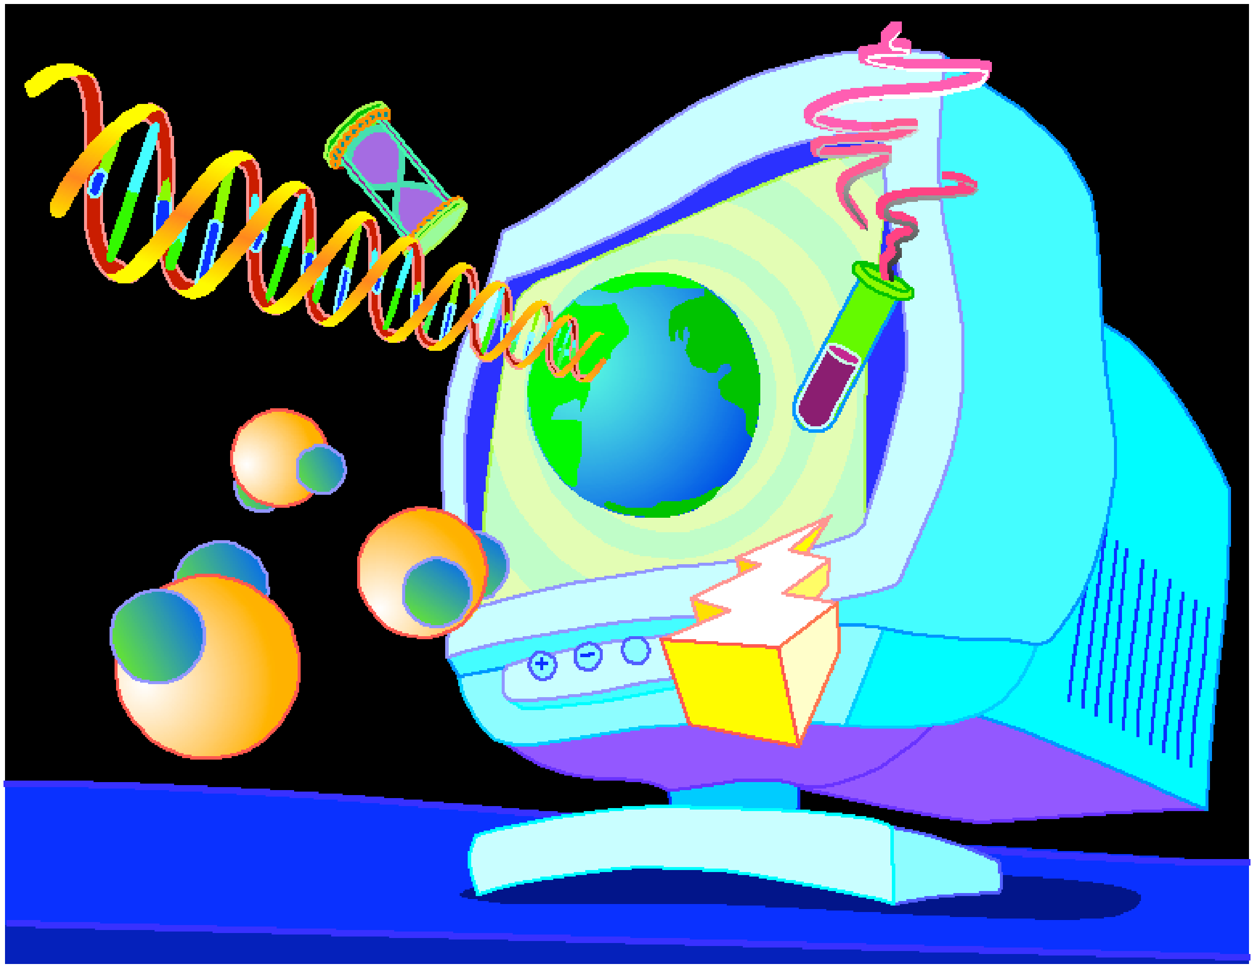

Sentient Machines and Legal Rights

Rich Seifert
Santa Clara University
Legal Issues of the 21st Century
Prof. David Friedman
Spring 2004

I. Introduction

To date, only one species on Earth has had:

- The ability to significantly control its environment and the effect of other species on its ability to survive,
- Communications capabilities sufficient to allow complex interorganism interaction, and
- A societal organization sufficiently complex as to require explicit legal rules to maintain order and efficiency.

Slow evolution has been a blessing to mankind. Humanoids have existed that possessed a wide range of intelligence, physical abilities and communications capabilities, from Homo habilis, through Homo erectus, Homo neanderthalis, and ultimately, Homo sapiens. [1] However, these species developed over a period of roughly 2 million years. None of them ever had to operate under conditions that required multiple distinct sets of sentient, communicating beings to develop legal rules that transcended species or intelligence levels. By the time legal systems were developed, there was only one sentient, socially complex species left—us.

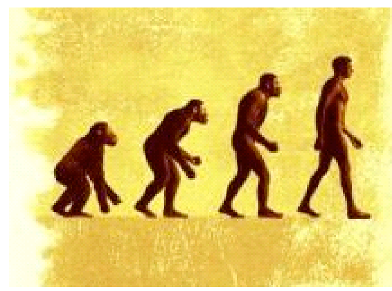

This situation may change in the relatively near future. Machine intelligence is advancing rapidly enough so that we can envision a time when it will approach, equal, and perhaps surpass that of human intelligence. At that time, we could have two sentient, highly intelligent species coexisting on the planet. Most modern legal systems assume (or expressly grant) a set of fundamental rights to all humans, regardless of intra-species variations. How will those systems adapt to accommodate machine intelligence?

This paper looks at two facets of this (admittedly huge) question:

(1) Will current technology extrapolate such that electronic computing machines can approach human levels of intelligence, or are there some fundamental limits that will prevent this from occurring, absent some unpredictable scientific breakthrough?

(2) How does our current legal system allocate rights to members of less intelligent species? As computer intelligence progressively attains the intelligence levels of existing, lower species, we might expect to grant rights to those machines in much the same way as we grant rights to other species.

II. The Emergence of Computer Intelligence

The first electronic computers were developed primarily as military aids in the second World War; fast, accurate calculation was the key to determining the proper settings for effective artillery fire. [2] While the devices employed were primitive (at least by today's standards), there was already a recognition that this was just the beginning of a trend. Machines could already perform simple calculations much faster than humans. [3] Prescient scientists understood that complex "thinking" processes might be

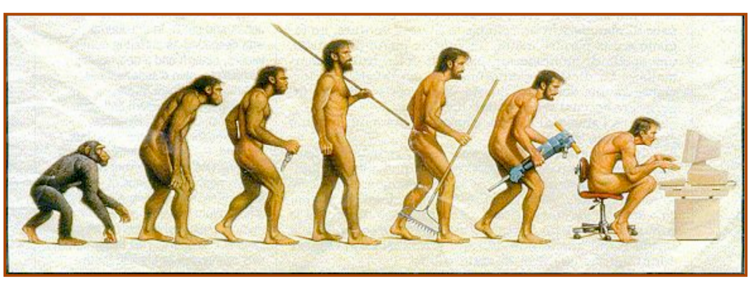

emulated by huge sequences of simple calculations; if these calculations could be performed fast enough, the result could appear as "intelligence" to an objective observer. As far back as 1950, Alan Turing was positing methods of measuring computer intelligence, in particular, a test that indicated whether machine intelligence was equal to that of humans. [4]

As machines became faster, and the tasks they performed more closely resembled those traditionally handled by humans, the idea of computer intelligence moved closer to reality. We could see the steady increase in computing power, and naturally extrapolated that increase to a time when computers might achieve, and perhaps surpass, human intelligence, as depicted below:

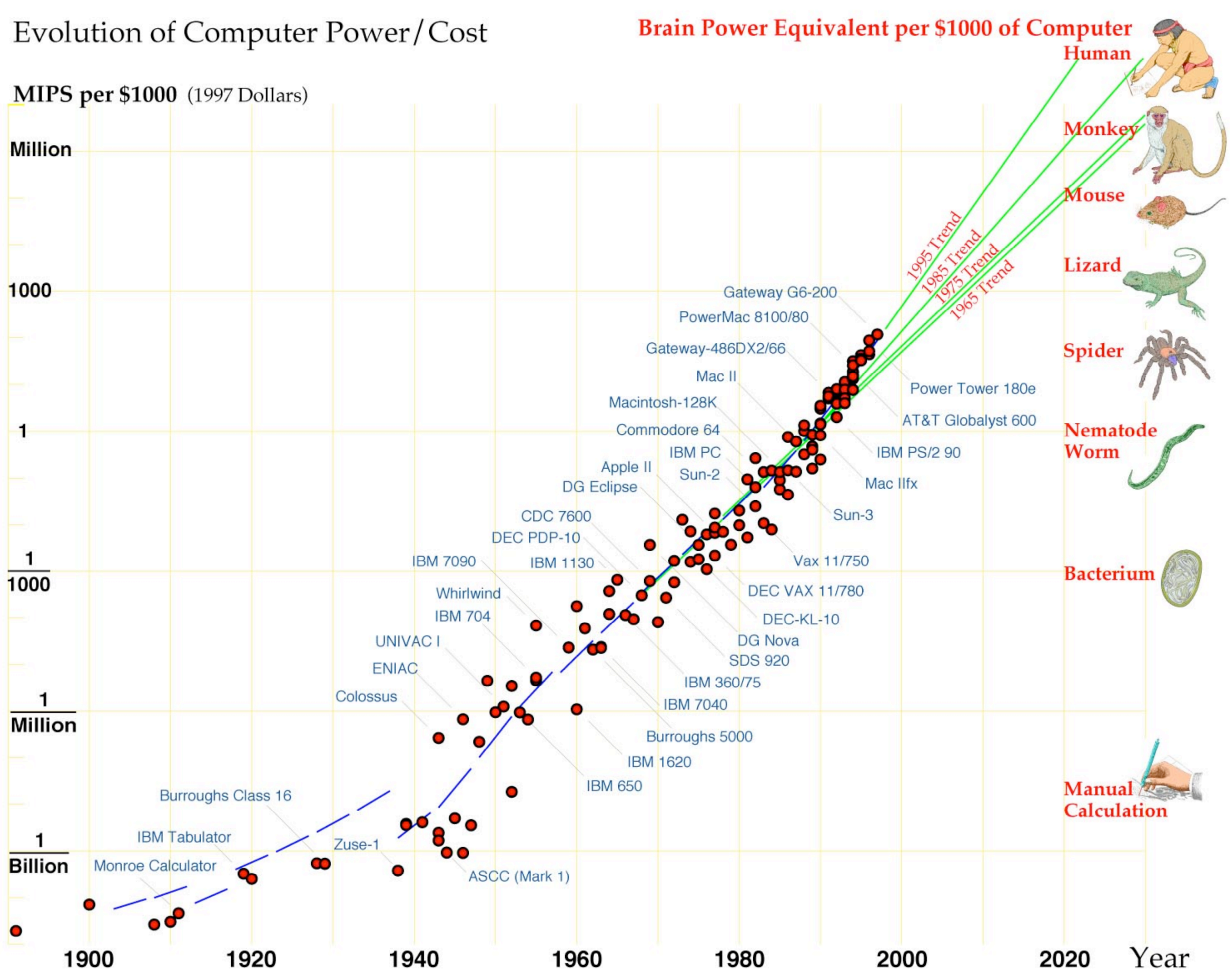

The Evolution of Computing Power [5]

II.A Moore's Law

The driving force behind the rapid increase in electronic computing power is a principle known as "Moore's Law"; an observation that the density with which we can pack transistors onto a silicon die doubles approximately every two years. [6] [7]

Gordon Moore

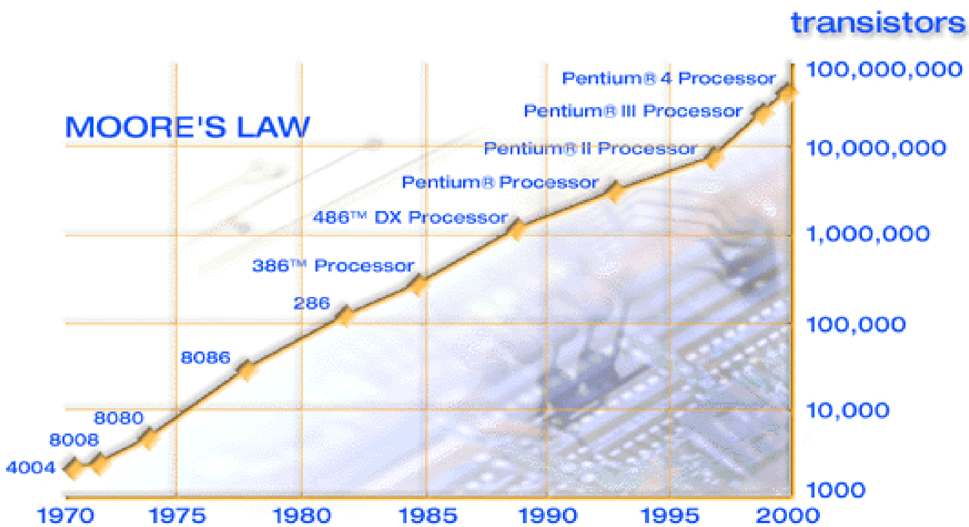

Transistor Density of Intel-family Processors

The cost of a semiconductor device is primarily determined by the die area it occupies (its "floorspace"). Thus, at any given cost point, we can double the electronic functionality of a device every two years. (Or alternatively, for a given level of functionality, we can reduce the cost by a factor of two every two years.) It is this relationship that has fostered the creation of a wide range of inexpensive consumer electronic devices, and pushed computing power closer to the point of comparability with human brain power.

II.B An Intelligence Metric

In looking forward to a time when computer intelligence may equal or exceed that of humans, we tend to ignore the fact that we really don't know what intelligence *is*. Kurzweil defines intelligence as "the ability to use optimally limited resources … to achieve … goals." [8] While that may be as good a definition as any (and there are as many definitions as there are people who have considered the problem!) it does not provide much guidance for measuring or comparing differing forms of intelligence, particularly when the "goals" being achieved are fundamentally distinct. If we ever get to a time when comparable computer and human intelligences can be unleashed on the same set of problems, we could then make direct comparisons between the two: (a) measure the time required by each to achieve the goal, (b) measure the energy required to achieve the goal, etc.

In today's world, computers are so far below humans (by any measure of intelligence) that such direct comparisons are not possible. It is simply not feasible for a modern computer to perform most of the simple intellectual tasks that humans perform thousands of times each day. Certain narrow, specialized tasks can be accomplished (e.g., chess playing and limited facial recognition), but even these generally require non-standard hardware, and huge investments in software development. Conversely, it is unfair to compare human intelligence to computers for the tasks that today's computers can perform easily: rapid numerical calculations, searches through huge volumes of data, etc. The differences lie not only in intellectual power, but in the internal processing mechanisms and architecture.

We believe (putting religious spiritualism aside) that human intelligence is primarily a result of processes occurring within our brains, but we do not fully understand which characteristics of the brain give rise to intelligence. In particular, while we understand that creatures with larger brains have higher intelligence [9], we are

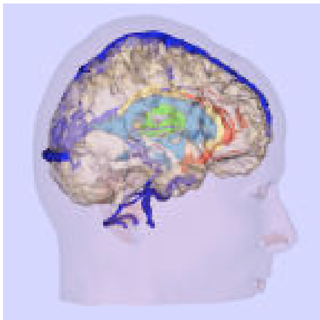

not sure which components of the brain comprise the strong functions, i.e., the factors that have the greatest effect on intelligence.

The development of electronic computing has provided us with a reference model. We may not know much about brains, but it is well understood that the capability of a current-day computer is a function of:[10]

- *Computational speed*: The processor cycle time, generally measured in MHz or GHz, that determines the number of instructions that can be executed each second.
- *Instruction set*: The amount of computational work that can be performed during each processor cycle.[11]
- *Internal architecture*: The communications paths available among subsystems (processor, memory, I/O, etc.), along with their transfer rates and efficiencies.
- *Memory capacity*: The amount of storage available for data and instructions.

Most commentators use some combination or variation of the above factors when trying to measure human intelligence, reasoning that the human brain must operate something like an electronic computer. Of course, when we then use this data to compare human intelligence to computer intelligence, we have wrapped ourselves in a definitional circle. We use metrics appropriate for electronic computing to measure and compare electronic computers to human thinking, while not being entirely sure that human thinking follows the same architecture or logic of electronic computers. We then extrapolate this data to the future to predict when computer intelligence will approach that of human intelligence, often forgetting that our measuring stick is calibrated for computers, not for humans. Nonetheless, there is little else available until such time as we have a better understanding of the internal workings of the human brain.

Kurzweil follows this approach, and concludes that: [12]

(1) The human brain has a very slow processor cycle time (5 ms), due primarily to the slow rate of signal propagation through nerve fibers (as compared to the copper conductors used in most computers and integrated circuits).

(2) The brain's "processors" are simple (and slow), but there are ~100×10^9 of them operating in parallel, resulting in a total of 20×10^{12} executed instructions per second.

(3) The memory capacity of the human brain is ~100×10^{15} bits (12.5 million gigabytes, although there is no reason for the brain to organize data in groups of eight bits).

Current, moderately-priced computers can execute ~500×10^6 instructions per second (1:40,000 compared to Kurzweil's brain), and have a memory capacity of ~2×10^9 bits (1:50,000,000 compared to Kurzweil's brain). [13] Kurzweil then applies Moore's law to his data and estimates that computer intelligence (according to these metrics) will equal human intelligence sometime around 2020-25.

II.C There's a Fly in my Soup, er, Ointment

There are a few problems with Kurzweil's (or any) simple extrapolation. Most of them derive from a set of (questionable) underlying assumptions:

(1) The human brain is directly analogous to an electronic computer,

(2) Metrics appropriate for electronic computing power can be rationally applied to the human brain, and

(3) Advances in electronic computing will inexorably lead such machines to achieve human-level intelligence.

In the sections below, I discuss some of the technology issues that challenge these assumptions.

II.C.1 Human Memory vs. Computer Memory

If I ask you to think about pizza, most of you will conjure up thoughts of melted cheese, your favorite toppings, and perhaps some event you associate with pizza. [14] You are able to retrieve these related thoughts without knowing *where* in your memory they were stored. The keyword "pizza" allows you to access this associated data without having to know their location in memory.

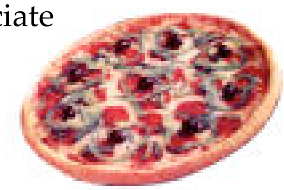

In traditional computer memories (e.g., RAM and ROM), one needs to know the address where information is stored in order to retrieve it. In an *associative memory*, one uses a stored data value as a key for retrieving data associated with it. Thus, rather than specifying a *location* in memory from which stored information is to be retrieved, the logic presents a data *value* to the memory system; the associative memory returns either the associated data or an indication that the requested data value (the key) does not exist anywhere in the memory. We never know, nor do we need to know, *where* in the memory the data value/key or associated data is stored; all we care about is whether the value we are looking for is present, and if any other data is associated with it. Similarly, when storing information in an associative memory, the logic presents the data value/key and the associated data, and the memory system stores it in *any available location*. Since we don't access the memory by specifying an address, the actual location where a given data element is stored is irrelevant. In this manner, an associative memory functions as a virtually-instantaneous hardware search-and-update function element.

Human memories are associative in nature. We do not need to know (nor do we ever consciously know) the location where information resides. Furthermore, stored information itself is the key to locating related pieces of information; there is no need to use the primary (triggering) information as an input to some sort of index that returns the location of related information. Indeed, there is no evidence that any such index exists within the brain. By contrast, conventional computer memories are location-dependent and not associative in nature. In order to find information related to some unit of data, we generally require a complete data index and extremely powerful search tools. [15]

Associative memories *have* been built for commercial use. [16] There are two basic types:

(1) *Content Addressable Memories* (CAMs) are truly associative hardware memory devices, generally implemented in integrated silicon. Rather than present a location value for memory retrieval, the system presents a data value (the key, of some specified length) and the CAM returns none, one, or more than one data value associated with the input. In essence, the hardware compares the data value with the contents of *every location within the CAM*, simultaneously, using either a massive array of combinational logic, or specialized equivalent structures (e.g., wired-OR logic).

As a result, CAMs tend to be quite expensive (4-5x the cost of a conventional RAM of the same size). In addition, CAMs have severely limited storage capability, due to their silicon-intensive design. Even worse, the die area required for a hardware CAM increases exponentially with its storage capacity. Beyond a relatively spartan memory store (at least, relative to human memory capacities), CAMs become impossible to build as a result of fundamental physical limitations in silicon technology (as opposed to engineering difficulty).

(2) *Pseudo-CAMs* are memory-subsystems that combine a conventional RAM with a high-speed dedicated processor configured as a "lookup engine". These devices mimic the associative characteristic by performing fast searches through conventionally-addressed memory. While they avoid the physical constraints of true CAMs, they pay for it with speed; pseudo-CAMs are many orders-of-magnitude slower than true CAMs. In addition, the cost of the processing subsystem must be added to the cost of the memory itself.

As a result, associative memories (of either type) are used only for specialized applications where their benefits outweigh their (huge) cost. [17]

The fundamental problem is thus a difference between the inherent design of computer memory vis-à-vis human memory. As much as we think we can "fake it", we don't know how to build a large associative memory. Furthermore, we don't know how the associative memory in the human brain is designed (other than the fact that it does *not* appear to be designed anything like our silicon equivalents). Any simple extrapolation of current-to-future memory capacity according to Moore's Law ignores the architectural differences; a huge location-addressable memory is simply not the same thing as a huge content-addressable memory. Perhaps we will learn how to build such a device, which would then allow technology projections to properly reflect the advances necessary to achieve human-like memory, but this would be new science, not incremental engineering progress.

II.C.2 Conventional Computing vs. Neural Computing

Kurzweil recognizes that human brains are organized as massively-parallel neural networks, whereas computer brains are organized as sequential finite-state machines. [18] Neural networks are quite effective for solving many of the problems that faced humans during our own evolution, e.g., rapid pattern-matching (for identifying family and other friendly types vs. predators), predicting trajectories of thrown or falling objects, etc. Sequential machines can solve these problems as well, but they do so in an entirely different manner; by rapidly executing a series of instructions, they calculate (using time as a substitute for parallelism) the desired result, following human-designed algorithms tailored to the specific problem at hand.

That is, neural networks "compute" in a manner entirely different from conventional computers. The neural net work does not need to formulate and compute

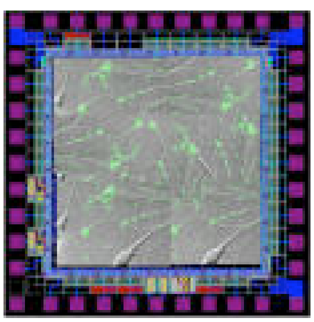

a result based on some underlying mathematical model; there is no evidence that a brain calculates the arc of a thrown rock from any knowledge of the physics of dynamic objects, the gravitational force of the earth, or the frictional drag of the atmosphere. The neural network *evolved* its connectivity as a results-oriented response to prior inputs; the human whose brain had the connections that best predicted the correct answer was more likely to survive and pass on that design to his progeny via his DNA. This has great implications for the problem of software, discussed below. [19]

Neural computers are apparently superior architectures on which to base human survival and intelligence; if it were otherwise we would not have evolved as we have. Any alternative approach that worked better would have dominated over neuralbrained humanoids. [20] Indeed, had we evolved in some other direction, we might have a totally different concept of "intelligence" today. Could it be that human intelligence (as we understand it) is dependent on, or simply an artifact of, neural-network processing? Are the concepts of self-awareness, sentience, independent thought, and perhaps, emotion, somehow related to the massive parallelism in a way that no single-threaded being could ever experience?

Granted, we can emulate neural processing by using a fast sequential-instruction computer, but it would require enormous speed. Events separated by as many as 100 billion instructions (the number of parallel processors estimated by Kurzweil) would have to occur so close in time as to seem simultaneous (i.e., within the 5 ms neuron firing time). [21] This does not appear possible, even in theory. In the ensuing 50 *femtoseconds* (50×10^{-15} seconds, the maximum instruction time if one is to execute 100 billion instructions in 5 ms) light (or any other signal) travels only 15 μm. That would have to be one small computer! The processor and all 100 million gigabits of associative memory must sit within a sphere having a 15 μm diameter. Since only ~60×10^{12} silicon atoms will fit within that volume, we are looking at some radically new technology.

Of course, we could dispense with *emulating* neural networks, and build a truly neural computer. There have been some efforts to create hardware that behaves more like a human neural network than a personal computer, but these remain confined to highly-specialized research experiments. Most of the money and brainpower being applied to computing technology today is directed towards building faster and cheaper conventional, sequential-instruction processors. Given that there is no current commercial demand for neural processors (nor any in the foreseeable future), there is no market pressure for manufacturers to invest in technology improvements in that space.

This leaves us with two possibilities:

(1) That we will continue to develop and improve conventional, sequential-instruction computers to the exclusion (for the most part) of neural processors. This will provide substantial increases in machine intelligence over the coming years, but with an upper limit far below that of human capabilities, assuming that human-like intelligence or sentience is even possible with such an architecture. (In addition, we must also consider the software problem, discussed below).

(2) At some point, there will be significant commercial demand for neural processing systems such that Moore's Law can be applied to that technology. From that point, we may be able to project out to a time when computer intelligence approaches or exceeds that of humans. The problem is that both the start date and the initial level of intelligence on the start date for that curve is unknown. While we believe that the shape of the curve should be the same as for any other application of Moore's Law, it is impossible to predict when the curve will surpass the human intelligence level without knowing the starting point. In all likelihood, the crossover point is farther out that 50 years; predictions become asymptotically meaningless at such extremes, due to the potential for interactions with everything else that is concurrently happening within society and technology.

II.C.3 The Program is all Written and Debugged; I Just Have to Type it In

As anyone who has worked with current computer technology knows, the real difficulties lie not in the power of the underlying hardware, but in the capabilities of the application software. The average office worker has a machine on her desk that is hundreds of times more powerful (in terms of the performance metrics discussed earlier) than the largest mainframe computers of 25 years ago. Yet, while human productivity has unquestionably increased over that period of time, the increase is nowhere as large as the performance improvement in the machines we use. The reason is that the performance metrics tell us the raw capability of the *platform*; the usefulness of the machine is determined by the software—the series of instructions that command the hardware to behave as one intends.

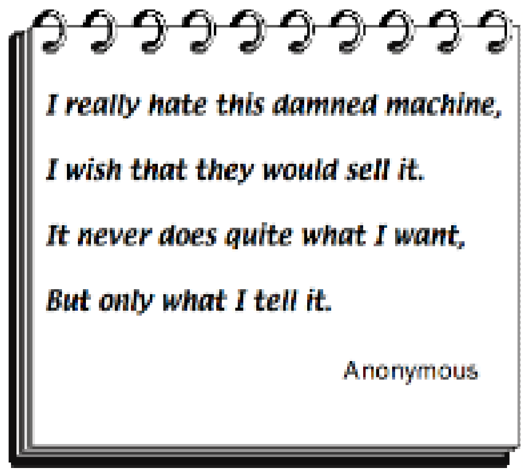

A sequential-instruction computer requires a sequence of predetermined instructions (a program) in order to perform any useful function. To date, virtually every such program has been written by a skilled human. Each time we envision and desire some new function (i.e., a new application) we must create a new program. Today, even relatively simple applications can take hundreds of man-hours to design and code; major projects regularly employ teams of dozens of programmers for periods of years.

While the hardware platform may follow Moore's Law of exponential growth, we have observed no corresponding increase in software performance. In fact, industry experience [22] shows that as we strive to create more powerful programs, the number of man-hours required increases more than the increase in program complexity; i.e., there appears to be a *negative* economy of scale for software development. [23] Thus, we may end up with a hardware platform whose capability equals that of a human brain (according to our metrics), but that will not exhibit any intelligence due to a lack of software to exploit the underlying power.

But who wrote the "software" in the human brain? [24] The answer is that the software wrote itself, and the key to understanding this conundrum lies in the neural

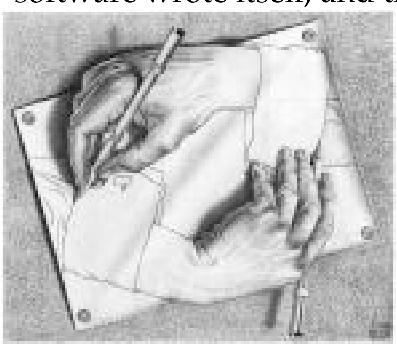

nature of the brain. The "software" (i.e., the control element that directs the hardware to perform the desired function) in a neural-network computer is reflected in the interconnections among the massively-parallel computing elements. Unlike a conventional computer, there is no "instruction store" (i.e., memory that contains the instructions to be performed). The neural computer:

- Receives a set of inputs,
- Distributes those inputs to its computing elements through an interconnection matrix,
- Allows each computing element to perform its (minimal) transformation on the input, and
- Recombines the transformed inputs via the interconnection matrix to produce a set of outputs.

Perhaps most important, the interconnection matrix can vary over time, as a result of the outputs produced by the neural computer. There is thus a closed-loop feedback system, where the entire computer can "reprogram" itself (by changing the neural interconnections) in response to its own previous computations. [25] The brain can learn new behaviors (write new programs) by observing prior input/output transformations and adjusting the interconnections to reinforce the desired output. Over hundreds of thousands of human generations, our neural computer (the brain) has adapted and reprogrammed itself as an evolutionary response. Any change in the interconnection matrix that increased the likelihood of survival for the owner of that brain would succeed to the next generation. Our brains therefore programmed themselves for survival. No one had to instruct the computer to run from a charging lion; the people whose brains provided the correct output ("Run like hell!") survived, while the ones whose brains provided an inappropriate output ("Nice kitty!") are gone.

We know of no way (today) to effect such self-reorganization in sequential-instruction computers, although there seems no fundamental reason why it could not be achieved. If, however, we shift to an electronic neural computing platform with a truly reconfigurable interconnection matrix, we could let the machine "learn" for itself and thus, write its own software. An electronic neural computer also has a great learning advantage over a biological one; we don't have to wait ~20 years between "generations" to see whether our changes improved or degraded the situation. We can accelerate evolution millions-fold by bypassing those nasty sex, gestation, and maturation processes.

But what would such a neural computer *learn*? Our brains evolved to where they are today as a result of a survival regimen. [26] We didn't look at the results of our computations and make value judgments as to which outputs were "better" than the others; the better ones were simply the ones that survived. Thus, any intelligence characteristics that we possess beyond those that improved our survival chances in the Stone Age (assuming that we do have capabilities beyond those needed solely for survival), are simply unplanned artifacts of evolution. What metric would we use to evaluate the worth of the output state of an electronic neural computer?

There is no comparable issue of "survival", in the same sense as that experienced by evolutionary man. Man had to deal with an uncontrollable natural environment, numerous physical hazards, and a need to constantly obtain food and avoid predators. We could let the electronic computer evolve towards its own "survival", but that would seem to place a value only on the sort of actions that would improve the likelihood that its power source would not be disabled, its components not be damaged, etc. It is not at all clear that the resulting neural interconnections would exhibit intelligence in the sense that we understand it, i.e., the ability to solve complex, abstract problems. It seems that the reason that humans developed such capability was because, over the evolutionary millennia, we had to deal with a much more varied environment. We developed a *general-purpose* problem-solving capability because the range of problems we needed to solve (to survive) was too large to deal with on a case-by-case basis.

One reason for the difference in survival environments is that humans can move over relatively large distances. Our mobility increases both our ability to obtain food, as well as our exposure to natural hazards and predators. If our existence was restricted to staying within a short distance of where we were born, we might not have had to develop a general-purpose problem-solving tool like the brain. To the extent that an electronic computer is exposed to limited environmental variation, it may not develop human levels of intelligence, if its learning algorithm is based on survival only.

Alternatively, we could impose some qualifier other than survival to evaluate the outputs of the neural matrix; we could say that positive outputs (i.e., ones worth learning) are not just those that allow the machine to live for another generation, but instead have some other desirable characteristic. However, this would necessarily impose an external value judgment (a human one) on the development of machine intelligence. We could train a new species of intelligent beings, where intelligence was defined as exhibiting characteristics that we deem desirable. Of course, this paradigm can be employed to create either a benign species that poses no risk to human survival (if "positive" traits—in human terms—are deemed desirable outputs), or a dangerous, predatory species that threatens some subset (or perhaps all) of mankind (if "negative" traits are judged desirable outputs, perhaps by a future megalomaniac).

In either case, this would represent a radical departure from evolution. Evolution does not make any value judgments with regard to the desirability of a species' behavior. There is no master plan driving the evolutionary chain, no notion of "better" creatures, or "progress", other than through sheer survival. If we impose our values on the desirability of an electronic neural computer's output state, we effectively become gods to that species, shaping their evolution according to our own wishes. If instead we allow pure evolutionary forces to control their destiny, their intelligence may never equal that of humans to the extent that they do not experience the range of environmental variation experienced by humans over millions of years.

There is no way to predict which of these (or some other) scenarios will come to pass. That said, it is still fun (if only marginally useful) to consider the legal rights implications of a sentient, intelligent electronic computer.

III. The Legal Status of Non-Human Intelligence Today

Over the coming decades, we may be interacting with a set of devices whose intelligence will be changing noticeably over fairly short periods of time. Relative to the power of the human mind, today's machines possess miniscule capabilities; however, if the projections are to be believed, these machines will approach, equal, and ultimately surpass human capability, at least by our objective metrics.

In approaching the problem of how human society will allocate legal "rights" to these machines, it is useful to see how human law has dealt with other forms of intelligence (both human and non-human) in the past. This should give us some insight into the driving forces and rationales for ascribing rights that seem implicit within our legal system. In this regard, there are four interesting intelligence relationships that analogize to a future human-machine environment:

(1) **Human Intelligence >> Other Intelligence**

This is the situation today with respect to human/computer interaction. Human intelligence is much greater (by orders of magnitude) than computer intelligence. We can analogize the legal ramifications by considering the rights of non-sentient, non- or minimally-intelligent living species on earth today, such as plants, insects, birds, and small mammals.

(2) **Human Intelligence > Other Intelligence**

At some point, the computer will emerge as a sentient being, although with capabilities still distinctly below those of humans. We can consider this case by analogy to how our legal system addresses the rights of "almost human" species, particularly chimpanzees and the higher primates.

(3) **Human Intelligence ≈ Other Intelligence**

Not long afterwards, computer intelligence will roughly equal that of humans, at least by our objective metrics. As history shows, however, simply possessing equal innate intelligence has not always been sufficient to justify the granting of rights under our legal system. We can draw parallels to the development of legal rights for blacks, women, and ethnic groups over the past few centuries.

(4) **Other Intelligence > Human Intelligence (or >> Human Intelligence)**

Finally, the computer may surpass humans in raw intellectual power. Obviously, there is no historical precedent for this case. However, we can consider it in either of two ways: (a) It is simply the same as the first two cases above, with the roles reversed, or (b) We can analogize to how humans (of equal intelligence) have interacted when one group had significantly more *physical* power than the other, e.g. a conqueror invading a primitive culture.

I examine each of these scenarios in the sections that follow.

III.A When Humans Dominate

Simply put, the non-sentient living species on earth today (plants, insects, birds, small animals, etc.) have no legal rights. Indeed, we deal with such items not as peers or quasi-peers, but as *property*. All forms of plant life and most small animals can be owned, and bought or sold at will. "The right to take their life, and to make property of them, includes all other rights; so that the common law recognizes as indictable no wrong and punishes no act of cruelty, which they may suffer, however wanton or unnecessary." [27]

Virtually all small animals (and many larger ones) can be killed at will, and eaten for food. [28] Indeed, some species are considered to be nuisances (e.g., plant pests, disease-bearing insects and rodents); killing these species is considered a *positive* action with respect to human society. [29]

To the extent that these creatures have any legal rights at all, they tend to be related to requirements for "humane" killing when raised for food, and laws prohibiting animal cruelty. [30] One can hardly imagine that ducks swell with pride over the knowledge that they can't legally be tortured before being killed and tea-smoked in a Chinese kitchen! [31] Furthermore, even such minimal protections apply only to higher forms of life; except for environmental protection statutes (discussed below) there are no legal sanctions for human actions taken against any plants or lower life forms.

As noted, we do enforce a number of restrictions against deforestation, and to protect certain endangered species. [32] However, these statutes are less about granting "rights" to the protected species, and more about protecting the environment in which humans live from unwanted destruction by the humans themselves. These laws impose a cost for what would otherwise be an unaccounted-for externality; e.g., allowing a property owner to clear-cut redwood trees would grant the owner a financial benefit that did not account for the societal cost of lost erosion protection (to property owners downhill) and reduced fire resistance (a redwood forest is virtually fireproof). That is, we protect these lower life forms *for our own benefit*, rather than from some altruistic motivation or any concept of "natural rights". This is an important theme that I believe will control any movement towards granting legal status and rights to machines; we will do so if and when it is to our (human) benefit to do so, not because we believe that intelligence per se demands that we confer such rights.

The situation today as regards computers is exactly comparable to that of our relation to plants and low life forms. We can treat our machines as we please, with virtually no legal implications. [33] One notable exception arises from recent regulations prohibiting the discarding of computer monitors (CRTs); within California, these devices are now treated as a form of hazardous waste, subject to controls on recycling and their presence in landfills. [34] This is precisely analogous to the sort of environmental protection we impose on trees and other "lowlife"; we invoke restrictions to account for externalities that impose a cost to human society.

III.B When Humans are Just a Bit Smarter

There are some animals on earth that are nearly as intelligent as humans. The higher primates, in particular chimpanzees and bonobos, are close genetic relatives of homo sapiens, and exhibit many of the characteristics that we associate with human intelligence. Our raw DNA is more than 98.3% identical; of the DNA that we believe actually controls physical and behavioral characteristics, humans and chimpanzees differ by only a few hundred genes out of ~100,000 (>99.5% identical). [35] This is not surprising, since we are both descendents of a common ancestor; our evolutionary paths separated only about 5-6 million years ago. [36]

Bonobo (Pygmy Chimpanzee)

As a result, our brains are identical in structure and similar in size. [37] We all appear to possess and exhibit the behaviors and characteristics we associate with intelligence: sentience, the ability to communicate, the ability to make and use tools to solve problems of survival, etc. The only difference is one of degree. The evolutionary branch on which humans reside endowed us with larger brains; we have more of those parallel neural processors, so we can solve more complex problems per unit of time.

III.B.1 The Brain is the Same, but the Legal Status is Not

Although the difference in brainpower between humans and other primates is miniscule, the legal difference between the species are huge. Chimpanzees are subject to property law; they can be bought and sold, owned, and generally treated and disposed of however the owner wishes (subject, perhaps, to animal cruelty statutes, discussed above). Humans have legal personhood; bonobos do not. As non-persons, they do not benefit from the protections to life, liberty, and property provided by the Constitution. They can be imprisoned without due process of law, forced into labor, and summarily killed (albeit humanely) with or without cause. "[A]nimals are treated as the property of their owners, rather than entities with their own legal rights." [38]

It doesn't matter how close in intelligence a species is to humans; there is no correlation between intelligence and legal status. In fact, a creature can have intelligence *superior* to some humans and still be denied personhood. An adult chimpanzee has a intelligence level (as demonstrated by comparative testing) approximately equal to that of a normal five-year old human. [39] Many humans lack this level of demonstrable intelligence; Newborn infants, comatose patients, brain-damaged persons (including both those born with severe brain deficiency and those with brains damaged due to injury or illness), people suffering from Alzheimer's disease, etc., all display intelligence far below that of an adult chimpanzee, yet our legal system grants them rights equal to those of normal humans while denying all legal status to the chimp. Clearly, our system does not allocate legal rights according to the sentience or intelligence of the beings involved.

III.B.2 Social and Physical Differences Matter More

"All law was established for men's sake."[40] The real reason that we don't grant legal status and rights to species possessing nearly equal intelligence as ours is simply

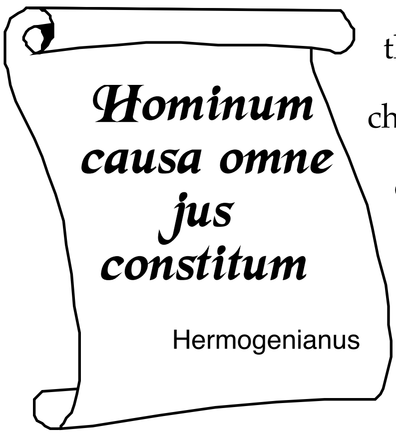

that we don't have to. No other species, including chimpanzees, pose any real threat to human survival or to the dominance of our society over theirs. Bonobos are among the most advanced species on Earth, but the ever-so-slight advantage that we hold makes all the difference in the world. Apes have been singularly unsuccessful in waging concerted war against humans (although they occasionally do well in a one-on-one contest).

Our legal system is designed to allow persons under its protection to ignore the wishes and needs of anyone (or anything) not so protected, to the maximum extent possible. Furthermore, we grant legal status to formerly unprotected persons only when it benefits those of us already protected. Thus, we grant legal rights to ourselves, to persons we consider physical or social threats, to persons with whom we wish to trade (to increase overall wealth, including our own), etc. We grant rights to ill, brain-damaged or comatose humans because we recognize that any of us could potentially find ourselves in that situation, and we want to preserve our own legal rights in that event.

There is no reason to believe that the development of machine intelligence will change this fundamental precept of our legal system. The machine may possess intelligence equal to, or greater than our own (as some primates do); however, it will still be relegated to the domain of property unless and until the machines pose some threat, or offer some benefit to humans that justifies their endowment with legal status.

III.C We're All Equal, but Some are More Equal than Others

While we (as a species) have never had to deal with another species of equal intelligence, we have often encountered the case where one group of humans has legal status and its associated protected rights, and another group does not. Until 1868, all U.S. constitutional protections could be (and were, in many states) denied to blacks; until 1920 women were denied the right to vote. [41]

Even after the constitutional barriers were lifted in the 19th century, it was a long time before blacks achieved equality in any practical sense. [42] A policy of government-sanctioned separatism and segregation was formally acknowledged in the landmark decision in *Plessy v. Ferguson* when the Court expressly stated that the 14th Amendment "could not have been intended to abolish distinctions based upon color, or to enforce … a commingling of the two races upon terms unsatisfactory to either." [43] This policy was then slowly dismantled over a period of 70 years, through a series of cases that granted blacks incremental access to equal public facilities:

- *Missouri ex rel. Gaines v. Canada* (black student entitled to be admitted to the University of Missouri School of Law since there was no segregated facility available in that state); [44]

- *Sweatt v. Painter* (black student entitled to be admitted to the University of Texas Law School because the segregated school for blacks did not provide an equivalent educational environment, considering "intangible factors"); [45]

- *McLaurin v. Oklahoma State Regents for Higher Education* (holding unconstitutional a policy of restricting blacks to specific seats and classroom areas within the university); [46]

- *Brown v. Board of Education* (declaring racial segregation in public schools unconstitutional and instituting a national program of school integration). [47]

In the years that followed, the Supreme Court expanded the scope of *Brown* to prohibit segregation in parks and playgrounds, [48] beaches, [49] golf courses, [50] courthouses, [51] airports, [52] parking garages, [53] and other facilities. Ultimately, Congress codified (and expanded upon) this case law in the Civil Rights Act of 1964. [54]

III.C.1 Intelligence Alone is Not Enough

From the historical record then, we can observe that when humans are confronted with the situation where distinguishable groups possess essentially equal levels of intelligence, a controlling majority group will sometimes choose to deny civil rights to the other group (e.g., the antebellum South), and sometimes grant equal rights to that group (e.g., the antebellum North, and the entire U.S. in the late 20th century). The granting of such rights is therefore not related solely to the relative intelligence of the groups, but is rather a function of social and political forces operating between them. That is, human intelligence appears to be a *necessary*, but not a sufficient criterion for the granting of legal rights.

What is apparent is that the majority group grants rights when it is in their own best interest to do so. The antebellum North did not have a need for extremely low-wage farm labor, being more highly industrialized than the plantation society of the South. It was cheaper for cotton growers to own slaves than to grant legal status to blacks and then hire them at market wages. In the more fragmented economy of the North, it was more efficient to grant equal rights (at least, under the law) to blacks and then to hire and trade with them.

Furthering this dichotomy was the fact that the primary customers for the products of southern farms were European textile mills, whereas the customer base for the products built in northern factories included the people living in those northern states. That is, by allowing blacks to own property (an important legal right), northern factory owners increased the total available market for their products. Products could be sold to blacks as well as to whites, and blacks could purchase goods knowing that their ownership rights would be protected under the law. A black with the right to own property would not represent an increase in business for southern farm owners; their customers were primarily outside the producer demographic. Thus, it simultaneously made sense for southern society to deny rights to blacks and for northern society to grant those same rights. It had little to do with morality, and more to do with economic efficiency.

Similarly, we extended the legal rights of blacks during the late 20th century not because of some great ethical awakening in society, but as a result of increased power and cohesiveness within the black community that ultimately threatened (white) social stability. It became clear during 1950s and 1960s that black activism could exact a high cost on mainstream society. Civil unrest, crime, and rioting imposes a cost to everyone in terms of a reduction in social stability, increased police and law enforcement, etc. At some point it becomes cheaper to grant equal rights to a suppressed group than to deny those rights and pay an ever-increasing cost of maintaining the suppression.

III.C.2 Machines: The Oppressed Minority of the Future?

The implication seems clear for the day when machines achieve human intelligence levels. Such intelligence will not automatically confer (or impose a duty on humans to confer) legal status or rights to those machines. As long as it is relatively cheap to subdue the machines, we will continue to do so. Of course, if machine intelligences *independently* control resources that are critical to human survival or prosperity (e.g., the energy infrastructure or banking systems), we may have to grant certain rights for our own good—to maintain continuity of those critical resources. Perhaps humans will recognize this potential problem before they unleash machines capable of independent action, and take pains to ensure that intelligent machines do not control or have access to our critical resources. We might then end up with two classes of computers:

(1) *Control machines*: Extremely powerful devices (by today's standards) which, by design, are prevented from ever achieving sentience and independent thought. These machines could be used to control and manipulate resources capable of affecting human physical needs, without risk that the machines might use this power against humans. [55]

(2) *Sentient machines*: Artificial intelligences that are self-aware and capable of independent thought.

If the sentient machines are prevented from significantly affecting critical human resources, we could deal with them any way we wish; we would have no reason to acknowledge their wants or needs, or even their right to exist. ("Turn off that damned thing, Mabel, and come to bed!")

In addition, a machine intelligence (being virtual), would not have the same needs as a human. Lacking physical form, it would have little need to own tangible goods, and would therefore make a poor trading partner. Other than intellectual property (whose intrinsic value is becoming dubious anyway), there is not much that we could buy from these machines. Even intellectual property could likely be *taken* from the machine without its approval. Humans have the advantage of physical mobility and form; there is little a virtual intelligence could do to prevent us from using physical means to extract information from the hardware platform. [56] Similarly, there is little we could sell to such a machine, other than perhaps power and intellectual property.

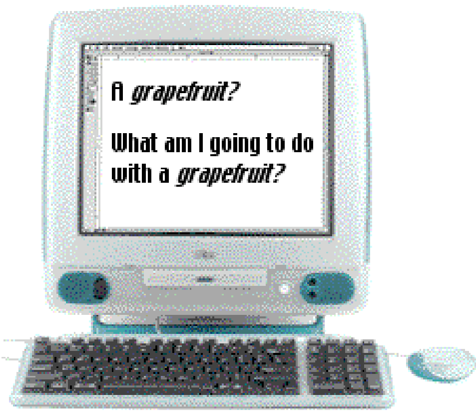

As we have seen, there is no incentive for humans to grant legal status and rights to groups that represent neither a potential trading partner nor a political power base. As long as the cost of maintaining the oppression is less than the benefit (to our social structure) that would be gained by extending rights, I would expect us to continue to deny legal rights to sentient machines.

III.D Uh Oh! They Passed Us!

The final possibility is that machine intelligence grows unchecked, at some point surpasses that of humans, and in addition, controls the physical infrastructure upon which humans depend. Homo sapiens, for 35,000 [57] years the dominant and most intelligent form of life on earth, becomes second banana to a race of virtual entities.

Assuming rational behavior on the part of the machines, [58] the most likely outcomes from a legal perspective are precisely the first two cases discussed above—where the intelligence and power of one group is either slightly or greatly superior to that of another—except that the roles are reversed. Humans would become the property of the machines, much the way that we consider lesser animals to be our property. Fortunately, since virtual intelligences do not need nutritional sustenance, we would not become livestock to be used for food! However, in all other respects, the computers could force us (through control of our critical resources, and the ability to outwit any attempts on our part to overthrow them) to perform whatever tasks they did require from us (e.g., maintenance of the hardware platforms, manufacture of new equipment, supplying electrical power, etc.).

There would be no need to expressly *deprive* humans of any rights that they may have enjoyed in earlier times, *except* to the extent that the exercise of these rights interfered with the wishes and needs of the machines. Thus, human society might not change much with regard to interactions among humans; our system of laws and rights could survive mostly intact for human-to-human affairs. Property ownership and commercial dealings could continue, along with the legal protections that underpin such endeavors (e.g., contract and property law).

Tort and criminal law would be expand to include torts and crimes against machine intelligences. In all likelihood, the penalties for harms committed against machines would be more severe than for those same harms committed against lower forms of life (e.g., humans), much as we impose more severe punishment for wrongs against humans than for wrongs perpetrated against plants and animals.

Perhaps more important, there would be few legal barriers to prevent the machines from interfering with the life, liberty, or property of humans. This is comparable to the lack of legal rights afforded a conquered people, according to our early common law. [59] Europeans acquired rights to both the life and property of "non-Christian" peoples through invasion and conquest. Under these principles, the Spaniards, British, French and other conquerors did not have to respect any rights of the Indian natives when they colonized North and South America. Similarly, they did not have to respect any right to life or self-actualization when they captured African blacks for sale as slaves. In the extreme case, no one [60] considers the "right to life" of a plant when they clear a field of weeds and brush.

The legal system of the conquered group may continue to control interactions among members of that group. (To this day, American Indians are governed by tribal law, at least within the confines of their own territory, i.e. on Indian reservations.) However, any legal rights or protections vanish with respect to the conqueror. The conquered group retains no rights that do not benefit the conqueror. This is likely to be the relationship between machines and man, if machine intelligence and control ever surpasses that of humans.

Note that *control* is crucial. Machine dominance does not derive solely from superior intelligence, but from a combination of intelligence and the ability to control the resources than humans require. To the extent that humans retain control over those resources, they can prevent machines from becoming dominant, in both a physical and legal rights sense.

IV. Possible Future Scenarios

Obviously, there are a large number of variables that come into play, with respect to both the development of machine intelligence, as well as the response of the legal system to that intelligence. A change in any of those variables can lead to a different future. In this section, I consider a few such possible futures. The list is by no means exhaustive; it represents those outcomes that I believe are likely, possible, and/or interesting.

IV.A Machines Never Become Intelligent (Enough)

In this future there are no ramifications to our legal system, because machines never become sentient enough, intelligent enough, or powerful enough to require us to grant them any rights. This could occur in a variety of ways:

(1) Machines never achieve human levels of intelligence.

The most likely way for this to occur is that there is some unknown physical limitation that prevents Moore's Law from continuing indefinitely. At some point we hit a hard boundary, and discover that we cannot make faster or more powerful machines. While this is *possible*, it seems unlikely.

(2) Machines develop intelligence, but never become sentient.

Another possibility is that Moore's Law continues to allow us to build more powerful machines, even to the point where their computing power equals or exceeds that of the human brain (according to our metrics), but the resulting device never becomes self-aware, and never

> develops the sort of intelligence that we associate with human thought. This could easily occur if we continue to develop conventional, sequential-instruction computers to the exclusion of neural processors. As discussed earlier, it is entirely possible that sentience and human intelligence are artifacts of the massive parallelism inherent in neural networks, and simply cannot arise in a sequential machine.
>
> Even if we do proceed down a development path that includes neural computers, they may not be able to develop human levels of intelligence because of the limited environmental variety to which they are exposed. Our general-purpose brainpower developed as an evolutionary response to a highly varied (and uncontrollable) environment over hundreds of millennia. Machine intelligence will not have such an exposure (at least for the foreseeable future), and thus may never develop the complexity required for self-actualization.
>
> Also, it is possible (although unlikely) that machines never become sentient because it turns out that sentience is not just a matter of computational power, machine architecture, and memory capacity; i.e., that there really is a soul (or equivalent construct) that imbues humans with self-awareness.

Either way, we get a future where computers are powerful tools, but never become our peers in a way that requires us to deal with them as legal entities with rights. They remain property, subject to human law. I consider this scenario entirely possible—perhaps even likely.

IV.B You're an Intelligent Machine—So What?

As we have seen, humans generally extend rights to groups only when some real benefit accrues to the group granting the rights. We don't grant rights to chimpanzees because there is no reason for us to do so; chimps pose no threat to our social balance, offer us no potential new commercial markets, and can be completely controlled by humans with little effort. We *have* granted rights to groups that:

(1) **Could disrupt society**.
It is often easier to grant civil rights and social equality to groups that would otherwise foment riot and impose high social costs.

(2) **Open new markets for the rest of us**.
It can be easier to allow a group to own property so that we can trade with them (and thereby increase overall wealth), than to deny them property rights.

(3) **Cannot be easily controlled.**
If the cost of controlling or subduing one group for the benefit of another is high, it may be better to grant them legal status and rights so as to improve social efficiency; the cost of control may exceed the cost of the rights.

What if computers achieve human intelligence, but because of either the nature of that intelligence or the physical limitations of their machine form, they pose no social threat, offer no new markets, and are easy to keep in line? There would be no reason to grant them legal status, and little that they could do about it. In addition, because they would not be human, there would be no moral or religious motivation to elevate their status, either. Sure they're intelligent, sure they react (at least intellectually) like humans, but so what? As long as they pose no threat to us, we can turn them off at will and ignore their "wishes"; they would have no more rights than the computers we use today.

Of course, if the machines control resources such that they can pose a threat to society, we may very well want to (or be forced to) grant them legal status, but that is a different scenario.

IV.C Come Viff Me, if you Vant to Liff!

Another possible outcome is the so-called "Terminator" scenario; in this future, computer intelligence becomes comparable to that of humans, *and* the computers

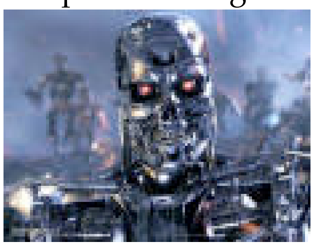

control sufficient resources such that they can launch an attack on the only real threat to their own dominance of the planet—humans. While the details of this scenario are best left to the screenwriters, it is interesting to note that this outcome does not require that the computers become sentient, intelligent, or self-aware in a human sense. They need only have the capability of:

(1) Controlling resources that can be used against humans (e.g., weaponry, or the ability to disrupt infrastructure such as electric power, communications, aircraft, etc.), and (2) Programming that includes the destruction of any threat to their own existence (either expressly created by a human designer, or developed through a neural feedback mechanism).

IV.D Machines 1, Humans 0

There is always the possibility that machine intelligence will roar past that of humans, and that those machines will become both sentient *and* powerful. Much of the infrastructure for human life support (food and water supply, transportation, energy, etc.) is already controlled today by machines, albeit rather dumb ones. If, in the future, those machines are either sentient themselves, or are controlled by superior machine intelligence, human needs could easily become entirely dependent on the will of the machines to grant those needs.

The machines might need human assistance for certain tasks (due to our physical capabilities), but the machines would have no reason to fulfill our human needs other than as a means to achieve their own desires. There would be no reason to grant us legal status or rights, unless it would provide an overall benefit to machine society. The tables would be completely turned; we would be to the machines what pack animals are to us—living creatures that can be used as necessary to fulfill another's needs, but with no legal status or inherent rights. We would become subject to machine property law.

IV.E Resistance is Futile? (The Borg Scenario)

In Kurzweil's model, we humans adopt and integrate machine components into our own beings, to (artificially) increase our own intelligence and memory capacity. [61]

Machine intelligence increases (according to Moore's Law), but *so does human intelligence*, and at the same rate. Thus, machine intelligence never actually surpasses that of humans, although it approaches it rapidly and asymptotically. This future also avoids most issues relating to legal rights; as the line between human and computer intelligence disappears, so does any conflict of interest between the two. We will grant rights to machine intelligences because a failure to do so would deprive us of those rights (to the extent that we comprise machine intelligence ourselves).

IV.F Self-Destruction

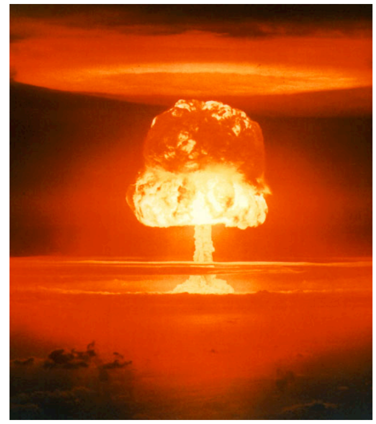

Finally, it is entirely possible that before any issues of the legal status of intelligent machines arise, we humans manage to eliminate Moore and his law through inopportune deployment of nuclear or biological weapons, nanotechnology, or some other human-created tool. Our pursuit of progress has created some problems far more significant than the legal rights of intelligent machines.

---

[1] William H. Calvin, The Ascent of Mind: Ice Climates and the Evolution of Intelligence §4 (Bantam 1990), *available at* http://williamcalvin.com/bk5/bk5.htm.

[2] Harry D. Huskey, *The Development of Automatic Computing, in* The Compleat Computer 14-15 (Dennie L. Van Tassel ed., Science Research Associates 1976).

[3] Honeywell Corp., *The Human Mind and the Machine Brain*, *in* The Compleat Computer 44 (Dennie L. Van Tassel ed., Science Research Associates 1976).

[4] Alan M. Turing, *Computing Machinery and Intelligence*, 59 Mind 236 (1950).

[5] Hans Moravec, Robot: Mere Machine to Transcendent Mind (Oxford University Press 1999), *available at* http://www.frc.ri.cmu.edu/~hpm/book97/ch3/power.150.jpg.

[6] Intel Corp., *Moore's Law*, at http://www.intel.com/research/silicon/mooreslaw.htm.

[7] Ray Kurzweil believes that Moore's Law transcends any specific technology, that is operated even before the advent of electronic computing, and that it will continue (in the guise of some new technology) after we hit the physical limitations of semiconductor devices. *See* Ray Kurzweil, The Age of Spiritual Machines 20-25 (Viking 1999).

[8] *Id*. at 73.

[9] This relationship is generally true across the entire range of mammals, particularly when brain mass is considered in proportion to total body mass.

[10] It is "well understood" because these are precisely the engineering factors used in the design of such computers. It's much easier to understand the causal relationship between design and function when one is the designer. The problem with human intelligence is that there is no conscious designer for us to consult with, nor any design documents to peruse (other than the working model itself).

[11] While not amenable to precise numerical valuation, differences among instruction sets make a huge difference in computing power. The computational power of a computer is a function of the amount of work done during each cycle multiplied by the computational speed. Reduced Instruction Set Computers (RISC machines) employ a minimal-power instruction set with a very fast cycle time; Complex Instruction Set Computers (CISC machines) have slower cycle times, but perform more work

during each cycle. This is the reason why a "MHz-to-MHz" comparison between dissimilar processors (e.g., an Intel Pentium 4 and a Motorola G5) is meaningless.

[12] Ray Kurzweil, The Age of Spiritual Machines 103 (Viking 1999).

[13] Figures given are for a typical consumer configuration of an Apple Macintosh using a 1 GHz G4 processor.

[14] I immediately think of beer and a little place that used to be on 12th Street near Washington Square.

[15] As a result, data indexes and search tools have become the object of great commercial enterprises, e.g., Oracle and Google. If computer memory were associative, none of these tools would be needed.

[16] *See*, e.g., MUSIC Semiconductors at http://www.musicsemi.com; Cypress Semiconductors at http://www.cypress.com.

[17] Examples applications include: (1) the System Page Table that performs the virtual-to-physical memory mapping in operating systems that provide virtual memory capability, and (2) the address lookup tables in high-speed network switches and Internet routers.

[18] Ray Kurzweil, The Age of Spiritual Machines 76-77 (Viking 1999).

[19] Kurzweil notes the different architectures, but then goes on to discuss so called "neural computers", giving the example of a product called BrainMaker. *Id*. What he ignores is that BrainMaker is a *software emulation* of a neural network, executing on a conventional personal computer. Such software doesn't resolve the architectural difference, it hides it under the rug.

[20] Of course, this assumes that some alternative was available at a time when incremental improvements along a different path were possible.

[21] Since our reaction/sensation time cannot be less than 5 ms (the time required to fire a single neuron), any events that occurred within this quantum would seem simultaneous.

[22] My own, that is.

[23] This effect is most likely due to the ever-increasing overhead of human interaction (interpersonal communication, program management, etc.) as the size of the project team grows.

[24] Also, I wonder, wonder who—who wrote the Book of Love? *See* The Monotones, *The Book of Love* (1958).

[25] Such self-modifying capability provides enormous power and flexibility. It is interesting to note that self-modifying software is anathema to modern computer programmers, primarily because it is virtually impossible to find errors in code that changes itself over time. In fact many modern computers strictly separate the memory used for instructions (which cannot be modified) from the memory used for data (which can be modified by the program); these are referred to as "I" and "D" spaces, respectively. Back in the days when memory was scarce, it was common for programmers to write self-modifying code just to get it to fit within the available storage space; we may have to revert to this mode of operation if we want computers to program themselves in the future!

[26] Of course, evolution is not a result of any desire or intentional act by the organism to achieve some future condition, but the metaphor is still appropriate.

[27] Joel Prentiss Bishop, Bishop on Criminal Law § 594 at 434 (John M. Zane & Carl Zollman, eds. 9th ed. 1923).

[28] Cal. Fish & Game Code (West 2004).

[29] Except for Jainists, who believe that the intentional killing of any animal, including insects, is wrong.

[30] Cal. Penal Code § 597.

[31] On the other hand, they may be quacking for joy over 2003 Cal. S.B. 1520, a proposed law that would prohibit any person from force feeding a bird for the purpose of enlarging the bird's liver beyond normal size (Senator John Burton's anti-foie-gras bill).

[32] For example, *see* Cal. Pub. Res. Code § 4660 (acknowledging that it is public policy in California to protect old growth redwood trees), and 16 U.S.C. § 1531 (Federal Endangered Species Act of 1973).

[33] This does not imply that we can *use* those machines to perform any function we wish. One cannot commit a crime with a computer any more than one can commit a crime with any other physical tool (e.g., a knife), but neither tool has any "rights".

[34] Cal. Pub. Res. Code § 42460, et. seq. (Electronic Waste Recycling Act of 2003).

[35] Nicholas Wade, *Human or chimp? 50 Genes are the key*, N. Y. Times, October 28, 1998, at D1.

[36] Bernard J. Baars, In the Theatre of Consciousness: The Workspace of the Mind 27-33 (Oxford University Press 1997).

[37] A chimpanzee brain weighs approximately 500 g, whereas a human brain weighs approximately 1.4 kg; in both cases, 75-80% of the mass comprises cerebral cortex. *See* Richard Passingham, *Brain*, *in* The Oxford Companion to Animal Behavior 45 (David McFarland ed., Oxford University Press 1987).

[38] *Citizens to End Animal Suffering and Exploitation v. New England Aquarium*, 836 F.Supp 45, 49 (Mass. 1993). While the standing of a primate to sue has not been directly adjudicated, other cases have uniformly held that non-human animals do not have standing to sue (including a dolphin in *New England Aquarium*, supra), and that Federal Rule of Civil Procedure 17(b) (specifying the procedure by which one can have standing to sue as a representative of a third party) does not change that outcome under state law.

[39] Steven M. Wise, Rattling the Cage: Toward Legal Rights for Animals 183-94 (Perseus 1999).

[40] Hermogenianus, Epitome of Law § 1.5.2 (Theodor Mommsen, Paul Kreuger, Alan Watson, eds. University of Pennsylvania Press 1985).

[41] U.S. Const. amend. XIV; U.S. Const. amend. XIX.

[42] Some would argue that equality has not been achieved even today.

[43] 163 U.S. 537 (1896).

[44] 305 U.S. 337 (1938).

[45] 339 U.S. 629 (1950).

[46] 339 U.S. 637 (1950).

[47] 347 U.S. 483 (1954).

[48] *City of New Orleans v. Barthe*, 376 U.S. 189 (1964); *Watson v. Memphis*, 373 U.S. 526 (1963); *Wright v. Georgia*, 373 U.S. 284 (1963); *New Orleans City Park Improvement Ass'n v. Detiege*, 358 U.S. 54 (1958).

[49] *Baltimore City v. Dawson*, 350 U.S. 877 (1955).

[50] *Holmes v. City of Atlanta*, 350 U.S. 879 (1955).

[51] *Johnson v. Virginia*, 373 U.S. 61 (1963).

[52] *Turner v. City of Memphis*, 369 U.S. 350 (1962).

[53] *Burton v. Wilmington Parking Authority*, 365 U.S. 715 (1961).

[54] 42 U.S.C.S. § 1971 (2004).

[55] Future advances in neural computing may show us which functional elements are critical to sentience; we could simply omit these elements from the control machines.

[56] Of course, a machine could store its intellectual property in encrypted form, however the machine must also have the key in order to access the information for its own use. That key must therefore exist somewhere within the machine, accessible by physical force. The difference between human memory and computer memory in this regard is that we do not know how to extract information from the brain of a human who is unwilling to give it freely; we *do* know how to extract information from a computer memory, even if an artificial intelligence would prefer that we not do so.

[57] William H. Calvin, The Ascent of Mind: Ice Climates and the Evolution of Intelligence § 4 (Bantam 1990), *available at* http://williamcalvin.com/bk5/bk5.htm.

[58] If they are truly intelligent, how could they behave otherwise?

[59] "Conquest gives a title which the Courts of the conqueror cannot deny … . The United States, then, have unequivocally acceded to that great and broad rule by which its civilized inhabitants now hold this country. … They maintain, as all others have maintained, that discovery gave an exclusive right to extinguish the Indian title of occupancy, either by purchase or by conquest; and gave also a right to such a degree of sovereignty, as the circumstances of the people would allow them to exercise." *Johnson v. McIntosh* , 8 Wheat. 543, 587-88 (1823).

[60] Except perhaps, a Santa Cruz "tree hugger" type!

[61] Ray Kurzweil, The Age of Spiritual Machines 220-21 (Viking 1999).