Managing BPM Toward the Singularity

Roy Altman, Memorial Sloan Kettering Cancer Center, USA

What is "The Singularity?"

"The Singularity" describes when ordinary computers exceed the capacity of the human brain. "Moore's Law," which states that computing power increases at an exponential rate, has held constant since the dawn of the computer age. If one extrapolates forward, experts agree that the Singularity will be reached between 2030 and 2045. Once The Singularity is reached and exceeded, we will find answers to age-old debates, such as whether machines can have consciousness.

Premise of This Paper

Between the present and when The Singularity is reached, we will have increasingly powerful technologies available to us. We need to be able to manage these technologies effectively.

We are entering "the age leading up to the Age of Intelligent Machines." Big Data Analytics and Artificial Intelligence are already having a profound impact on business and society.

One can safely predict that the current trends will continue, and computers will get smarter and smaller. Digital devices will permeate everything we do and assist in every decision we make. The delineation between our digital and physical lives will continue to blur, as wearable devices will enable ubiquitous connection to the web. The nature of work is already changing, and technology will support a blurring of the lines between our work and private lives.

BPM's role has been to provide structure: initially for predictable processes and later for uncertain ones. The Internet of Everything, in which all devices and services are connected, still requires structure if it's to be utilized to its greatest extent. Given that some of the "things" will be advanced Big Data Analytics and Intelligent Agents, means that more than ever these resources need to be managed to achieve our goals. BPM will be more essential than ever to orchestrate resources while complexity is rapidly increasing. Essentially, the role of BPM will be to manage this complexity.

Traditionally, we manage the technology we have at hand. Most companies are late adopters of newer technology, preferring to mitigate the risk of early adoption before the market for that technology matures. Forward-thinking management is looking around the corner to plan to leverage leading-edge technology of the next three to five years (while mitigating risk). No doubt we will continue to be surprised and astounded by the technological advances over the next few decades, at a micro level. But since we *know* the general direction we're heading at a macro level, it only makes sense to plan to manage the powerful technologies of the next 5-30 years, with BPM being a central component of that plan.

HOW NEAR IS THE SINGULARITY?

The human brain is the most complex machine we know of. Computing power is usually measured in millions of instructions per second (MIPS). This measurement doesn't really relate to the human brain, because a computer does one thing at a time very fast and a brain does trillions of things at a time very slowly. This is because most computers are designed in what's called the Von Neumann architecture[1], named for the computing pioneer John Von Neumann. The Von Neumann architecture is designed to do one thing at a time very fast, with the ability to quickly switch between tasks[2] and remember where to proceed from when returning. The brain is a massively parallel architecture, whereby the billions of neurons and synapses can fire simultaneously. Von Neumann-architecture machines are good and deriving a finite answer very quickly. Parallel architectures are good at pattern recognition. Nonetheless, experts conclude that the activity of the neurons and synapse in the brain equate to between 100 million and 2.8 trillion MIPS. A large number, considering today's home computer is roughly 80 MIPS.

Moore's Law has stated that processor speed doubles roughly every 18 months. Computing speed is a function of how small you can build transistors, as their speed increases when they're closer together. This equation has held up throughout computing history, so far. There is a physical limit, however, when transistors are only a few atoms apart. Much of the increase has come from clever engineering, which will present challenges, but not necessarily roadblocks when the physical limit is near. However, we're nowhere near the physical limit. Extending the lines of Moore's Law will get us in the brain's MIPs range between 2030 and 2045, and there's still a lot of range for increasing power past that.

Calculating MIPS on a home computer misses the point, though. Today, computers are massively networked, so the power of a single computer doesn't represent the computing power available. Using the open-source software NEST, the scientists simulated a network consisting of 1.73 billion nerve cells connected by 10.4 trillion synapses. The process took 40 minutes to complete the simulation of 1 second of neuronal network activity in real, biological, time. Although the simulated network is huge, it only represents 1% of the neuronal network in the brain. Using a network of computers, The Singularity may be reached quicker than the estimates above.

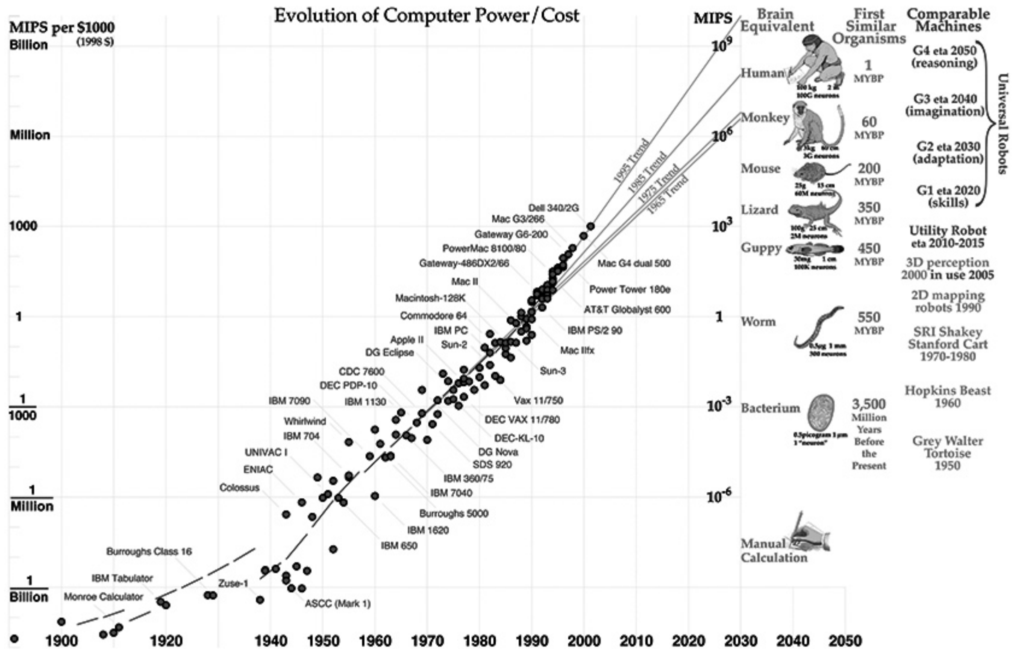

Figure 1: Evolution of Computer Power/Cost.

WHAT WILL HAPPEN AT THE SINGULARITY?

When The Singularity is reached, nothing significantly different than the day before will occur. That is the nature of scientific revolutions; they tend to be evolutionary rather than revolutionary[3]. That's the point of this paper: between now and The Singularity (and beyond), we will have to manage increasingly powerful software. Let's start planning for that now, and seize the opportunity to leverage powerful technologies to solve our goals, rather than let the technology increase unplanned.

Just because we have computers with the technical capabilities of the human brain doesn't mean we can create an artificial mind. We don't understand the software of the brain. However, we can attempt to reverse-engineer the brain by recording the state of every neuron and placing it into computer memory. The Singularity actually represents a confluence of technologies, including non-invasive brain mapping. Today we can do CAT scans which can map large areas of the brain without having to dissect one. This technology will improve to the point where it can map the state of individual neurons. Neural networks have the ability to learn, so the artificial mind will have all of the memories of the source, but their experiences will diverge after the transfer. At that point, we will find out if consciousness can arise from the sum of its parts, or if human consciousness is something different.

A LOOK AT THE BRAIN

The brain is organized very differently than the architecture of conventional computers, which makes modeling the brain on the computer challenging.

- *Brain "circuits" (neurons) are very slow in comparison to digital circuits, but it's massively parallel:* on the order of one hundred trillion interneuronal connections, each with the potential of processing information simultaneously[4].

- *The brain combines analog and digital phenomena.* Although it is analog by nature, a neuron either fires or it doesn't. The firing of a neuron strengthens a connection, and the lack of firing weakens it. That's how we learn: by neurons firing to reinforce information that we want to preserve.
- *The brain can rewire itself.* The brain has a property called plasticity, which means it's constantly adapting to environmental changes or damages. The nervous system is self-organizing. Computers don't rewire themselves, but we're beginning to see "self-healing" software that can mimic this phenomenon.
- *The brain can contradict itself.* It is possible to hold conflicting beliefs simultaneously.
- *The brain uses emergent properties.* Intelligent behavior is an emergent property of the brain's highly complex, and seemingly chaotic activity. It will be interesting to see if this emergent property can arise from man-made computers.

We are beginning to see more computers with massively parallel architectures, and software intent on modeling the way neurons behave.

WHERE WE ARE NOW

CURRENT STATE OF OPERATIONAL SYSTEMS

We have had wonderful advances of technology that has made our lives better, but still many companies struggle to run payroll correctly. Operationally, many IT shops are dealing will common issues:

- *Processes:* many processes are broken, not consistently applied, poorly documented and understood, or just not followed. The Sierra-Cedar HR Systems Survey[5] has reported that the leading area that clients intend to spend time on is Business Process Improvement. This has topped the list the last three years.
- *Data:* access to data and reporting is often at a low state or maturity. Most companies can only provide basic reporting, not data visualizations, self-service or predictive analytics.
- *Integrations:* for the most part, "integrations" between systems are typically batch, flat-file interfaces; despite that web services are mature technologies. One reason that vendors often don't support modern integrations is that customers don't request them.
- *Silo'd Organizations:* a single, hierarchical management structure only reinforces the silo'd nature of organizations; there is duplication of processes for each reporting structure, and no coherent strategy for integrations. Management often doesn't take a holistic view of the highly matrix relationships that determine how work actually gets done in an organization.

Many companies are struggling with the basic challenges and aren't prepared to integrate newer technologies into the mix.

ENTERPRISE SOFTWARE

Enterprise Software is in the throes of transition. Over the last couple of decades, companies would license Enterprise Resource Planning (ERP), Human Capital Management (HCM), Customer Relationship Management (CRM) or Financial software and customize it to adhere to their company's processes. However, the customizations made upgrades to the software more complex, so they occurred infrequently. Also, it was expensive to maintain the on-premise software with dedicated IT teams.

Enter Cloud, or Software-as-a-Service (SaaS). Cloud is a deployment model: it means that neither your data nor the software is stored on-premise, instead managed by the vendor. You don't know where your data is physically stored, nor should you care if you trust the vendor to keep it safe. Software-as-a-Service is a delivery model whereby use of the software is provided as a service, and client has a subscription for use. The architectures to support this means that every client is running the same iteration of the software, and all clients share a single, multi-tenant database. Since all clients are running the same code, it cannot be customized for the needs of a single client. It is therefore configurable, rather than customizable, whereby the configurations for each client are stored in the database, so that the software can react differently based on each client's needs, and still share the same code and database. Modern software has many very flexible configuration points. This means that the vendor can easily make updates to the software, which is then available to all clients. Since updates are quicker and easier, the vendor can make them available more frequently; so the vendor can innovate faster.

In certain industry sectors and cultures, companies have been reluctant to adopt cloud software out of concern that the vendor won't keep their data safe. However, more and more companies are concluding that these fears are unfounded and recognize the advantages to moving administrative applications to the cloud.

MANAGING HUMANS

Processes invariably involve humans, so software to manage those workers is critical to an organization's ability to function effectively. Human Resources (or Human Capital) Management software is a system of record for the workforce, and manages all of the transactions relating to workers (hiring, terminating, promoting, transferring, etc.). A recent category of software called Talent Management has emerged, which is specifically concerned with engaging the workers and optimizing their work. This includes functions such as Talent Acquisition, Performance, Compensation, Learning and Succession Planning.

Managing humans is complex in that they are expensive to recruit and hire, difficult to keep motivated, develop their careers, and retain them in the workforce. However, humans provide special skills that are difficult to automate – although that is changing.

MANAGING PROCESSES – THE STATE OF BPM SOFTWARE

BPM software is general-purpose (meaning not application specific) software to manage the interactions between people and automated processes. It typically includes an integration engine to integrate the software agents and a flexible workflow engine to manage the human interactions in a process.

BPM software has been very successful in automating and more effectively managing complex yet predictable processes.

ENTER ADAPTIVE CASE MANAGEMENT

Adaptive (or Dynamic) Case Management (ACM) recognizes that increasingly workers are knowledge workers – requiring special skills; whose work does not necessarily follow prescribed paths. The bulk of automated support thus far has been geared toward automating administrative tasks, rather than the more unpredictable and complex work knowledge workers engage in.

Adaptive Case Management is well aligned with the trends toward more sophisticated software in that it addresses the needs of humans working the way humans do. Advanced automation often takes the form of an Intelligent Digital Assistants (IDA's), and ACM software is in line with this development in that it helps workers manage unpredictable tasks.

MOBILE AND SOCIAL

When the iPhone was launched in 2007, it ushered in a new age of mobile computing, and mankind never looked back. It has gotten to the point where people are inseparable from their devices. This is in line with Moore's Law stating that computing power will get faster, cheaper and smaller.

Social computing taps into a basic need of humanity to interact with others. It allows people to bridge physical space while at the same time polarizing people from interacting in the real world. Regardless, it's clear that social computing is eroding the division between our physical and digital lives.

CONNECTED DEVICES

The Internet has given rise to a trend whereby devices can be connected and communicate with one another to share information relevant to the goal of each device. This is called "the Internet of things." An example is the Nest thermostat (the company has been acquired by Google), which becomes familiar with the patterns of your life to more effectively and efficiently heat your home.

THE STATE OF ARTIFICIAL INTELLIGENCE

After early attempts fell short of expectation in the 1980's, artificial intelligence (AI) now pervades everyday life. The Siri voice recognition facility in the Apple iPhone usually understands what I'm saying, and often offers a useful response. GPS software quickly computes the best way to get from point A to point B. MS Word indicates usage errors in sentence construction. Google searches relevant results. Software is now being used to write earnings statements in business publications. Facial recognition software allows identification of people. Self-driving cars are well on their way to becoming mainstream.

Early attempts at AI were based on "Expert Systems," whereby the rules for decision-making were embedded in the process. What differentiates humans from machines is that humans are good at pattern recognition, whereas machines are good at fast computations. As stated earlier, the human brain is a massively parallel architecture, whereby design common computers are better geared for fast computation. For instance, it is an extremely complex task for a robot to navigate through a crowded room, such as Grand Central Station in New York, yet any self-respecting cat can do it with aplomb. Current and future AI is better, and will continue to improve, in tasks that require pattern-recognition. Google and others are making great strides with self-driving cars, a task heavily dependent on pattern recognition. Occupations that have traditionally fallen into the realm of knowledge-workers are in danger of being supplanted by machines. Highly compensated fields like evaluating market trends may be among the casualties.

BIG DATA

Big Data, or advanced analytics techniques, have garnered much attention of late. Some definitions of Big Data require that the three "V's" must be satisfied:

- *Volume:* there must be a large volume of data ("large" is a relative term)
- *Velocity:* the data must be changing rapidly
- *Variety:* there must be several data sources

In the year 2000, mankind produced 15 exabytes of data (which is 15 billion gigabytes). In 2012, we produce 2.5 exabytes *per day*. Needless to say, the volume of data is increasing exponentially, which is a trend that is unlikely to reverse. Much of this is considered *unstructured data*, meaning that it isn't generated from systems used by organizations, but rather exists on the web. We can mine that information to produce correlations, or similarities between two occurrences. We can use statistical models to indicate strong correlations which are likely to be causal in nature. Using this, we can extend these models to make predictions about what is likely to occur, based on data analysis of what has occurred, and applied feedback mechanisms. Big Data analytics works by filtering out "noise" and recognizing patterns in vast amounts of information. It has already made an impact in targeted advertising, and is quickly making its way into the information strategies of leading companies.

THE ROLE OF BPM

BPM should play a significant role in managing increasingly sophisticated technologies. The role of BPM is to be a framework within which all other technologies will operate. As technologies become more geared toward the personal intelligent assistant concept, BPM and ACM will merge into a way to manage both structured and unstructured processes. We are entering an era of great complexity – among the available technologies and their impact on our lives. The concept of process *orchestration* will take on a new meaning as ACM/BPM becomes the framework for managing that complexity.

THE FUTURE OF WORK

We are heading in a direction where administrative work will continue to shrink, while knowledge work will be increasingly growing. Companies will "get their act together" and master operational systems with varying degrees of success, thereby achieving short-term strategies. *All* administrative systems will be in the cloud: the commoditization of infrastructure will be complete and there will only be sporadic vestiges of on-premise systems, as the laggards hold out, apart from very application-specific systems developed in-house. As a result, areas that we now see as strategic will become operational and commoditized, and the focus of what is "strategic" will center on effectiveness rather than efficiency. We will be able to do things quickly and accurately, however we still need to decide what we want to accomplish.

BPM will follow suit and traditional BPM/ACM balance will increasingly shift toward the latter. More sophisticated Intelligent Assistants will augment the decisions of the knowledge worker, with the help of Big Data.

Social and collaboration tools will mean that we will increasingly become more connected from the confines of our home offices. Will coffee room gossip survive as it migrates to social media? The 9-5 workday will be a thing of the past. Our work and leisure lives will become intertwined, just as our virtual and physical worlds are merging.

The challenges will be as always, personalities, culture and politics, now with the added task of managing a plethora of increasingly intelligent and complex technologies.

As more types of jobs will be automated by intelligent agents, there will be a greater divide between the knowledge “haves” and “have not’s.” Even if we are able to raise the education level of our population, there exists the danger of pervasive, structural unemployment if we are not able to resolve this issue.

Where are we Going – near term 5-15 years

What I consider to be the mid-term period will see rapid advances in technology. Big Data will grow to the point that every action and every decision we make will involve analytics. Predictive Analytics will become key to our decisions, and we will gain more insight into root causes through better data mining and analysis techniques. As a society we will have to make decisions as to the trade off between personal privacy and connectivity and convenience. Work will become more virtual, as collaboration software will be better able to communicate non-verbal cues missing in emails, phone calls, and to an extent, video. More of our interactions will take place on social media (adding to the data that can be harvested for predictions).

Miniaturization will continue. Devices that we now carry in our pockets will be worn, perhaps on a shirt button. Google Glass, the famously unsuccessful attempt by Google at optic wearable device, was merely ahead of its time by a few years. Devices such as this will be ubiquitous, allowing us to superimpose generated digital information with that of the real world. With devices getting smaller, input methods will have to change. Voice recognition will continue to improve to the point where keyboards become unnecessary and impractical. Voice recognition technology will encompass semantic meaning rather than just syntactic understanding of an utterance, thereby enhancing the relevance of the response. Experts differ, but I think that implants will not catch on, as they are too invasive to install and update, so we will continue to use wearable devices.

Our devices will play the role of Intelligent Digital Assistant, and we won’t leave home (literally or figuratively) without them. I’m unsure as to whether the Turing Test[6], whereby one can’t discern between a digital or human intelligence, will be passed by then, but surely the interactions between human and device will become deeper. Imagine going through our lives with the equivalent of a genie on our shoulder filling in all the blanks of what we don’t know? Still, humans are not just rational decision makers; emotion plays a large part of how we act and what we decide. Our IDA’s will help us sort through the emotional side of our reasoning as well. This is not to say that we will become more rational decision makers, just that we will have the tools to augment what we know and are, for whatever ends we make of it. But since knowledge-workers add the most value by making decisions, and information supporting decision-making is readily available, people will spend a good deal of time making high-level decisions. There should be a great economic benefit from increased, and higher level automation, as productivity per (human) worker should soar.

Managing a Human/Software Workforce

The workforce will consist of a mixture of humans and intelligent agents, each requiring their own considerations. While those most cantankerous of workers (humans), require a great deal of management to keep them engaged, compensated and productive, machines have fewer needs (apart from electricity). Since intelligent agents "learn" as humans do, one area of commonality is training. Robots need to be trained just as humans do. However, the way they learn will undoubtedly be different and will require different training strategies.

Robotics

Cognitive and spacial intelligence are two very distinct things. It is much easier to automate cognitive tasks, and non-moving technology will outstrip mobile robots. There are some applications that require robots that can move and manipulate their environments. There will continue to be advances in this field. Humana-form robots, as we're used to in Science Fiction movies, is probably further away.

- *Self-driving cars:* Google is making great advances with their self-driving car. Over the next 5-15 years, this technology will be commonplace. Human-driven cars will probably not disappear by then, but as people recognize the benefits of no accidents and no traffic jams, driverless cars will be rapidly adopted. Seniors and the disabled will be granted newfound mobility due to this technology.
- *Defense:* Many advances in AI come from Defense Advanced Research Projects Agency (DARPA) research, which has had many civilian applications. Drone airplanes are becoming commonplace already. Specially designed reconnaissance robots will rival cats in their ability to navigate terrain. Robotic soldiers are not long in coming. Currently, humans, who are responsible for making the kill decision, pilot drones remotely. By then, will we be ready as a society to have autonomous robots making kill decisions?
- *Home care:* Seniors and the disabled will be the big winners in this timeframe. Not only will self-driving cars afford a new mobility, but also domestic household robots can cook, clean and assist so as to enhance independence and quality of life.
- *Mining:* In general, robotics will be used for dangerous occupations. Mining, or rescue operations are good applications.
- *Space exploration:* It seems obvious that robots will conduct space exploration. They lack the boredom that human explorers would be subject to in long journeys, and don't require life support. As AI becomes more human-like, it will make perfect sense that future space exploration be unmanned.

The Role of Humans

As machines increase in sophistication, more emphasis will be placed on the qualities that differentiate people from machines: expert thinking, interpersonal skills such as the ability to explain and persuade, creativity and applying common sense. Less emphasis will be placed on technical tasks, as they are easier to automate. In a sense, advanced automation democratizes the role of humans; a person coming from an underprivileged background is just as likely to have interpersonal skills and common sense than their privileged counterparts.

The arts offer an opportunity for differentiation. Artists perceive the world from a different perspective, interpreting the world abstractly. True creativity may be the final frontier of differentiation between humans and intelligent machines. Today, people bemoan that there aren't enough students studying STEM (Science, Technology, Engineering, Math) disciplines to meet the demand (particularly with representation of all demographics of the population). Perhaps the correct acronym should be STEAM (adding Art) for a more accurate depiction of our future needs.

Regardless, we can be assured that the needed skill sets will change over the coming decades, and today's students will be applying for jobs that currently don't exist. Futurist/author/inventor Ray Kurzweil (who now works for Google) predicts that the most in-demand job in the IT field in 20 years will be Personality Designers[7].

WHAT SHOULD WE BE DOING?

THE PLATFORM – PLAYGROUND OF THE FUTURE

One thing we do know about the future is that the basic protocols for machine-to-machine communication will not fundamentally change. While previous attempts at integration standards have been supplanted, XML, which has been around since 1998, is the fundamental building block to integrations going forward.

The "as-a-Service" paradigm extends to Platforms and Infrastructure as well as Software. We need to create a Platform-as-a-Service (Paas) that will be the key to leveraging advancing technologies in the decades to come. Such a platform's primary function should be to integrate the various and disparate digital agents that exist in the ecosystem of the 'Internet of Everything' and the enterprise. This platform would yield benefits today as well as being a platform for leveraging the future technologies. You can think of this platform as a "Playground of the Future," providing a framework for consuming technologies yet to come. Essentially, the challenge becomes one of managing a plethora of advanced technologies and how they interact in an ecosystem of more connected devices and intelligent agents.

Such a platform should offer the following services[8]:

- *Integration platform:* Integrations are key. The platform should support multiple API's, but particularly modern integrations based on XML and web services. The platform should have the ability to integrate both data and processes.
- *Data taxonomy:* All data should be identified as to the source and access rights. As the amount of data proliferates, this will be essential to managing large and disparate datasets, both structured and unstructured, and also key to improving data quality.
- *Common security model:* The platform should enforce a single common security model throughout the organization. Placing it on a platform reduces the need to replicate the security model in different places and ensures consistent application.
- *Multi-organizational model:* The management model is slowly shifting away from the single, hierarchical view of the enterprise to one with myriad connections, with a mix peer-to-peer network relationships along with mini-hierarchies, all dependent on the process or context. It is essential to manage these centrally to simplify the daunting task

of maintaining the connections as well as the ability to re-use relationships wherever they are needed.

- *Orchestration services:* prescribed processes should execute as efficiently as possible, although they often involve software agents across multiple systems. These "traditional" BPM services should be integrated into the platform so as to manage processes across all participants.
- *Big Data Analytics:* Big Data is bigger than any one application, as it encompasses information from several sources, both structured and unstructured. There should be a common analytics engine across the enterprise.
- *Adaptive Case Management:* ACM should be embedded within all processes available to the enterprise. There will be many intelligent resources available, but ACM will provide the framework by maintaining a goal orientation; keeping all resources focused on those goals.

THE SINGULARITY AND BEYOND – THE AGE OF INTELLIGENT MACHINES

Once The Singularity is reached, we will discover the nature of intelligence. Perhaps attempts at reverse engineering will be successful and we will be able to transfer the state of a human brain at a point in time to a computer. Will the machine demonstrate emergent intelligence? Or is intelligence something greater than the sum of its parts, and we will be left with an inert repository of neural states?

Our brains are biological organs that evolved over millions of years to be the centralized control of our bodies. Perhaps a disembodied brain would not function the same way. We are motivated because of our biology, and achieve what we do because we have wants and needs. Striving for survival is a powerful motivator. Would a machine without biology work the same way? If a computer received the neural states from a person at a point in time it would have all of the memories from that person up to that point. The machine would have memories of what it was like to have a body, much like an amputee remembers a missing limb. Perhaps that is enough to satisfy the brain's biological heritage.

Or maybe we'll figure out the software of the brain by then and create digital life forms other than by reverse engineering, that are different from humans in that they are not dependent on biology and derive their motivations from other sources?

One implication of the non-invasive scanning technology described earlier is that we'll then use thought control as an input device to computers; voice recognition will no longer be necessary. People will have to develop skills to control devices while protecting private thoughts. We may then be able to communicate with each other seemingly telepathically.

As mentioned before, The Singularity represents a confluence of technologies. Continued miniaturization will have profound impacts on health care. Nanobots will course through our blood streams, identifying cancer cells and zapping them before they have a chance to grow. Nanobots will be able to clean up the decay around cells associated with aging, so longevity technology will be available to those that are healthy enough (and wealthy enough, initially). As lifespans increase, we will naturally reproduce less, as historically populations have always self-regulated. The turnover of generations (death and new life) will happen slower, or perhaps disappear altogether. Before we opt in to this situation, we must consider the consequences of less generational turnover. How many of us want to live in a world where the only sure thing is taxes?

Humana-form Robot technology will come later, but the digital “brain” could be implanted in a robot and the body parts upgraded as the technology improves. Thus, humans can be fully immortal when they exist in purely digital/robotic form.

ETHICAL IMPLICATIONS

If digital intelligence is emergent, and we have digital sentient life forms, the notion of robots doing our bidding should be reexamined. If a machine can think and feel, does it have rights? Is dominion over sentient robots akin to slavery? Perhaps it would be more appropriate to think of them as our children or colleagues, rather than our servants?

If a person copies their brain into a computer, which “version” of that person is entitled to rights, such as ownership?

Should the computer version be considered the same person? If we replace an arm or leg, or even a heart with a prosthetic device, we still think of them as the same person. Why should it be different with a prosthetic brain?

WILL THEY TAKE OVER THE WORLD (…AND IS THAT A BAD THING?)?

As computing power continues to increase, and digital life forms have access to all of the knowledge of humanity, it seems inevitable that they will quickly surpass our capabilities. Some people have suggested that once they can create other robots, there will be no longer any need for biological humans and we will be subjugated, or eliminated!

It seems obvious that if we have such concerns, we simply shouldn’t build machines that are capable of subjugating us. We can modify the programming to prevent this from happening, much like Asimov’s Three Laws of Robotics[9]:

1. A robot may not injure a human being or, through inaction, allow a human being to come to harm.
2. A robot must obey the orders given it by human beings, except where such orders would conflict with the First Law.
3. A robot must protect its own existence as long as such protection does not conflict with the First or Second Law.

Can we build intelligent machines that are not sentient, or once we build machines that can learn, will we have no control over how they process the knowledge they acquire?

When Watson famously beat the human champions on Jeopardy, Ken Jennings wrote along with his Final Jeopardy answer: "I for one welcome our new computer overlords." Perhaps sentient machines are the next evolutionary step for humanity. If Earth becomes uninhabitable due to climate change or wars, maybe robotics is the only way humanity can continue?

UTOPIA OR DYSTOPIA?

The “Brave New World” we are creating can result in a Utopian or Dystopian result (or, of course, a combination). The Utopian view, as exemplified in Star Trek, is that automation creates all the economic value we need, so people are free to pursue what enriches them, such as the arts or space exploration. The Dystopian view, as is commonly shown in Philip K. Dick novels, is where technology is in the hands of a few powerful people, and our very existence is for someone else's benefit. The future we create for ourselves depends on the decisions we make over the next few decades.

EPILOG

The purpose of this paper is not to draw conclusions about the technology of the near future, but to pose whether we are asking the right questions. With exponential advancement, very powerful technologies will be upon us very quickly. We need to manage them to achieve our goals, and avoid the undesirable outcomes that can result from corruption or neglect.

I'm fascinated by the notion that as software designers, we can rearrange some electrons, and as a result, work gets easier, more value is produced, and maybe *we can create emergent life!* This is modern day alchemy: creating things of great value from basic ingredients. It is our sacred trust to use that power wisely and responsibly for the benefit of humanity.

REFERENCES AND NOTES

[1] – Von Neumann Architecture: computer architecture devised by computer pioneer John Von Neumann. Almost all conventional computers today use this architecture.
http://en.wikipedia.org/wiki/Von_Neumann_architecture

[2] – Multitasking is a term originally applied to how a computer operating system quickly switches between tasks by placing the memory location of the last task on a stack, and then retrieving it from the stack when that task is resumed. Human brains, which have a massively parallel architecture, are notably bad at true multitasking. When a person claims they are “multitasking” it is more likely that they are “getting distracted.”

[3] – (Kuhn 1962) Thomas Kuhn. The Structure of Scientific Revolutions, 1962.

[4] – (Kurzweil 2005) Ray Kurzweil. The Singularity is Near, 2005. Viking.

[5] – Sierra-Cedar HR Systems Survey: http://www.sierra-cedar.com/research/annual-survey/

[6] – Turning Test: a test devised by Alan Turing to determine if an intelligence is human or artificial.

http://en.wikipedia.org/wiki/Turing_test

[7] – (Kurzweil 1992) Ray Kurzweil. Age of Intelligent Machines, 1992. MIT Press.

[8] – (Altman 2013) Roy Altman. Creating an Integrated Platform for Process Intelligence. From Intelligent BPM, 2013. Future Strategies.

[9] – From Science Fiction writer Isaac Asimov’s “Robots” series.