

Chapter 3

The Law of Artificially Intelligent Brains

Placing an Exponent on Intelligence

Benefiting from exponentially improving technologies, in numerous examples what was once considered a task distinctly requiring human intelligence is now being done much faster and more efficiently by artificially intelligent machines. For example, while driving a car requires complex cognitive, perceptual, and motor skills, artificial intelligence is quickly mastering the art of driving and doing so in highly congested traffic. In fact, based on the law of accelerating returns automated cars are improving to the point where public policy may dictate that a person born today may not be able to legally drive when they reach their teenage years. Given the rate of advances in information technologies, within a few years automated cars will become so “smart” that the only necessary response from a human will be a voice activated destination.

Of course, even though sensors collect information and transfer it to an onboard computer, the “mechanical” car itself isn’t becoming smart, the computer directing the car, that is, its brain. And because the raw processing power and capabilities of artificially intelligent machines is directly related to the software, algorithms, and architecture which together comprise its brain, laws that relate to its ability to store information, compute, and communicate will contribute to an emerging law of cyborgs, a central topic of this book. Further, since the hardware, software, and algorithms of an artificially intelligent brain will continue to improve, some computer scientists predict that within a few decades artificial intelligence may exceed human levels of intelligence and pose an existential threat to humanity. For this reason a comprehensive understanding of how the law might apply to an artificially intelligent machine and particularly to the architecture and capabilities of its brain may be essential to the survivability of the human species.

Interestingly, in comparison to human driving performance, after 6 years and 2.7 million km’s driven, the director of Google’s self-driving project reports, “Not

once was the self-driving car the cause of the accident.”¹ However, that’s not to say that self-driving cars haven’t been in an accident, in fact, there have been about a dozen minor accidents during the past 6 years, but in every case, a human driving another car was the cause of the accident. In fact, artificial intelligence is getting so good at what it does, the idea of keeping a human out of the decision making loop in systems involving artificially intelligent machines is being seriously considered. On this point, a few courts have actually found humans negligent for failing to follow the advice provided to them by a computer. Two early cases on this point was *Wells v. U.S.* and *Klein v. U.S.* In *Wells*, a court inferred negligence on the part of a human pilot based on evidence he switched from autopilot to manual control in a crisis situation.² In this example, the brain of a machine was considered the better decision maker than that of the human. And in *Klein*, the court found that in cases of negligence, while the pilot is not required to use autopilot on a landing, his failure to do so was thought inconsistent with good operating procedure and evidence of a failure of due care. Can we conclude from these above examples that there ought to be a law protecting artificially intelligent machines from humans? That’s an interesting question, but I don’t really mean to imply that artificial intelligence is always superior to humans and will always be benevolent. In fact, I am more concerned with the potential dark side to artificial intelligence, than I am living in a world where artificial intelligence serves humanity.

Of course, “artificially intelligent brains” do far more than drive cars, now days semi-autonomous drones deliver packages, some robots assist physicians in surgery, and “artificial intelligence” writes sports and weather reports, or makes stock trades. All of these tasks require an impressive amount of intelligence and in some cases complex motor skills by the machine; however, no one would seriously think robots with these abilities are anywhere near human levels of intelligence. Instead, we humans think that robots with the cognitive and perceptual abilities in the above examples are simply remarkable tools to serve us, and we have the general notion that as advances in technology continues, the future will give us an even better set of tools to meet our needs. I believe this is a naïve view of the future, with dangerous implications for humanity. Agreeing with this position, Elon Musk, CEO of SpaceX and Telsa Motors, describes advances in artificial intelligence as “summoning the demon” and thinks that by creating a rival to human intelligence we are simultaneously building the biggest threat facing the world.³ If we accept the viewpoint advocated by Nick Bostrom, director of Cambridge’s

¹Adrienne LaFrance, 2015, When Google Self-Driving Cars Are in Accidents, Humans Are to Blame, at: <http://www.theatlantic.com/technology/archive/2015/06/every-single-time-a-google-self-driving-car-crashed-a-human-was-to-blame/395183/>.

²*Wells v. U.S.*, 16 Av.Cas. 17914 (W.D. Wash. 1981); *Klein v. U.S.*, 13 Av.Cas. 18137 (D. Md. 1975).

³Samuel Gibbs, 2014, Elon Musk: Artificial Intelligence is our Biggest Existential Threat, at: <http://www.theguardian.com/technology/2014/oct/27/elon-musk-artificial-intelligence-ai-biggest-existential-threat>.

Future of Humanity Institute, that artificial intelligence could pose an existential threat to humanity, then how do we, through our courts and legislators, respond? Some propose completely banning research on artificial intelligence, others propose coding “friendliness” into the “minds” of artificial intelligence (likewise will a future artificial intelligence breed docile humans given our aggressive nature?), while others propose government regulations designed to give artificial intelligence certain rights, and to deny it others. On the last point I believe that there already is an emerging body of law, primarily in the field of intellectual property and constitutional law that speaks to the issue of regulating the architecture and output of an artificially intelligent brain, including the thoughts and speech produced by an artificial intelligence. While these laws and government regulations were enacted to protect the rights of humans and not self-aware machines, I believe they may also contribute to an emerging law of cyborgs, that is, they represent a set of laws that could serve as precedence for future artificially intelligent machines that have reached human levels of intelligence and then argue for rights.

Elon Musk is not alone in his warnings about the potential threat that artificial intelligence could pose to humanity. Cambridge cosmologist Martin Rees, the former Astronomer Royal and President of the Royal Society, addressed similar topics in his 2004 book, *Our Final Hour: A Scientist's Warning*,⁴ as did computer scientist, Bill Joy, co-founder of Sun Microsystems in his 2000 article published in *Wired*, “*Why the Future doesn't Need Us*.”⁵ Yet another concern expressed by some prominent researchers in artificial intelligence and robotics is that by the end of this century we will either be serving the artificially intelligent machines that we are in the process of creating now (who eventually will take charge of their own design), or we could be inconsequential to them and relegated to being the second most intelligent species on the planet. But there may be a third alternative, as proposed by robotics expert Hans Moravec, Google's Ray Kurzweil, and by this author discussed throughout this book- and that alternative is to merge with “them,” thus becoming the product of our technological future and not relegated to the status of bystander.

For reasons discussed below, many of the public are unaware how close we are to a future consisting of machines with human-or-beyond levels of intelligence, and still others (including some prominent AI researchers and philosophers) outright dismiss a future with strong artificial intelligence as either impossible, the subject of science fiction, or too far in the future to give serious thought now.⁶ I think those among the public and academia, who dismiss the dramatic rise of artificial intelligence and its implications for humanity fail to realize that the basic technologies necessary to create artificially intelligent machines are here now, improving exponentially, and leading to the design of machines that will match

⁴Martin Rees, 2004, *Our Final Hour: A Scientist's Warning*, Basic Books; Nick Bostrom, 2014, *Superintelligence: Paths, Dangers, Strategies*, Oxford University Press.

⁵Bill Joy, 2000, *Why the Future Doesn't Need Us*, *Wired* 8.04.

⁶Miguel Nicolelis, 2013, *The Brain is Not Computable*, MIT Technology Review.

humans in intelligence and motor skills, and possibly within 20–30 years. As we get closer to human-like artificial intelligence, I argue that a “law of artificially intelligent brains” will be necessary for our legal institutions to develop and that such an approach will provide a framework in which to discuss many of the social and legal questions that will be shaped by the rise of artificial intelligence. Such issues will speak to the law as it applies to tort liability, contract rights, and criminal culpability for artificially intelligent machines operating autonomously from humans. But a “law of artificially intelligent brains” will also focus on the software, operating systems, and computer architecture of the artificial brain itself. Given the importance of determining the role of artificial intelligence in society, this chapter is not the first to address these issues, there are a number of law review papers and books written on law and robotics (for example, see Gabriel Hallevy, *When Robots Kill: Artificial Intelligence Under Criminal Law* and papers by Law Professor Ryan Calo), and I expect there will be more interest by nations at the forefront of the robotics revolution and the European Union as artificial intelligence becomes more tightly integrated into society, and more autonomous of humans, while asserting claims to be self-aware, and arguing for rights.⁷

With these observations in mind, this chapter examines some of the legal and policy issues that relate to the design of artificially intelligent brains. To discuss these topics I borrow heavily from current law which relates to the software written for computers, and the law relating to the computer architecture which is essential for the machines ability to compute and thus to reason and “think”—these areas of law can be thought of as a law of artificially intelligent brains.

The Numbers Behind Brains

As I lecture on the topic of our cyborg future to merge with artificially intelligent machines, I inevitably get the following question- how close are we to computer hardware and software that matches the human brain in performance? Several prominent roboticists and inventors seem to have settled on a timeframe that is unsettling to some- before midcentury. But first, what computing resources are necessary to reach the goal of human-like artificial intelligence? That is, what storage capacity and raw processing power must an artificially intelligent brain have to match the human brain? And if human and machine brains had similar functionalities and architectures, would the same laws apply to both entities? When thinking about this question, recall that chimpanzees have brains with architectures that are similar to a human brain and a chimp’s behavior is that of a distant relative; yet chimps receive none or almost no individual rights in most jurisdictions.

⁷Gabriel Hallevy, 2013, *When Robots Kill: Artificial Intelligence Under Criminal Law*, Northeastern Press; Ryan Calo, 2015, *Robotics and the Lessons of Cyberlaw*, 103 *California Law Review*.

To receive human-like rights, is a high hurdle to pass; humanity is not generous affording rights to other animals. Doing some “back of the napkin” calculations, we can answer the question posed about the computational resources necessary to match a human brain by looking at the numbers involved in reverse engineering the brain.

For nearly a decade, neuroscientists, computer engineers, and roboticists have been working to reverse engineer the human brain so they can ultimately create a computing architecture based on how the mind works. The key to reverse-engineering the human brain lies in decoding and simulating the cerebral cortex—the seat of cognition. The human cerebral cortex has about 22 billion neurons with trillions of synapses. A supercomputer capable of running a software simulation of the human brain, according to some researchers, would require a machine with a computational capacity of at least 36.8 petaflops and a memory capacity of 3.2 petabytes.⁸ All interesting and technologically possible, but an important and pragmatic question is how many lines of code would be required to simulate a brain? Terry Sejnowski, head of the computational neurobiology lab at the *Salk Institute for Biological Studies* agrees with Ray Kurzweil’s assessment that about a million lines of code may be enough to accomplish that task. Intuitively this number seems low to me, but I did say we are doing “back of the napkin” calculations, so let’s see how the math works. According to Kurzweil: “The design of the brain lies in the blueprint provided by the genome. The human genome has three billion base pairs or six billion bits, which is about 800 million bytes before compression.”⁹ Kurzweil notes that “eliminating redundancies and applying lossless compression, that information can be compressed into about 50 million bytes.”¹⁰ About half of that information is about the brain, which comes down to 25 million bytes, or roughly a million lines of code.¹¹ I have read rebuttals to this number as being far too low, but what amazes me, is that even if we increase the lines of code necessary to simulate the brain even by orders of magnitude; given exponentially accelerating technologies, we are already creating the technology and gaining the knowledge to unlock the mysteries of the brain, so it’s just a matter of time before we can simulate the brain with a million lines of code, or even 100 million lines of code if necessary.

So how close are we to an artificial intelligence with human-like abilities, that is, a being that might argue for legal protection for the software, algorithms, and integrated circuits that allow it to think, problem solve, and to control the motion of its body? In his books and papers, robotics expert Professor Hans Moravec put the 2020s as the time period of human-like robots, this estimate also corresponds to Ray Kurzweil’s prediction. In a paper Hans Moravec published in 1998 (author of

⁸Priya Ganapati, 2010, Reverse-Engineering of Human Brain Likely by 2030, Expert Predicts, at: <http://www.wired.com/2010/08/reverse-engineering-brain-kurzweil/>.

⁹*Id.*

¹⁰*Id.*

¹¹*Id.*

Mind Children: The Future of Robot and Human Intelligence and *Robot: Mere Machine to Transcendent Mind*), which was based on his seminal work on computer vision for robots he estimated that about 100 million MIPS of computer power would be necessary to match “overall human behavior” and that about 100 million megabytes were necessary to match the capacity of the 100-trillion synapse brain.¹² Both of these numbers are less than those provided in the material above (the brain computes in the petaflop range), but since we are already in the petaflop computing range with super computers, we have matched the raw processing power of the brain based on estimates of its raw processing ability to computer. Further, the speed of progress in artificial intelligence is also accelerating as a few years after Moravec’s comments, IBMs Deep Blue defeated the world chess champion using chips designed to operate at 3 million MIPS, or 1/30 of Moravec’s total estimate of human performance (as an aside- if a computer completes 200,000 instructions in 0.02 s, then $200,000/0.02$ would equal 10 MIPS).

Of course we know from neurophysiology that the cerebral cortex with its 22 billion neurons is critically important for human cognition. And to emphasize here the difficulty of building the architecture to create a brain and thus why developing a “law of artificially intelligent brains” will be extremely challenging, if a reasonable estimate of the number of synapses per neuron is 12,500 (some estimate the number to be about 10,000) then the 22 billion cortical neurons alone would require something on the order of 275 trillion transistors to match the number of synapses in the cortex (we are several years out from creating such chips but we will get there). But this level of complexity doesn’t take into account the changing structure of neural networks as we learn and create new memories, and that there may be subcellular processing occurring moving the brain from the paradigm of a single computer to a self-contained Internet with billions of simpler nodes working together in a massively parallel network. So, I view estimates that are specific to when we may build a machine that reaches human levels of intelligence, with a strong interest, and believe that it will happen, but I would not be surprised that as we learn more about the brain and specifically the cortex we may find levels of complexity that will move back the date for the Singularity. However, what’s a few years, or decades, or even centuries after all the planet is 45.4 million centuries old! Clearly the human brain is incredibly complex, but operating under Moore’s law we are now at petaflop (10^{15}) computing with supercomputers, and eventually will reach exaflop computing (10^{18}), at that point computers will be much faster than us based on raw processing power, performing a quintillion calculations per second. My point is this, I fully acknowledge how complex the human brain is, and that modeling its performance will be extremely difficult (the grandest challenge yet for humanity), but the difficulty of creating human-like artificial intelligence, should be considered against the backdrop of exponentially accelerating technologies, and particularly the rapid progress being made in neuroscience and machine learning.

¹²Hans Moravec, When Will Computer Hardware Match the Human Brain? *Journal of Evolution and Technology*, Vol. 1, 1998.

Law and Brains

Acknowledging the challenge of creating an artificially intelligent brain, let's now discuss more specifically laws that might apply to an artificial intelligence. If we (and our machines that help us) can write code (combined with the necessary computational resources) to create an artificially intelligent brain, what current laws relate to the computer code and algorithms that comprise the brain of an artificially intelligent machine? Several areas of law are relevant to this question and thus to our cyborg future, including copyright and patent law, trade secret law, and Constitutional law on the speech output and algorithms of an artificially intelligent brain. Both copyright law and patent law are applicable to some extent to the protection of software and algorithms (meaning the owner of the software has certain rights under copyright law) that contribute to an artificially intelligent brain, and as I have stated throughout this book, contribute to an emerging law of cyborgs. Specifically, an owner of a copyrighted software program has the right (with some exceptions) to: copy the software, create a derivative or modified version of it, and distribute copies of the software to the public by license, sale or otherwise. Are these rights relevant for an artificially intelligent brain? Anyone exercising any of these exclusive rights without permission of the software copyright owner is an infringer and subject to liability for damages or statutory fines.¹³ Interestingly, one could literally steal the mind of an artificially intelligent brain by copying its software and algorithms and if so would be an infringer under copyright law. Similarly, in the cyborg future brain scanning technologies could also be used to copy the thoughts generated by a human mind but since the "software of the mind" is not copyright protected (the output of the mind can be), the person would not be an infringer; maybe we need to change the law.

Out of necessity, absent direct statutes and case law involving artificial intelligence, we look to "human law" as a way to frame issues involving artificial intelligence. In the early days of computing, software developers turned to the statutory protection offered under the Copyright Act to protect the intellectual property rights associated with their programs, arguing that the writing of code was similar to other forms of writing. Similar logic should apply to the code written for the operating system and programs under the direction of an artificially intelligent brain. In the U.S. Copyright Act, the general requirements for copyright are: "original works of authorship fixed in any tangible medium of expression, now known or later developed, from which they can be perceived, reproduced, or otherwise communicated, either directly or with the aid of a machine or device."¹⁴ Further, the Copyright Act defines a computer program (think of the software used by a robot to parse images in a scene), as "A set of statements or instructions to be used directly or indirectly in a computer to bring about a certain result."¹⁵ Examples include software which enables computer vision, robot navigation, or trial-and-error learning.

¹³17 U.S. Code § 102—Subject matter of copyright.

¹⁴U.S. Copyright Act, 17 USC §§ 100 et. seq.

¹⁵U.S. Copyright Act, Sect. 101.

When a computer program is written out on a piece of paper, copyright exists in that work upon its creation so long as the traditional copyright requirements are met (basically, the work must be original, that is, the work must have been developed independently by its author, and there must have been some minimum creativity involved in the work). So, software, which is a central component of an artificially intelligent brain, is clearly copyrightable subject matter. We know that the source code, that is, the language used to write programs for an artificial intelligence is copyright protected, but what about object code? This question was answered in *Apple Computer, Inc. v. Franklin Computer Corp.*, in which a U.S. Circuit Court ruled that programs in both source code and object code are copyright protected.¹⁶ Interestingly, the court rejected the argument that because object code only communicates directly to a machine it should not be protected, this raises the possibility that machine-to-machine communication in an abstract language (a form of machine telepathy), would be copyright protected.

But what about software not written on paper but etched on a chip that compromises the hardware architecture of an artificially intelligent brain? In *Franklin*, the question of whether programs encoded on chips, are utilitarian objects- and thus not subject to copyright protection was addressed. The *Franklin* court rejected the argument that programs were solely “utilitarian” noting that the medium on which the program was encoded should not determine whether the program is subject to copyright. This bodes well for an artificial intelligence arguing for rights to the content of its mind; based on copyright law the memories and thoughts of an artificially intelligent brain written in source or object code may be copyright protected when they are fixed within the integrated circuits compromising an artificially intelligent brain.

Of course, software is more than just individual lines of code, collectively, code performs functions that are essential to the operations of an artificially intelligent brain. In *Franklin* the court rejected the argument that operating systems are not copyright protected because they are “processes, systems or methods of operation.”¹⁷ Instead the court ruled that an operating system is to be considered a work of authorship under the Copyright Act. This holding by the court is directly relevant to an artificially intelligent brain; for example, consider robots, a technology leading us towards the Singularity, clearly robots are getting smarter from one generation to the next based on improvements in their software, algorithms, and the “physical design” of the architecture of their brain. Since an operating system, that is, the backbone of an artificially intelligent brain, is eligible for copyright protection, it cannot be reproduced by another, or a derivative of the artificially intelligent brain made without permission from the copyright holder. The owner of the copyright for software comprising an artificially intelligent brain may (or may not) be the owner of the robot (the owner of the robot may be licensing the software), but in the future the ownership of an artificial intelligence by a human may

¹⁶*Id.*

¹⁷*Apple Computer v. Franklin Computer*, 714 F.2d 1240 (3d Cir. 1983).

be questionable (the 13th Amendment prohibits slavery or involuntary servitude). And if in the future artificially intelligent robots are emancipated from a human owner, copyright law may be used by them as one form of protection for their speech; and software is considered a form of speech.

Summarizing the discussion to this point, the main features of an artificially intelligent brain consists of programs, an operating system, and algorithms, and courts have established that the literal elements of a programs code are protected by copyright law, but there are more issues to discuss for artificially intelligent brains. For example, in *Computer Associates International, Inc. v. Altai, Inc.*,¹⁸ the issue for the Second Circuit Court of Appeals was whether and to what extent copyright protects the non-literal elements of program code, that is, the structure, sequence, and organization of the program. As a basic point, copyright protects the expression of an idea but not the idea itself (ideas are protected by patents or to some extent trade secret law). So for an artificially intelligent brain how the software is written makes all the difference in terms of acquiring protection for its code.

So where do we draw the line between the expression and idea in programs? In *Baker v. Selden* the court stated that things that “must necessarily be used as incident to” the idea are not subject to copyright protection.¹⁹ This opinion, however, gave no advice on how to separate an idea from its expression. Facing a similar issue, *Whelan v. Jaslow*, a landmark case in defining principles that apply to the copyright of computer software, the Court attempted to delineate the differences between “idea and expression by saying that the function of the work is the idea and everything else not necessary to the function is the expression of the idea.”²⁰ But other courts have found this approach unworkable, and have adopted the filtration approach taken by the Second Circuit in *Computer Associates Int’l. v. Altai*.²¹ That approach separates the code’s ideas and public domain elements from its expression and then extends copyright protection only to the expression.²² In the abstraction-filtration-comparison test the court first determines the allegedly infringed program’s constituent structural parts. Then, the parts are filtered to extract any non-protected elements. Non-protected elements include: elements made for efficiency (i.e. elements with a limited number of ways it can be expressed and thus incidental to the idea), elements dictated by external factors (i.e. standard techniques), and design elements taken from the public domain.²³ Any of these non-protected elements are thrown out and the remaining elements are compared with the allegedly infringing program’s elements to determine

¹⁸*Computer Associates International, Inc. v. Altai, Inc.*, 982 F.2d 693 (2d Cir. 1992).

¹⁹*Baker v. Selden*, 101 U.S. 99 (1879); Copyright Act § 102(b) (the subject matter of copyright).

²⁰*Whelan Assocs., Inc. v. Jaslow Dental Lab. Inc.*, 609 F. Supp. 1307, 225 U.S.P.Q. (BNA) 156 (E.D. Pa. 1985).

²¹*Id.*, note 18.

²²*Id.* note 20.

²³*Id.* note 18.

substantial similarity. In my view the above approaches to determining which aspects of software are copyright protected will be difficult to apply to the software of an artificially intelligent brain, just as determining which aspects of human thinking represent function versus expression would be difficult. For that reason courts may need to devise another test suitable for the cyborg age in which to decide what aspects of code are copyrightable subject matter especially the code comprising the input and output of an artificially intelligent brain.

More About Artificially Intelligent Brains

Repeating a basic point, an artificially intelligent brain will consist of the computer architecture, software, and algorithms to direct its behavior and to make sense of the world. Will such a brain with appropriate computational resources and software reach consciousness? This is a question of great debate, but there is an established neuroscientific consensus that the human mind is largely an emergent property of the information processing resulting from the 100 billion neurons comprising its architecture. And we know from the above discussion that much of the artificially intelligent brains software is “protected” by intellectual property law. So an emerging law of artificially intelligent brains is beginning to take shape. And interestingly, while we can conclude that there is no “law of neurons,” there is a law of software. Thus, the brain of an artificially intelligent machine can be scrutinized under the law in ways a human brain cannot.

Another question of interest to those designing the cyborg future and wondering how the law might apply is whether artificially intelligent brains will surpass the human brain in capabilities. As we did above, let’s think like an engineer for a moment and focus on the quantitative aspects of brains. Given that the electrochemical signals that human brains use to achieve thought travels at about 150 m/s, this is orders of magnitude slower than the speed at which electronic signals are sent by computers. Therefore, a massively parallel electronic counterpart of a human biological brain will be able to think millions and eventually trillions of times faster than our naturally evolved system. Also, consider that neurons can generate a maximum of about 1000 action potentials per second, whereas the clock speed of microprocessors reached 5.5 GHz in 2013, which is about five million times faster, this means that in some respects, computer brains are already superior in performance to human brains.²⁴ But supercomputers also can have energy requirements that compete with some municipalities, and have grown larger than the laboratory-sized calculating machines at the infancy of computers.²⁵ The human brain, meanwhile, uses roughly 20 watts and occupies a small volume.

²⁴Mind Uploading, Wikipedia, at: https://en.wikipedia.org/wiki/Mind_uploading.

²⁵Geoffrey Mohan, 2014, Cognitive computer chip apes brain architecture, at: <http://www.latimes.com/science/sciencenow/la-sci-sn-brain-chip-computer-20140807-story.html>.

However, since computers are improving in computational performance exponentially, we can expect a laptop computer to have the computational power of the human brain within a few years.

The recent rise of artificial intelligence has been spurred by many factors, including a tremendous decrease in the price of information technologies combined with an exponential increase in performance. With so much computing power available, algorithms are more-and-more capable of understanding languages, recognizing images, and performing more autonomously from humans.²⁶ For example, artificially intelligent machines in manufacturing are not only getting smarter, but their costs are tumbling- and this price performance relationship is pervasive in all information technologies. On this point, consider that the labor costs for an essential technology for the future development of artificially intelligent machines- electronics manufacturing has plummeted. For electronics manufacturing, robots are becoming so cost effective that in many cases it already costs just a few dollars an hour to use a robot for a routine assembly task versus six times more for an average human worker. How long will it be before the robots are designing their next generation based on their own criteria and displace even more humans from the workplace?

Interestingly, Hans Moravec commented that human-like performance from machines will only make economic sense when their “brains” cost about \$1000- and when can we expect that? Our evidence suggests around 2029 for replicating the human brain. It is important to note for our cyborg future, following Moore’s law, the price-performance of computers will continue to double every 18 months or so at least for the next decade, and once we reach human levels of performance for robots, they will continue to get smarter, after all, their evolution is not based on biology and thus does not rely on random mutations in genes to work their way into the human genome. Still some critics argue that Moore’s law is running out of steam, but if the past trend in computing technology continues, another technology will take over for current chip design techniques, and will continue the exponential improvement in computing power to midcentury and beyond. The amazing power of exponential growth in information technologies is experienced by people every day, in fact, just consider, we all carry the proof of exponentially improving technologies in our hand as the cell phone we use now is a million times cheaper and a thousand times more powerful than a supercomputer of the mid-seventies. And by the way, every cell phone call is routed using artificial intelligence. Many don’t realize it, but we are completely dependent on artificial intelligence now, from air traffic control systems to home appliances, artificial intelligence is in the background, silently doing its job.

Continuing Hans Moravec’s comments above on the desirability of a \$1000 computer brain, as far back as 1999 and in his recent writings, Ray Kurzweil predicted that by the 2020s a \$1,000 laptop would have the computing power and

²⁶Rise of the Machines, 2015, The Economist, at: <http://www.economist.com/news/briefing/21650526-artificial-intelligence-scares-peopleexcessively-so-rise-machines>.

storage capacity of a human brain (100 billion neurons, 100 trillion synapses).²⁷ In fact, we are well on our way to creating a low-cost computer with this amount of raw processing power. Relying on exponential growth curves, Kurzweil predicted that the hardware needed to emulate the human brain could be ready as early as 2020—this could be done using technologies such as graphics processing units which use a massively parallel architecture, which I might add is an ideal architecture for brain-software algorithms. While critics that are opposed to the idea of artificially intelligent machines gaining human-like intelligence worry that this outcome could prove disastrous to humanity, many also acknowledge that we are now entering a time when computers have the processing power necessary to match the brain’s computational abilities. But more than computational power is needed to create human-like artificial intelligence. For example, critics point out that current software is nowhere near being able to model the human brain in its ability to process information and make decisions.

While the critics are right that computational resources are necessary but not sufficient to create human-like intelligence, still the software for artificial intelligence and the algorithms to mimic the decision making of the brain (that is, to simulate neuronal networks) are also making great strides. In fact, Kurzweil predicts that software to accurately model the brain will take only a little longer to develop than acquiring the processing power of the brain, putting the date at 2029.²⁸ But in my view while creating artificial intelligence that matches a human in ability will be a landmark event in humanity’s history, what humanity really needs to focus on is what happens after the Singularity is reached- how do we survive in the shadow of intelligent beings far superior to us? A major thesis of this book is that for humanity’s survival, we need to merge with our technological progeny, or as Hans Moravec puts it- our “mind children.” For us to merge with artificially intelligent machines, enter cyborg technologies (which will be a key factor for our future survivability). As computing technology keeps advancing at an exponential rate, within a few decades, we will have the combined intelligence of the human race accessible by a neuroprosthetic device implanted within our brain. This capability will be essential for the survival of our species once the Singularity is reached by artificial intelligence.

Even though the Singularity is predicted to be only a few decades away, computers already have a big advantage over us: they are interconnected via the Internet and share information with each other billions of times faster than we humans are able to do using the limited communication bandwidth provided to us by nature. This means that a law of artificially intelligent brains needs to consider how a collective form of artificial intelligence shares liability and other responsibilities under the law. The most accurate predictor of the future (at least within the timeframe of a few decades), is Google’s Ray Kurzweil who says that by the

²⁷Ray Kurzweil, 2006, *The Singularity is Near: When Humans Transcend Biology*, Penguin Books.

²⁸*Id.*

2040s, non-biological intelligence will be a billion times more capable than biological intelligence, that is, us.²⁹ The reader may be wondering how is this even possible and why so soon? After all, current artificial intelligence, while remarkably smart in limited domains, lacks the general intelligence and common sense displayed by a 4 year old. But the problem that people have in understanding the future, as pointed out by Kurzweil in his fascinating books about what the future might offer (see for example, *The Age of Spiritual Machines* and *The Singularity is Near*), is that people are linear thinkers, they extrapolate the world they live in now along a straight line to predict where technology will be in the future. Centuries ago linear thinking about technology worked quite well, but around the middle of the nineteenth century, the rate of technological advancements began to noticeably speed up. Considering computing resources, plotting the exponential growth of many computing-based technologies has shown that the growth rate for information technologies is decidedly not linear.

Believe me a liner scale versus exponential scale for technological progress and particularly the brain of an artificial intelligence makes all the difference. I'll prove that to you with an example. Do you want to be rich? Tongue-in-cheek, I argue it's easy to do. Here's how- let's use our spare change and allocate a 31-day month to reach our goal of riches. The first day of the month, place a penny on day one of the calendar, and double the amount placed on the calendar for each additional day until the end of the month. What happens? The second day you have 2 cents, and the third day 4 cents, the fourth day 8 cents and by the end of the week you have 64 cents total. One week of stacking pennies has gone by, are you feeling rich yet? Continuing, day 14 you have \$81.92. At that point I say to you, you are about half way to the end of the month, do you still believe me that you are going to be very rich in 17 more days based on doubling pennies from one day to the next? Most people respond no, they use linear thinking to scale the problem in their mind, and claim that they will end up with a few hundred dollars at most (which is a lot more than they thought they would have at the beginning of the month). So, I continue the exercise (with your pennies!). By day 21 we have \$10,485.76 and it's definitely getting interesting, but only 10 days to go until the end of the month and we're not rich yet. By day 25 you have accumulated \$167,772.16, and by now you are likely fascinated with the concept of exponential growth. Finishing the exercise, day 31 we have \$10,737,418.24. To reiterate the point about exponential growth, we achieved that amazing result due to the doubling of pennies from one day to the next. It turns out that doubling pennies has a lot to do with the rise of artificial intelligence, the performance that their brains will be able to achieve, and our future to merge with artificially intelligent machines. Given that information technologies are improving exponentially (note that the magnitude of the exponent signifying the growth is important), I think you will agree with me that remarkable technologies await us. In fact, you may be thinking, a world of incredibly smart tools to serve us, that's the future which awaits humanity. But not so fast, the

²⁹*Id.*

problem as I see it is that our tools will become much smarter than us (they may even look like us, see Chap. 7: *The Law of Looks and Artificial Bodies*), and then who will be master and who will be the servant?

While the doubling of technological resources is important so too is the time frame between doublings (which also makes all the difference). In our get rich example, instead of using days on a calendar, let's say we used 10 year time periods, so after 10 years, we have 2 cents, and in 20 years we have 4 cents, eventually we will get to over ten million as we did in the above example, but who wants to wait three centuries? I argue that exponentially improving technologies will change everything due to the power of doublings and the short time intervals between the doublings. Let's say starting now that humans doubled their intelligence from one generation to the next. Clearly that's impossible, but if so, about 18 years from now, a person with twice the intelligence of the general population would be born. The problem is we can't significantly alter our intelligence in that time period. Because we are products of the exceedingly slow process of evolution, the cycle time for improving human intelligence is measured in millennium. But there's a solution to keeping up with increasingly intelligent machines- exponentially improving technologies integrated into our body. According to Moore's law the cycle time to double the number of transistors on a chip is about 18 months. Using the above example, while the human had to wait 18 years to double its intelligence, in the same time period, a computer would have experienced 12 doublings of computational power (note that computer power is necessary but not sufficient to produce artificial intelligence, so I am not implying that the AI would be 12 times smarter). Referring back to our example using penny's to get rich, 12 doublings is the difference between 2 cents (one time period) versus \$20.48 (12 time periods); so you can see if we remain as we are now, the brain of our "competition" will rapidly leave us behind. Remember, because supercomputers are already computing in the petaflop range (a quadrillion floating point operations per second)- so double 20 petaflops 12 times and then compare that to the processing power of the brain (which is a petaflop computing machine and without technical enhancements will continue to be so based on biology). The important point is this, with the use of exponentially accelerating technologies, we are essentially placing an exponent on the increase in computational power of the technologies which may lead to an artificial intelligence that matches then exceeds us. But in theory the same principle could also work for the human brain, that is, once the brain is wirelessly connected to the cloud through a neuroprosthetic device. In fact, by 2045, (as predicted by Ray Kurzweil, we could multiply our intelligence a billion fold by linking wirelessly from our neocortex to a synthetic neocortex in the cloud.³⁰

As should be clear by now, predictions about the future are a byproduct of understanding the power of Moore's Law, and more generally of the "Law of

³⁰Peter Diamandis, 2015, Ray Kurzweil's Mind-Boggling Predictions for the Next 25 Years, at: <http://singularityhub.com/2015/01/26/ray-kurzweils-mind-boggling-predictions-for-the-next-25-years/>.

Accelerating Returns”.³¹ If the cycle time to double computational processing power is about 18 months, then given the power of super computers now, a few doublings represents an incredible improvement in computing power over where we are now. For example, if we start with 40 petaflops, after three doublings (less than 6 years) we are already at 320 petaflops; 8x the computing power in less than 6 years is an amazing increase (I’ll take that in the stock market!). And given such a short cycle time between doublings, it is no wonder that Kurzweil and others have been so successful predicting remarkable engineering advances, they only have to postulate about what’s possible by exponentially improving technologies 20 years out or less. Generally, information technologies follow an exponential growth curve based on the principle that the computing power that enables them doubles every 18 months or so. In fact, information technology has seen exponential growth for decades. This has led to vast improvements in memory, processing power, software algorithms, voice recognition and overall machine intelligence.³² And with the increased raw processing power for computers, so too have advancements been made in algorithms to emulate thinking, and the design of chips which process information more as the brain does (mostly parallel processing) compared to the computers of past decades designed to process information based on the von Neuman computer architecture. The law which relates to algorithms and computer chips is, of course, part of the emerging law of cyborgs and artificially intelligent brains.

Machine Learning and Brain Architectures

Information technologies improving exponentially are not the only area of science and engineering making tremendous strides leading to artificial general intelligence and our future to merge with artificially intelligent machines- so too is neuroscience generating exponentially growing volumes of data and knowledge on specific aspects of the brain.³³ In fact, thousands of neuroscientists are working to map the brain across all its levels and functions. It is likely that research in neuroscience will ultimately reveal the detailed mechanisms which led from genes to cells to neuronal circuits, and ultimately to cognition and behavior—the biology that makes us human and conscious. This knowledge will help transform computing making artificial general intelligence all that more possible. But of course developing human-like artificial intelligence will be extremely difficult and challenging, in fact, the human brain performs computations inaccessible to the most

³¹*Id.*

³²See generally, 2029 timeline contents, at: http://www.futuretimeline.net/21stcentury/2029.htm#VXxble_bJjo.

³³European Commission, 2014, From lighter airplanes to new treatments for brain diseases, at: http://europa.eu/rapid/press-release_MEMO-14-531_en.htm.

powerful of today's computers—all while consuming no more power than a light bulb. According to Europe's digital agenda for the future which is part of Europe's Human Brain Project, understanding how the brain “computes reliably with unreliable elements, and how different elements of the brain communicate, can provide the key to a completely new category of hardware, neuromorphic computing systems; and to a paradigm shift for computing as a whole.”³⁴ The economic and industrial impact is potentially enormous but the ultimate result will likely be an artificial intelligence that exceeds us unless our destiny is to merge with our technological progeny.

The phrase “artificial intelligence” often brings to mind futuristic visions of human-like machines; however the ability of a machine to learn is a concept that is already in play today. And the machine's ability to learn is a direct result of its brain architecture, software, and algorithms. So how do current computers learn and acquire the knowledge to be intelligent? One approach to creating human-like artificial intelligence is to take a “machine learning approach” which allows a computer program to discern the key features of one dataset and then apply what it has learned to make predictions about another.³⁵ Familiar examples of this machine learning approach includes according to *Biome*, “optical character recognition, spam filtering, automatic face recognition, and various data mining applications.”³⁶

While a super computer has the raw processing power of a brain (in the range of petaflop computing), without implementing the rules/algorithms which enable thinking this amount of processing power cannot lead to artificial general intelligence. But clearly advances are being made using a variety of approaches to create computers that think and reason more as humans do; some of the techniques rely on algorithms, and others on the design of the architecture of the computing hardware itself. For example, one of the techniques being used to create a computer that “thinks” is an approach termed “deep learning” which is actually a refinement of the field of machine learning.³⁷ With deep learning, machines teach themselves without human intervention by crunching large sets of data and then statistically analyzing the data looking for patterns. This type of machine learning is especially powerful because it represents a way of getting computers to know things when they see them, by producing for themselves the rules programmers cannot painstakingly specify for every event and contingency that may occur in the world. Here I should make the point, no matter what the algorithm, software, or computing architecture, these components of “thinking” contribute to a developing field of jurisprudence relating to a law of artificially intelligent brains.

³⁴The Human Brain Project, The European Commission, at: <https://ec.europa.eu/digital-agenda/en/human-brain-project>.

³⁵Peter, Flach, 2012, *Machine Learning: The Art and Science of Algorithms that Make Sense of Data*, Cambridge University Press; Nikhil Buduma, 2015, *Fundamentals of Deep Learning: Designing Next-Generation Artificial Intelligence Algorithms*, O'Reilly Media.

³⁶The rise of machine learning: how to avoid the pitfalls in data analysis, 2014, at: <http://biome.biomedcentral.com/the-rise-of-machine-learning-how-to-avoid-the-pitfalls-in-data-analysis/>.

³⁷*Id.* note 35.

Discussing whether algorithms, basic components of an artificially intelligent brain, are a form of speech, Duke University Law Professor Stuart Benjamin points out that “many human activities involve the transmission of bits, according to the algorithms and protocols created by humans and implemented by machines.”³⁸ In my view, Benjamin’s use of the phrase “created by humans” is a qualifier applied to speech that may disappear within a few decades as artificial intelligence gets smarter.³⁹ Benjamin poses the question- “Are these algorithm-based outputs speech, under the First Amendment?” We know that computers “think” by manipulating bits, done by using algorithms such as those that statistically analyze data, for example, to detect lines and edges in a scene to identify an image. In fact, computer code is basically a set of instructions and algorithms (is the human mind the same?). According to Benjamin, even if algorithms are not speech their output may be and thus subject to at least some First Amendment protection. Of importance for a law of artificially intelligent brains using software and algorithms to produce behavioral outputs, is the case of *Sorrell v. IMS Health, Inc.*, in which the Supreme Court held that the creation and dissemination of information are speech within the meaning of the First Amendment.⁴⁰ According to Professor Benjamin, by “extending the First Amendment to messages produced by artificial intelligence, we would be treating the products of machines like those of human minds.”⁴¹ Thus, in his view we could then say that speech was truly created and not just transmitted, or aided, by a machine. In fact, the issue of whether the output of a computer is speech was addressed in *Brown v. Entertainment Merchants Ass’n*, where the output of video games were considered speech because the court concluded they communicated ideas like a literary device.⁴² I believe this holding serves as precedence for a future artificial intelligence claiming rights to its speech. However, the issue of granting “free speech” to computers is problematic, according to Columbia Law Professor Tim Wu.⁴³ According to Wu, “computer programs are utilitarian instruments, meant to serve us.”⁴⁴ He points out that “the First Amendment is intended to protect actual humans against state censorship.”⁴⁵ Wu argues that nonhuman or automated choices should not be granted the full protection of the First Amendment, and often should not be considered speech at all. In Professor Wu’s view, to give computers the rights intended for humans is to “elevate our machines above ourselves.”⁴⁶ Responding to Wu’s

³⁸Stuart Benjamin, Algorithms and Speech, University of Pennsylvania Law Review, Vol. 161, No. 6, May 2013, 1445–1494.

³⁹Pamela Samuelson, Allocating Ownership Rights in Computer-Generated Works, 47 *Pitt. L. Rev.* 1185, 1985.

⁴⁰*Sorrell v. IMS Health Inc.*, 131 S.Ct. 2653 (2011).

⁴¹Stuart Benjamin, *id.*, note 38.

⁴²*Brown v. Entertainment Merchants Ass’n*, 131 S.Ct. 2729 (2011).

⁴³Tim Wu, Machine Speech, University of Pennsylvania Law Review, Vol.161, 1495–1533.

⁴⁴*Id.*

⁴⁵*Id.*

⁴⁶*Id.*

argument, I believe we are a few decades away from having to confront the issue of whether artificial intelligence is superior to humans in intelligence because information technologies are improving exponentially; at that point courts may have no choice but to determine the boundaries of protection for speech produced by artificially intelligent machines.

A central aspect of an artificially intelligent brain, is an algorithm, and when considered solely as a mathematical formula expressing a universal principle of nature (e.g., gravity), is not patentable, because the patent would create a huge and fundamental monopoly over laws of nature. This general rule against patenting algorithms was at one time applied to computer software, because software largely consists of procedural instructions in mathematical form that makes a computer accomplish a certain and definite result. Now days, the U.S. Patent and Trademark Office will allow patents on that aspect of an algorithm that accomplishes a useful and concrete result, and provided the software patent is tied to a particular machine or transforms an article into a different state. For example, in an important case about the patentability of business methods expressed in code, *State Street*, the court ruled that mathematical algorithms are nonpatentable only when they are “nothing more than abstract ideas consisting of disembodied concepts that are not useful.”⁴⁷

Like the human brain, deep learning algorithms are used by artificially intelligent machines in an attempt to learn multi-level representations of data, embodying a hierarchy of factors that may explain them. Such algorithms have also been demonstrated to be effective both at uncovering underlying structure in data, and have been successfully applied to a large variety of problems ranging from image classification, to natural language processing and speech recognition. Interestingly, MIT researchers discovered that a deep-learning system designed to recognize and classify scenes also learned how to recognize individual objects.⁴⁸ To discover this they used a deep learning system to train a successful scene-classifier, which proved to be between 25 and 33 % more accurate than its best predecessor. This result implies that scene-recognition and object-recognition systems could work in concert or could be mutually reinforcing; this is one of many steps being made by thousands of researchers in the direction of creating machines with human-like thinking abilities.

Another promising approach that mimics human learning, and thus may constitute a critically important aspect of knowledge acquisition for an artificially intelligent brain, is being investigated by Professor Pieter Abbeel at UC Berkeley who with colleagues has developed a type of reinforcement learning which works by having a robot complete various tasks—putting a clothes hanger on a rack, assembling a toy plane, screwing a cap on a water bottle, and more—without pre-programmed details

⁴⁷*State Street Bank and Trust Company v. Signature Financial Group, Inc.*, 149 F.3d 1368 (Fed. Cir. 1998).

⁴⁸Yoshua Bengio, Ian Goodfellow, and Aaron Courville, *Deep Learning*, MIT Press, In preparation.

about its surroundings.⁴⁹ “Most robotic applications are in controlled environments where objects are in predictable positions,” said UC Berkeley faculty member Trevor Darrell, director of the *Berkeley Vision and Learning Center*.⁵⁰ According to Darrell, “the challenge of putting robots into real-life settings, like homes or offices, is that those environments are constantly changing.”⁵¹ To be intelligent, the robot must be able to perceive and adapt to its surroundings. Conventional, but impractical, approaches to helping a robot make its way through a 3D world include pre-programming it to handle the vast range of possible scenarios or creating simulated environments within which the robot operates.⁵² Instead, the UC Berkeley researchers are using deep learning techniques, which is loosely inspired by the neural circuitry of the human brain when it perceives and interacts with the world.⁵³ The techniques for machine learning described here are a clear departure from the brittle method of having to program every rule into the mind of a machine else it doesn’t know the rule, imagine parents having to do that with their kids.

In the world of artificial intelligence, deep learning programs create “neural nets” in which layers of artificial neurons process overlapping raw sensory data, whether it be sound waves or image pixels.⁵⁴ This helps the robot recognize patterns and categories among the data it is receiving. According to Sarah Yang, “People who use Siri on their iPhones, Google’s speech-to-text program or Google Street View might already have benefited from the significant advances deep learning has provided in speech and vision recognition.”⁵⁵ However, applying a deep reinforcement learning approach to motor tasks in unstructured 3D environments has been far more challenging, since the task goes beyond the passive recognition of images and sounds. UC Berkeley’s Trevor Darrell pointed out that “We still have a long way to go before our robots can learn to clean a house or sort laundry, but our initial results indicate that these kinds of deep learning techniques can have a transformative effect in terms of enabling robots to learn complex tasks entirely from scratch.”⁵⁶ Based on Darrell’s work and other researchers exploring the use of deep learning for robots, in the next 5–10 years, significant advances in robot learning capabilities may occur.⁵⁷ This observation coincides with my view that based on the law of accelerating returns, we are entering a time period in

⁴⁹Amy Jiang, 2015, UC Berkeley Researchers Enable Robots to Learn Through Trial, Error, The Daily Californian, at: <http://www.dailycal.org/2015/05/24/uc-berkeley-researchers-enable-robots-to-learn-through-trial-error/>.

⁵⁰Sarah Yang, 2015, New ‘deep learning’ technique enables robot mastery of skills via trial and error, at: [deep-learning-robot-masters-skills-via-trial-and-error](#).

⁵¹*Id.*

⁵²*Id.*

⁵³*Id.*

⁵⁴*Id.*

⁵⁵*Id.*

⁵⁶UC Berkeley Robot Learns By Trial and Error, 2015, Robot Magazine, at: <http://www.botmag.com/uc-berkeley-robot-learns-by-trial-and-error/>.

⁵⁷Sarah Young, *Id.*, at note 50, discussing the work of Trevor Darrell.

which noticeable improvements will occur between one version of a robot and the next (similarly to cell phones). So before the public is fully aware, the age of artificially intelligent robots with human-like intelligence may be upon us.

I should add, the principal forms of intellectual property protection for artificially intelligent machines which use neural networks in the United States include patents, copyrights, trade secrets, and mask works (see next sections). As with previous forms of new technology, some aspects of neural networks and the software of an artificially intelligent brain transcend existing legal categories. This is primarily due to their dynamic nature, as well as the impossibility of predefining the trained state of the system. As a result, these aspects of neural network technology may be left with limited protection until Congress or the courts respond by customizing current laws to fit this technology, much as they have already done with computer software.

Brain Architecture

Generally, the architecture of a machine's brain in combination with software and algorithms will determine its ability to compute and therefore to exhibit intelligence. One of the factors driving increased intelligence in machines is Moore's law- however the physical limits possible by etching circuits on a silicon chip is beginning to be reached, so will Moore's law run its course, and will the day of exponential growth for computing resources be over? I don't think so, there are numerous techniques being investigated which if successful, will continue the exponential growth of computing resources. In fact, IBM is studying the use of fully integrated silicon chips using high-speed pulses of light to transmit information. This means the chip will be able to move data at rapid speeds and longer distances than current computing systems. Since the silicon photonic chip is wavelength-multiplexed, it can transmit multiple wavelengths of light thus increasing the bandwidth of information transmission compared to technology which exists today. This discussion highlights the fact that while the human brain is based on a particular architecture and a relatively slow transmission rate of signals, an artificially intelligent brain has the ability to dramatically change along with advances in technology. For this reason, I wonder whether a law of artificially intelligent brains will always lag behind technological developments? If so, our role as human legislators will continually be challenged and perhaps in the future an artificial intelligence will get involved in the rule making.

One of the most interesting technologies for computing being investigated now is quantum computing.⁵⁸ Instead of encoding information as either a zero or a one, as today's computers do, quantum computers will use quantum bits, or qubits,

⁵⁸Elenor G. Rieffel and Wolfgang H. Polak, 2014, *Quantum Computing: A Gentle Introduction*, MIT Press.

whose states encode an entire range of possibilities by capitalizing on the quantum phenomena of superposition and entanglement. If quantum computers are successfully developed, computations that would take today's computers thousands of years to perform, would take only a few minutes; imagine if our artificially intelligent progeny had this capacity to think, imagine if we did. Another promising area of research for computing is being led by IBM and the Defense Advanced Research Projects Agency (DARPA) and is aimed at the development of cognitive-computing chips using new materials, such as gallium arsenide, carbon nanotubes, and graphene. In fact, an IBM-led research team has created a computer chip that is designed to mimic the brain's architecture. At the time of the writing, the "TrueNorth" chip is a 5.4 billion transistor chip with one million programmable neurons and 256 million synapses, but in contrast, remember, the brain has about 100 trillion synapses. However, in less than 20 years a neuromorphic chip may reach the brain's level of complexity; further, the TrueNorth chip is currently 1,000 times as energy efficient as a conventional chip.⁵⁹

For all the exponential advances in processing speed, materials, and manufacturing, digital computing today relies on an architecture rooted in the 1940s and with a well-known "bottleneck" between the processor and memory. Specifically, the von Neumann computer architecture is the standard platform of computing and includes three components: a CPU; a slow-to-access storage area, like a hard drive; and a secondary fast-access memory (RAM). A computer with a von Neumann computer architecture stores instructions as binary values (creating the stored program concept) and executes instructions sequentially—that is, the processor fetches instructions one at a time and processes them in sequence.⁶⁰ In terms of thinking and reasoning about the world, an artificially intelligent brain uses integrated circuits to perform calculations and to manipulate the symbols representing "computer thought;" this is done using circuits consisting of resistors, transistors, capacitors, etc.—all etched onto a tiny chip, and connected together to achieve a common goal. Integrated circuits come in all sorts: single-circuit logic gates, voltage regulators, motor controllers, microcontrollers, microprocessors, the list just goes on-and-on, but think of these components as features comprising the architecture of an artificially intelligent brain.

The von Neumann sequential method of information processing has limitations, not the least of which is that it fails to perform anywhere near the capability of the three pound brain setting on our shoulders (computers beat us with brute force computing not with eloquent massively parallel processing). But much research is being done to determine how the brain functions, to reverse engineer the brain's neurocircuitry, to fabricate chips that perform like the human brain does, and to

⁵⁹Dharmendra S. Modha, Introducing a Brain-inspired Computer, TrueNorth's neurons to revolutionize system architecture, (accessed 2015) at: <http://www.research.ibm.com/articles/brain-chip.shtml>.

⁶⁰Irv Englander, 2004, *The Architecture of Computer Hardware, Systems Software, & Networking: An Information Technology Approach*, Wiley.

write the software and algorithms to mimic human thinking.⁶¹ For example, neuromorphic computing, a concept developed by Carver Mead in the late 1980s, involves the use of very-large-scale integration (VLSI) systems containing circuits to mimic neuro-biological architectures present in the nervous system. Specifically, the VLSI systems are used to model perception, motor control, and multisensory integration.⁶²

So, to summarize this brief discussion on the architecture of an artificially intelligent brain, integrated circuits form a main component of the architecture of the machines brain, and consist of billions of tiny inter-connecting electrical paths meticulously arranged onto a single piece of material, such as silicon. Designing an integrated circuit chip is not a simple feat, in fact, as chips become even smaller, issues such as hot spots, leakage etc., make an effective, power-efficient design extremely difficult to achieve.⁶³ Successful designs usually result from the enormous effort of highly qualified experts coupled with huge financial investments. However, copying each layer of an integrated circuit and preparing “pirated“ integrated circuits can be done with comparatively little effort. According to Charl Goussard, “taking into account the enormous effort and cost to develop an integrated circuit design, the wide industrial applicability, the constant demand for improvement, and the ease at which such designs can be copied, it seems logical that some form of statutory protection should be afforded for the designers or owners of these designs.”⁶⁴ But where do we find these rights? And of course, by now, you may be thinking as I do, that any laws which relate to the software, algorithms, and architecture of a computer, serves as precedence for a law of artificially intelligent brains.

Hardware Protection for Artificially Intelligent Brains

Over the past few decades during which software development has become more sophisticated, courts have pointed out the difference in purpose between copyright and patent laws for software. The broad protection for software as provided by patent law, must meet the standards of novelty and nonobviousness in order for a patent to be granted; the standards for copyright protection are originality and some level of creative expression. For software, the purpose of copyright is to protect particular expressions of an idea that are written in source code by a

⁶¹Ludmila, I. Kuncheva, 2014, *Combining Pattern Classifiers: Methods and Algorithm*, Wiley.

⁶²See generally, *Neuromorphic Computing*, Wikipedia, at: https://en.wikipedia.org/wiki/Neuromorphic_engineering; NAIP PATENT BLOG, At: <http://naipblog.blogspot.com/2009/08/brief-overview-of-ic-design-protection.html>.

⁶³Peter McCrorie, *On-Chip Thermal Analysis Is Becoming Mandatory*, at: <http://chipdesignmag.com/display.php?articleId=2171>.

⁶⁴Charl Goussard, 2009, *What is Integrated Circuit Design?* at: <http://naipblog.blogspot.com/2009/08/brief-overview-of-ic-design-protection.html>.

programmer (which is then compiled into object code), not the idea itself; an idea is the subject of patent law. Both copyright and patent law have a role to play in protecting software as intellectual property and thus contribute to a law of artificially intelligent brains.⁶⁵

Patent law which protects ideas is clearly relevant for the components of an artificially intelligent brain, for example, circuits designed to model the properties of neurons have received patent protection. One example is a “silicon neuron” patent (U.S. patent 5648926 A) that describes an integrated circuit that is designed to emulate the functions of a biological neuron; many other patents have been awarded in this area. For software, the U.S. issues patents if the patent application describes the code in relation to computer hardware and related devices and limits the software to specific uses- this may include software that connects to and runs hardware components. This description of patent protection for software seems directly applicable to an artificially intelligent brain as the software running the brain is used to control the effectors and actuators of the machine.

As noted earlier in this chapter, of particular importance to our cyborg future is that copyright also extends to programs etched on chips. Once chips are fabricated, they are plugged into the computer and become part of the computer’s brain architecture. This means that a computer’s brain has rights under copyright law that is not afforded human brains which of course consist of billions of neurons. Generally utilitarian objects are not the subject of copyright protection and chips are clearly utilitarian, but as stated earlier in this chapter, in a case dealing with software the *Franklin* Court rejected the argument that software encoded on chips was to be considered “utilitarian” and thus not copyright protected noting the medium on which the program is encoded should not determine whether the program itself is protected under copyright.

To provide the legal protection for the architecture of an artificially intelligent brain, we could look to rights under patent law to grant a limited monopoly to the designer of the different hardware components comprising the artificially intelligent brain. For example, with integrated circuits, provided that their design displays satisfactory inventiveness and meets the required standard of uniqueness, patent protection is an option for the protection of the intellectual property rights embodied in an integrated circuit design. However, the lion’s share of integrated circuit designs is considered obvious under most patent systems given that they typically lack any improvement (inventive step) over their predecessors (prior art).⁶⁶ Further, integrated circuits are comprised of numerous building blocks, each “building block could potentially be patentable. However, since an integrated circuit contains hundreds or thousands of semiconductor devices, a patent claim to an integrated circuit would have to cover hundreds or thousands of individual elements- this would be like trying to write a patent on the neuronal circuits in the

⁶⁵Copyright v. Patent: A Primer on Copyright and Patent Protection for Software, at: <http://www.law.washington.edu/lta/swp/law/copyvpatent.html>.

⁶⁶*Id.*, note 62, NAIP PATENT BLOG.

brain.⁶⁷ Consequently, a patent claim that attempts to describe an entire integrated circuit may be hundreds of pages long. Clearly, such a narrow claim would provide almost no protection, and especially for an artificially intelligent brain consisting of billions of circuits.

Even if one sought such narrow protection, writing a patent application supporting a claim with thousands of elements would be extremely tedious and expensive.⁶⁸ As indicated by Rajkumar Dubey, writing for *Mondaq*, “Obviously, integrated circuits are not easily describable in a patent specification or the claims. Also, it may take several years to obtain an integrated circuit patent from most patent offices worldwide. This is unacceptable given that an integrated circuit’s useful commercial life may be less than 1 year.”⁶⁹ What if the same principle of obsolescence applied to the human brain such that every 1–2 years a person had to apply for patent protection of the neuronal circuitry of their brain? Or imagine that in the coming cyborg age the human brain is equipped with neuroprosthesis with billions of integrated circuits. That is, imagine the human brain becoming obsolete every 2 years or so due to the necessity of having to integrate (or update) new technology within the brain. The cumbersome, time-consuming nature of patent filing combined with extremely narrow protection would make patent law an insufficient form of protection for the brains neuroprosthetic devices and therefore the brain of an artificially intelligent machine.

Other forms of protection for intellectual property are also inapplicable to the integrated circuit layouts, which, will represent a major component of an artificially intelligent brain. Design patents protect the ornamental, but not the functional aspects of an article of manufacture described in its drawings. Since an integrated circuit layout is more functional than ornamental, design patent protection is generally inapplicable to integrated circuits. Finally, in many cases trade secret law cannot be used to protect most integrated circuits because an integrated circuit layout may be reverse-engineered. But what if an artificially intelligent brain is writing its own programs which are stored internally on its integrated circuits, and what if the programs have commercial value (that is, are trade secrets)? Once a program is stored on a tangible medium of expression it may still remain a trade secret but once communicated to the public trade secret protection is lost. However, since an artificially intelligent brain communicates in object code, and keeps the source code “locked in its mind,” it is simultaneously possible to maintain both trade secret and copyright protection for the program. Here I should point out that the reverse engineering of the human brain is one of the main techniques that some researchers are using to try to create an artificially intelligent brain.

Rajkumar Dubey writing about integrated circuits has commented that “The layout of transistors on the semiconductor integrated circuit, or topography of

⁶⁷*Id.*, note 62.

⁶⁸*Id.*, note 62.

⁶⁹Rajkumar Dubey, *Semiconductor Integrated Circuits Layout Design in Indian IP Regime*, 2004, at: <http://www.mondaq.com/india/x/28601/technology/Semiconductor+Integrated+Circuits+Layout+Design+In+Indian+IP+Regime>.

transistors on the integrated circuit, determines the size of the integrated circuit as well as its processing power.”⁷⁰ He states “That is why the layout design of transistors constitutes such an important and unique form of intellectual property fundamentally different from other forms of intellectual property like copyrights, trademarks, patents and industrial designs” and therefore in my view is of interest to a law of artificially intelligent brains.⁷¹ Given that patent, copyright, and trade secret law cannot adequately protect integrated circuit design, an exclusive protection for semiconductor integrated circuits layout-design has become necessary to the semiconductor industry. This level of protection represents a body of law that has significance for our cyborg future. So what protection may be available for the hardware of an artificially intelligent brain? In 1984 the U.S. passed the *Semiconductor Chip Protection Act* which provides statutory protection for integrated circuit design rights.⁷² Although codified under the same title as Copyrights, the Act is clearly intended to provide integrated circuit designs with *sui generis* (“of its own kind”) rights. It has some aspects of copyright law, some aspects of patent law, and in some ways it is completely different from either.

Providing legal protection for the physical components comprising the architecture of an artificially intelligent brain will also form a part of an emerging law of cyborgs, and is similar to the idea of protecting “bodily integrity” for humans.⁷³ Semiconductor chips are massed produced from multi-layered three-dimensional templates that are called “chip masks” in the trade, and “mask works” under the Act. The main purpose of the *Semiconductor Chip Protection Act* is to prohibit “chip piracy”—the unauthorized copying and distribution of semiconductor chip products copied from the original creators of such works. But the Act could also provide protection for the architecture of an artificially intelligent brain given that it is constructed with integrated circuits.

According to the Act, just like with copyright, integrated circuit design rights exist when they are created; this is unlike patents which confer rights after application, examination, and issuance of the patent. However, the exclusive rights afforded to the owners of integrated circuit designs are more restricted than those afforded to both copyright and patent holders. Modification (derivative works), for example, is not an exclusive right for owners of integrated circuit designs (this has implications for mind uploads, see Chap. 7: *The Law of Looks and Artificial Bodies*). Furthermore, the exclusive right granted to a patentee to “use” an invention, cannot be used to exclude an independently produced identical integrated circuit design.⁷⁴ Thus, reproduction for reverse engineering of an integrated circuit design is specifically permitted by most jurisdictions.

⁷⁰*Id.*

⁷¹*Id.*

⁷²Semiconductor Chip Protection Act, 17 U.S.C. §§ 901–914.

⁷³Gowri Ramachandran, 2009, *Against the Right to Bodily Integrity: Of Cyborgs and Human Rights*, 87 Denver University Law Review, No. 1, p1.

⁷⁴*Id.*, notes 64, 69.

Given the importance of protecting integrated circuits from piracy, several nations, including Japan and the European Community have followed the example set in the U.S. and endorsed their own similar statutes/directives recognizing and protecting integrated circuit designs (also referred to as the “topography of semiconductor chips”). And in 1989, a Diplomatic Conference among various nations was held, at which the *Treaty on Intellectual Property in Respect of Integrated Circuits* (IPIC Treaty) was adopted internationally.⁷⁵ This treaty has been partially incorporated into the TRIPS agreement of the World Trade Organization (WTO).⁷⁶ I believe that the potential threat that artificial intelligence could pose to humanity is serious enough, that just as with the semiconductor industry, international law should be crafted to create a common response to the potential threat that artificial intelligence could pose to society.

Further, other issues of law dealing with computer chips are also applicable to an artificially intelligent brain. For example, an important consideration for protecting the brain of an artificially intelligent machine concerns its memories and how they are stored and loaded to different devices. Memory chips such as an EPROM chip (erasable programmable read only memory), are chips that retains its data when its power supply is switched off. EPROM chip topographies are protectable under the *Semiconductor Chip Protection Act*, but such protection does not extend to the information stored on the chips, such as computer programs.⁷⁷ Such information is protected, to the extent that it is, by copyright law applicable to software which was discussed earlier.⁷⁸ Interestingly, the Court of Appeals for the Ninth Circuit in *MAI Systems Corp. v. Peak Computer, Inc.*, held that loading software into a computer’s random access memory (RAM) created a “copy” and a potentially infringing “reproduction” under the Copyright Act.⁷⁹ What that holding meant is that even if no hardcopy was made, temporally storing a program in RAM was a reproduction and potentially infringing act. So turning on a computer constitutes a reproduction of the operating systems programs because they are automatically stored in RAM whenever the computer is activated, or for that matter whenever a file is transferred from one computer network user to another. The *MAI* court held that the program temporarily stored in RAM represents a reproduction, although the U.S. Congress subsequently enacted an amendment to the Copyright Act to specifically carve out exceptions to this court decision in several circumstances.⁸⁰

⁷⁵TRIPS-The areas of intellectual property that it covers are: copyright and related rights (i.e. the rights of performers, producers of sound recordings and broadcasting organizations); trademarks including service marks; geographical indications\ including appellations of origin; industrial design; patents including the protection of new varieties of plants; the layout-designs of integrated circuits; and undisclosed information including trade secrets and test data.

⁷⁶Overview: the TRIPS Agreement, World Trade organization, at: https://www.wto.org/english/tratop_e/trips_e/intel2_e.htm.

⁷⁷Semiconductor Chip Protection Act of 1984, Wikipedia, at: https://en.wikipedia.org/wiki/Semiconductor_Chip_Protection_Act_of_1984.

⁷⁸See generally, *id.*

⁷⁹*MAI Systems Corp. v. Peak Computer, Inc.*, 991 F.2d 511 (9th Cir. 1993).

⁸⁰See, 17 U.S. Code § 117—Limitations on exclusive rights: Computer programs.

Our Competition Against Better Brains

Based on the above discussion, if artificially intelligent brains continue to get faster under Moore's law, and their brain architecture more sophisticated, there might conceivably come a point-in-time when artificial intelligence is capable of performance comparable to that of human intelligence. From that point on, artificially intelligent computers would not stop the process of getting smarter, but instead would accelerate the process of acquiring knowledge and connecting to the world through the Internet. In fact, in the last several decades the steady trend has been for computers to get faster, have greater memory capacity, and be networked to each other and to the emerging *Internet of Things*. And while not increasing exponentially (according to some authors), developments in artificial intelligence and knowledge in brain science is still rapidly increasing. Therefore, I believe it is just a matter of time before artificially intelligent machines claim to be self-aware and argue for rights. Based on that observation, the more we develop laws and policies which relate to the functioning, software, algorithms, and architecture of an artificially intelligent brain, the more we may be able to control our own destiny and shape the future as we approach the Singularity. Ray Kurzweil has convincingly made the point that once human levels of artificial intelligence is reached, artificially intelligent brains will then keep developing based on exponentially improving technologies until they are far more intelligent and capable than humans.⁸¹ The rate of development of artificially intelligent machines in terms of physical design will also show improvement because they will take change of their own development from their slower-thinking and less intelligent human creators. That is, when the Singularity is reached, and then surpassed, it is thought that artificial intelligence will work incredibly quickly at improving itself. What will a law of artificially intelligent brains be then?

According to James Barrat, after the Singularity, it's impossible to predict with certainty the behavior of these smarter-than-human intelligences with whom we might one day share the planet or that we might one day merge with (through a steady process, not all at once like a step function).⁸² But by merging with artificially intelligent machines, we may become super-intelligent cyborgs (or some other to-be-determined entity), using computers to extend our intellectual abilities (see Chap. 7 on *The Law of Looks and Artificial Bodies*). If we don't merge with the technology we are creating, and remain the biological product of evolution, maybe artificial intelligence will be benevolent and help us treat the effects of old age, prolong our life spans, and "fix" poverty and other forms of human suffering. But, in contrast, maybe artificially intelligent machines will turn on humanity and attempt to exterminate us or to control us in undesirable ways. But if we do merge with our artificially intelligent progeny we may avoid extinction and we may begin

⁸¹Ray Kurzweil, *supra*, note 27.

⁸²James Barat, 2015, *Our Final Invention: Artificial Intelligence and the End of the Human*, St. Martin's Griffin.

the process that will transform our species into something that is no longer recognizable as such to humanity.⁸³ This transformation has a name: Posthumanism and is a development discussed throughout this book. Of course many oppose the idea that humanity could someday transform into something “new”; and for those strongly opposed to the Singularity, now is the time to mount opposition. Clearly, artificial intelligence as it exists today doesn’t produce the kind of intelligence we associate with humans and clearly we are still human and not machine entities.⁸⁴ However, after we are more-and-more enhanced with cybernetic technology we will blur the line between human and machine, and as we move towards late century or early next century, we may have completely transformed to become the technology.

The artificial intelligence of today tends to be able to master only one highly specific domain, like interpreting search queries or playing chess. They typically operate within an extremely specific frame of reference and lack common sense. They’re intelligent, but only if you define intelligence in a narrow and limited way. The kind of intelligence Ray Kurzweil is talking about when he describes future artificially intelligent beings, which is called strong artificial intelligence doesn’t exist yet. Why not? Obviously we’re still waiting for the exponentially growing computing power made possible by Moore’s law to be combined with advances in algorithms, knowledge learned from neuroscience about the circuits of the brain, and improving architectures of artificially intelligent brains. But as Lev Grossman writing in Time magazine states- “it’s also possible that there are things going on in our brains that can’t be duplicated electronically no matter how many MIPS we throw at them.”⁸⁵ Grossman further says that “the neurochemical architecture that generates the ephemeral chaos we know as human consciousness may just be too complex and analog to replicate in digital silicon.”⁸⁶ Further, the biologist Dennis Bray is a voice of caution about the desirability of the cyborg future stating- “Although biological components act in ways that are comparable to those in electronic circuits,” he argued, in a talk titled “*What Cells Can Do That Robots Can’t*,” “they are set apart by the huge number of different states they can adopt.”⁸⁷ Multiple biochemical processes create chemical modifications of protein molecules, further diversified by association with distinct structures at defined locations of a cell.⁸⁸ Bray points out that the “resulting combinatorial explosion of states endows living systems with an almost infinite capacity to store information

⁸³Lev Grossman, 2011, 2045: The Year Man Becomes Immortal, quoting Ray Kurzweil, at: <http://content.time.com/time/magazine/article/0,9171,2048299-4,00.html>.

⁸⁴*Id.*

⁸⁵*Id.*

⁸⁶Miguel Nicolelis, *id.*, note 6.

⁸⁷Dennis Bray, What Cells Can Do that Robots Can’t, Youtube video, at: <https://vimeo.com/18143991>.

⁸⁸*Id.*

regarding past and present conditions and a unique capacity to prepare for future events.”⁸⁹ The complexity of biology makes the binary language that computers use to manipulate data look crude and it remains to be seen whether digital technology can simulate a brain.

As Grossman notes “Kurzweil admits that there’s a fundamental level of risk associated with the Singularity that’s impossible to refine away, simply because we don’t know what a highly advanced artificial intelligence, finding itself a newly created inhabitant of the planet Earth, would choose to do.”⁹⁰ It might feel like competing with us for resources, then again, it might not, but if it does eventually we will lose. If the Singularity is coming, these questions will have to be addressed whether we like it or not, and Kurzweil thinks that trying to put off the Singularity by banning technologies is not only impossible but also unethical and probably dangerous for humanity.⁹¹ Kurzweil argues that “It would require a totalitarian system to implement such a ban,”⁹² continuing, he states “It wouldn’t work. It would just drive these technologies underground, where the responsible scientists who we’re counting on to create the defenses would not have easy access to the tools.”⁹³

Kurzweil does not see any fundamental difference between flesh and silicon that would prevent the latter from human-like thinking. However, the law does distinguish between neurons and integrated circuits; primarily because one can own circuits, but not another person’s neurons. Kurzweil defies biologists to come up with a neurological mechanism that could not be modeled or at least matched in power and flexibility by software running on a computer.⁹⁴ If Kurzweil is correct, an artificially intelligent entity arguing for rights is an eventuality, therefore humanity would be wise to establish a regulatory scheme to protect humanity and to ensure all intelligent beings that emerge to join society have basic rights.

To summarize, artificially intelligent brains are improving rapidly based on exponentially accelerating technologies. They may match humans in general intelligence by midcentury, therefore the emerging law of cyborgs, and particularly the laws discussed in this chapter which relate to an artificially intelligent brain, could provide important protections not only for the future rights of artificially intelligent beings, but of humans either merging with them, or living amongst them as less intelligent beings relying on their sense of fairness to treat humanity with respect and justice.

⁸⁹*Id.*

⁹⁰Lev Grossman, *id.*, note 83.

⁹¹Lev Grossman, *id.*, note 83.

⁹²Lev Grossman, *id.*, note 83, quoting Ray Kurzweil.

⁹³Lev Grossman, *id.*, note 83, quoting Ray Kurzweil.

⁹⁴Singularitarianism, posted by PZ Meyers, 2011, quoting Ray Kurzweil, at: <http://scienceblogs.com/pharyngula/2011/02/13/singularitarianism/>.