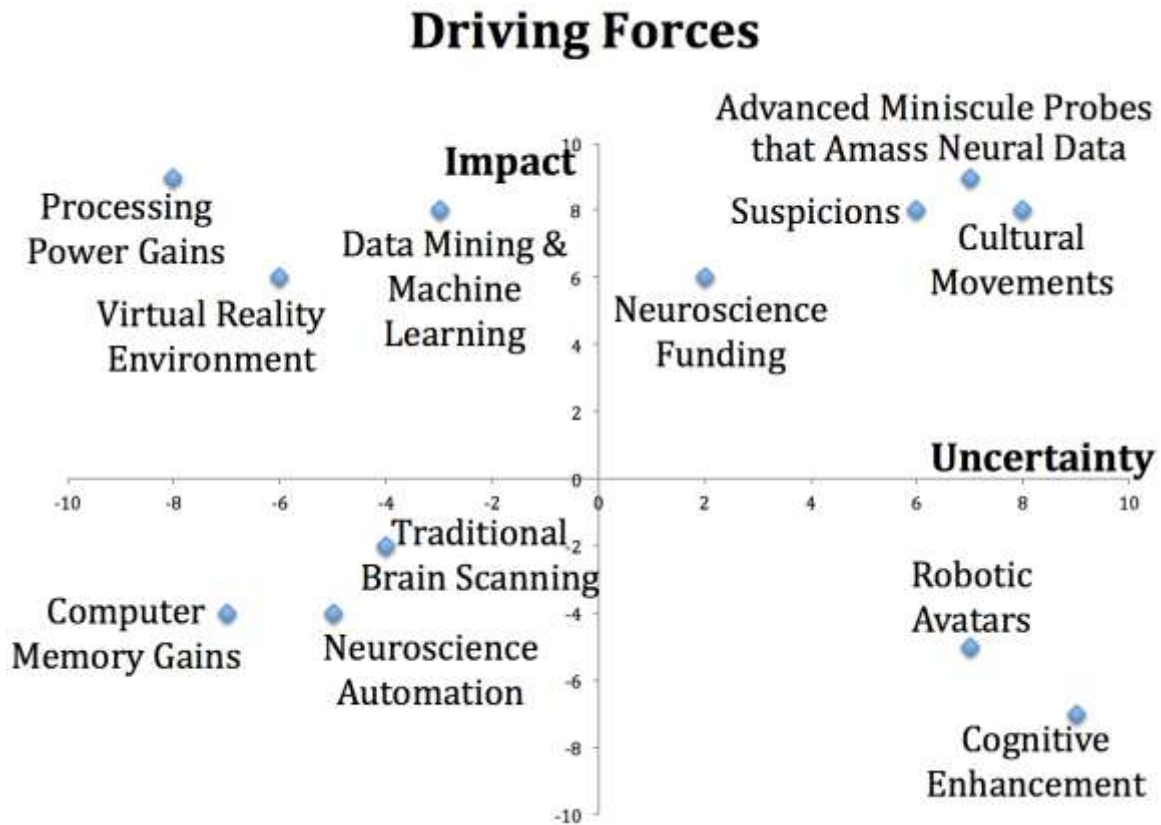


Figure 5: Driving forces by our estimated relative uncertainties and impacts. Advanced Miniscule Probes that Amass Neural Data (AMPAND) are quite important for WBE, and their development by 2063 is very uncertain. Cultural Movements surrounding WBE and Suspensions regarding them will also affect the technology and its use by influencing whether WBE causes more cooperation or more competition, and they too have an uncertain outlook.

Judging by Figure 5, there are three forces that are both very high impact and very high uncertainty: AMPAND, cultural movements, and suspensions. The latter two primarily deal with how competitive versus cooperative the culture of WBE is. Due to the potential power of the technology, in addition to its intrinsic questioning of personal identity, WBE could bring people together or pit them against each other. We believe the highest impact and most uncertain driving forces (our *critical uncertainties*) can be summarized as the existence of AMPAND and how competitive versus cooperative the culture surrounding WBE is.

5. Scenarios

Given the fact that we have two distinct critical uncertainties, we can imagine four scenarios, each consisting of a different combination of polarizing cases of our critical uncertainties. Note that our description of each of the scenarios is not a description of the manner in which the scenario must be, but simply a plausible manner in which it could be, with a bit of creative liberty on our behalf. Our four scenarios (for the year 2063) are shown in Figure 6.



Figure 6: Critical axes and scenarios that we identified, according to our critical uncertainties. If AMPAND (Advanced Miniscule Probes that Amass Neural Data) exist and the culture is cooperative ("Peace of Mind"), then WBE is developed, causing people to believe in less rigid views of individuality and leading to a more unified world in which highly capable SIMs address the grand problems facing humanity. If AMPAND exist but the culture is competitive ("Mind Your Own Business"), then WBE is developed, but tensions between nations and massive unemployment due to cheap labor from SIMs leads to cyber-wars and poverty. If AMPAND do not exist and the culture is cooperative ("Open Minded"), then WBE has not yet been developed, but the prospects are looked at optimistically and unify people. If AMPAND do not exist and the culture is competitive ("Out of Sight, Out of Mind"), then people fear that WBE will be destructive, and nations are competing to develop WBE first for geopolitical advantage.

5.1 Peace of Mind (top right)

This scenario is the most optimistic. Here, technology reaches the requisite level for WBE, and the culture surrounding the technology is cooperative. In this scenario, WBE presents a challenge that unifies scientists across the globe. Following successful development, initial protests and anti-uploading sentiments are quickly dismissed as people see the benefits of WBE. Most upper middle class people have life insurance plans that cover uploading if they near death, and the rapidly falling cost of the technology is leading many to believe that uploading will soon be available to the masses. SIMs can engage in a mind meld in which they share memories, and they can later separate into their individual identities. This experience creates a profoundly unifying experience for those involved. Many people (in particular SIMs) have rejected the apparently true concept of individual identity, and instead taken to believing that all consciousness is fundamentally the same. This feeling has led many towards utilitarianism, and as the SIMs who tend to feel this way have great capabilities due to cognitive enhancements made available to SIMs (such as increased clock-speed), they are swiftly addressing the problems facing humanity.

5.2 Mind Your Own Business (bottom right)

In this scenario, mind uploading is feasible, but there is much social competition surrounding it. The world is centered around a few major nations who have the technology. Unlike the Cold War of the 20th century, however, conflicts here are fought in cyberspace, and these big players don't care about the happenings of smaller countries that don't have capabilities to compete on the same level. Each of these major countries has "armies" of SIMs that work around the clock designing better, more precise weapons, which are autonomously guided. Many of the weapons are computer viruses that disrupt, control, or spy on the SIMs of the enemy side. Other SIMs spend their time improving processing power and computer memory as much as they can to increase the potential size of their armies. Since each of these countries only uses uploading for the smartest, hardest working, most patriotic citizens for their armies of SIMs, the governments have no use for the rest of the citizens, so they are ignored. Massive unemployment due to SIMs taking the majority of jobs has led to anti-uploading movements. Anti-government and anti-uploading protests are routine, yet the governments simply ignore them, not bothering to put an end to these protests as the biological humans pose no threat to their power.

5.3 Open Minded (top left)

In this scenario, WBE has not yet been realized, but there is much cooperation towards advancing neuroscience and creating the technology. An open, international movement to create WBE is underway. WBE is being portrayed often in the media, usually in positive light. Initial large-scale projects to map the brain faced much criticism, but these voices of criticism faded once medical advances became evident. While still somewhat fringe, a fad has spread for people to have their brains cryopreserved upon death with hopes of being uploaded in the future.

5.4 Out of Sight, Out of Mind (bottom left)

In this scenario, WBE has not yet been realized, but there is still a lot of competition surrounding the technology. A number of countries have determined that the first country to develop WBE will be able to upload multiple copies of its smartest citizens and have clear economic and militaristic advantages. This realization has led the militaries of these countries to conduct such research in secret. Different countries have worked out different pieces of the puzzle of creating

AMPAND, but since they don't collaborate, WBE has not been realized. Due to occasional leaks, the public knows governments are developing WBE and is incredibly fearful of it. Leading neuroscientists fear that they may face repercussions from their governments if they don't stay with the research or if they leak information.

6. Wildcards

This section concerns wildcards, events of low/ uncertain probability that, if they occur, will completely change the landscape of WBE. The wildcards we have selected are not meant to be all encompassing. Instead, we chose to include wildcards that we considered some combination of plausible, impactful, and thought provoking.

6.1 Singularity Arrives

If a technological singularity—the arrival of iteratively self-improving artificial general intelligence—occurs, then the effects to technology in general, including WBE, may be profound. Such a smart AI might be able to increase computational resources and create the necessary neural data gathering tools more quickly than in the absence of such AI – if it wants to. The social effects of WBE in this wildcard could also be more jarring, as WBE might appear more rapidly than in other circumstances, and without much warning.

6.2 A Mind is a Terrible Thing to Waste

Before we get to WBE, we will advance neuroscience considerably and develop radically new and highly appealing technologies. Under this wildcard, technology reaches a standstill in the near future as culture and human nature fail to evolve at comparable rates. For instance, people easily “plug-in” their brains to devices that induce a state of pure ecstasy with “no permanent side effects.” While in this state, however, they are totally unproductive. People decide to plug-in to these states very often, and as the drive to progress decreases, technological improvement grinds to a halt before WBE is realized.

6.3 Amazing Brain Freeze

In this report so far, we have assumed that AMPAND are necessary for gathering the knowledge for constructing a model of the human brain. This wildcard considers the situation where that is not the case. Here, the structure of the brain can be determined from the traditional brain scanning technologies, and the rules of operation for the activity of neurons are determined by tests on single neurons, *in vitro*. Therefore, AMPAND are not necessary, and WBE is realized much easier.

6.4 BlueGene Vista

Supercomputers in the coming decades might not have the computational resources necessary for WBE, for one of two reasons. First, if exponential gains in processing power end in the near future, computers won't have the necessary requirements for emulating a brain. Second, if there is no scale separation in the brain (or one only on a very low level), the computational requirements needed for WBE would be too large to emulate the brain. If we want to emulate the

brain by emulating individual molecules, even if exponential gains continue, we will not have the necessary computational power for around 100 years. (Sandberg and Bostrom, 2008)

6.5 Virtual Planet of the Apes

Before we start testing potentially dangerous procedures (such as those involving AMPAND) on humans, we will likely test them on animals. Additionally, since animals have smaller brains that are computationally easier to emulate than ours, we will likely upload a nonhuman primate before a human. Once uploaded, adjusting the characteristics of this primate might prove quite easy. For instance, increasing its intelligence significantly is plausible. In this wildcard, we increase the intelligence of this nonhuman SIM significantly past that of humans. In response, the SIM hatches a plan, which involves tricking us into uploading more of the same type of animal or copying itself multiple times, and then acquiring some outside power, such as through a robotic avatar or through the Internet. Once they have this access, they gain power, until they are the ones controlling world events instead of people. Eventually, they more or less ignore people (the way we do with other animals). WBE would exist, but the super smart primate SIMs see no reason for uploading dumb humans.

7. Forecast

We believe the most likely scenario for the year 2063 is a less extreme version of “Peace of Mind.” In this scenario, WBE exists, and the culture is generally cooperative (though not completely). While the cultural forces for cooperation and competition are both strong, the power resides mostly with those who press for cooperation. Therefore, the development of WBE brings people together in most ways, while increasing tension in others. Most people think WBE is a good thing due to the economic benefits. Naturalists and certain religious people think WBE is an abomination, and while they protest and lobby for bans, they are unsuccessful. The economic benefits of WBE is a tide that lifts most boats, and humanity’s larger problems are being addressed more easily due to the help SIMs provide.

References

- Alivisatos, A. P. *et al.* 2013. The Brain Activity Map. *Science*, 339: 1284 – 1285. doi: 10.1126/science.1236939
- Alivisatos, A. P. *et al.* 2012. The Brain Activity Map Project and the Challenge of Functional Connectomics. *Neuron*, 74: 970 – 974. doi 10.1016/j.neuron.2012.06.006
- Anthony, Sebastian. 2012. *IBM creates 9nm carbon nanotube transistor that outperforms silicon*. ExtremeTech. Available electronically at <http://www.extremetech.com/computing/115657-ibm-creates-9nm-carbon-nanotube-transistor-outperforms-silicon>
- Avatar*. 2009. Cameron, James. 20th Century Fox. Film.
- Boston Dynamics. 2009. *PETMAN Prototype*. Youtube. Available electronically at <http://www.youtube.com/watch?v=67CUudkjEG4>

- Brodkin, Jon. 2009. *Moore's Law has decades left, Intel CTO predicts*. NetworkWorld. Available electronically at <http://www.networkworld.com/news/2009/092209-intel-cto-moores-law.html>
- Carell, Thomas. 2011. Molecular computing: DNA as a logic operator. *Nature*, 469: 45 – 46. doi:10.1038/469045a
- Center for Brain Science. 2010. *ATLUM*. Harvard University. Available electronically at <http://cbs.fas.harvard.edu/science/connectome-project/atlum>
- Coinnews Media Group. 2013. US Inflation Calculator. *Coinnews Media Group, LLC*. Available electronically at <http://www.usinflationcalculator.com>
- Corrado, M., Simon S., and Harris C. 2011. *3M and IBM to Develop New Types of Adhesives to Create 3D Semiconductors*. IBM News Releases. Available electronically at <http://www-03.ibm.com/press/us/en/pressrelease/35358.wss>
- The Creation of Humanoids*. 1962. Barry, Wesley. Emerson Film Enterprises. Film.
- Dean, Thomas *et al.* 2013. On the Technology Prospects and Investment Opportunities for Scalable Neuroscience. *ArXiv:1307.7302*
- Douglas, S. M., Bachelet, I., Church, G. M. 2012. A Logic-Gated Nanorobot for Targeted Transport of Molecular Payloads. *Science*, 335: 831 – 834. doi: 10.1126/science.1214081
- Duke University. 2013. *The Video Game Industry*. Duke University. Available electronically at <http://sites.duke.edu/soc142-videogames/international-trade-patterns/main-players/>
- Dvorsky, G. 2013. *New \$1.6 billion supercomputer project will attempt to simulate the human brain*. io9 Science. Available electronically at <http://io9.com/5980117/new-16-billion-supercomputer-project-will-attempt-to-simulate-the-human-brain>
- Freejack*. 1992. Murphy, Geoff. Warner Bros. Film.
- Hanson, Robin. 2008. *Economics of the Singularity*. IEEE Spectrum. Available electronically at <http://spectrum.ieee.org/robotics/robotics-software/economics-of-the-singularity>
- Jean-Baptiste Michel*, Yuan Kui Shen, Aviva Presser Aiden, Adrian Veres, Matthew K. Gray, William Brockman, The Google Books Team, Joseph P. Pickett, Dale Hoiberg, Dan Clancy, Peter Norvig, Jon Orwant, Steven Pinker, Martin A. Nowak, and Erez Lieberman Aiden*. *Quantitative Analysis of Culture Using Millions of Digitized Books*. **Science** (Published online ahead of print: 12/16/2010)
- Kodandaramaiah, S. B. *et al.* 2012. Automated whole-cell patch-clamp electrophysiology of neurons *in vivo*. *Nature Methods* 9: 585 – 587. doi: doi:10.1038/nmeth.1993

- Koene, Randal A. 2012a. Fundamentals of Whole Brain Emulation: State, Transition and Update Representations. *International Journal of Machine Consciousness* 4: 5 – 21. doi: 10.1142/S179384301240001X
- Koene, Randal A. 2012b. Experimental Research in Whole Brain Emulation: the need for Innovative *in-vivo* Measurement Techniques. *International Journal of Machine Consciousness* 4: 1250004. doi: 10.1142/S1793843012500047
- Koomey, Jonathan. 2012. *The Computing Trend that Will Change Everything*. MIT Technology Review. Available electronically at <http://www.technologyreview.com/news/427444/the-computing-trend-that-will-change-everything/>
- Kording, Konrad P. 2011. Of Toasters and Molecular Ticker Tapes. *PLOS Computational Biology*, 7: e1002291. doi: 10.1371/journal.pcbi.1002291
- Kurzweil, Ray. 2010. How My Predictions Are Faring. *KurzweilAI*. Available electronically at <http://www.kurzweilai.net/predictions/download.php>
- Kurzweil, Ray. 2012. *How to Create a Mind: the Secret of Human Thought Revealed*. New York: Penguin Group.
- Kurzweil, Ray. 2005. *The Singularity is Near*. New York: Penguin Group.
- Lytton, William W. 2002. *From Computer to Brain: Foundations of Computational Neuroscience*. New York: Springer.
- Marblestone, Adam H. *et al.* 2013. Physical Principles for Scalable Neural Recording. *ArXiv:1306.5709*
- Markoff, J. 2013. *Obama Seeking to Boost Study of Human Brain*. The New York Times: 17 Feb 2013.
- The Matrix*. 1999. Wachowski, L. and Wachowski, A. Warner Bros. Pictures. Film.
- NIH. 2012a. Actual Obligations by IC, FY 1983 – FY 1999. *National Institute of Health*. Available electronically at http://officeofbudget.od.nih.gov/pdfs/spending_history/Actual%20Obligations%20by%20IC,%20FY%201983%20-%20FY%201999.pdf
- NIH. 2012b. Actual Obligations by IC, FY 2000 – FY 2011. *National Institute of Health*. Available electronically at http://officeofbudget.od.nih.gov/pdfs/spending_history/Actual%20Obligations%20By%20IC%20FY%202000%20-%20FY%202012.pdf
- PhysOrg. 2012. *Researchers build first physical ‘metatronic’ circuit*. *PhysOrg*. Available electronically at <http://phys.org/news/2012-02-physical-metatronic-circuit.html>

Pohl, Frederik. 1954. *The Tunnel under the World*. Galaxy Magazine.

Sandberg, Anders. 2013. Feasibility of whole brain emulation. In *Theory and Philosophy of Artificial Intelligence* ed. Vincent C. Müller. SAPERE; Berlin: Springer. Available electronically at http://shanghailectures.org/sites/default/files/uploads/2013_Sandberg_Brain-Simulation_34.pdf

Sandberg, A. and Bostrom, M. 2008. Whole Brain Emulation: A Roadmap, Technical Report #2008-3, Future of Humanity Institute, Oxford University. Available electronically at <http://www.fhi.ox.ac.uk/brain-emulation-roadmap-report.pdf>

Seo, D., Carmena, J. M., Rabaey, J. M., Alon, E., Maharbiz, M. M. 2013. Neural Dust: An Ultrasonic, Low Power Solution for Chronic Brain-Machine Interfaces. *ArXiv:1307.2196*

Shelley, M. 1818. *Frankenstein*. London: Lackington, Hughes, Harding, Mavor & Jones.

The Sixth Day. 2000. Spottiswoode, Roger. Columbia Pictures. Film.

Trafton, A. 2012. *Robots that reveal the Inner Workings of Brain Cells*. MIT News. Available electronically at <http://web.mit.edu/newsoffice/2012/robots-recording-neurons-0507.html>

Tripp, S. and Grueber, M. 2011. *Economic Impact of the Human Genome Project*. Battelle Memorial Institute. Available electronically at http://battelle.org/docs/default-document-library/economic_impact_of_the_human_genome_project.pdf

2045 Initiative. 2013. *2045 Initiative*. Available electronically at <http://2045.com>