

When will computer hardware match the human brain?

(Received Dec. 1997)

Hans Moravec

Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213-3890, USA
net: hpm@cmu.edu
web: <http://www.frc.ri.cmu.edu/~hpm/>

ABSTRACT

This paper describes how the performance of AI machines tends to improve at the same pace that AI researchers get access to faster hardware. The processing power and memory capacity necessary to match general intellectual performance of the human brain are estimated. Based on extrapolation of past trends and on examination of technologies under development, it is predicted that the required hardware will be available in cheap machines in the 2020s.

Brains, Eyes and Machines

Computers have far to go to match human strengths, and our estimates will depend on analogy and extrapolation. Fortunately, these are grounded in the first bit of the journey, now behind us. Thirty years of computer vision reveals that 1 MIPS can extract simple features from real-time imagery—tracking a white line or a white spot on a mottled background. 10 MIPS can follow complex gray-scale patches—as smart bombs, cruise missiles and early self-driving vans attest. 100 MIPS can follow moderately unpredictable features like roads—as recent long NAVLAB trips demonstrate. 1,000 MIPS will be adequate for coarse-grained three-dimensional spatial awareness—illustrated by several mid-resolution stereoscopic vision programs, including my own. 10,000 MIPS can find three-dimensional objects in clutter—suggested by several "bin-picking" and high-resolution stereo-vision demonstrations, which accomplish the task in an hour or so at 10 MIPS. The data fades there—research careers are too short, and computer memories too small, for significantly more elaborate experiments.

There are considerations other than sheer scale. At 1 MIPS the best results come from finely hand-crafted programs that distill sensor data with utmost efficiency. 100-MIPS processes weigh their inputs against a wide range of hypotheses, with many parameters, that learning programs adjust better than the overburdened programmers. Learning of all sorts will be increasingly important as computer power and robot programs grow. This effect is evident in related areas. At the close of the 1980s, as widely available computers reached 10 MIPS, good optical character reading (OCR) programs, able to read most printed and typewritten text, began to appear. They used hand-constructed "feature detectors" for parts of letter shapes, with very little learning. As computer power passed 100 MIPS, trainable OCR programs appeared that could learn unusual tpestyles from examples, and the latest and best programs learn their entire data sets. Handwriting recognizers, used by the Post Office to sort mail, and in computers, notably Apple's Newton, have followed a similar path. Speech recognition also fits the model. Under the direction of Raj Reddy, who began his research at Stanford in the 1960s, Carnegie Mellon has led in computer transcription of continuous spoken

speech. In 1992 Reddy's group demonstrated a program called *Sphinx II* on a 15-MIPS workstation with 100 MIPS of specialized signal-processing circuitry. *Sphinx II* was able to deal with arbitrary English speakers using a several-thousand-word vocabulary. The system's word detectors, encoded in statistical structures known as Markov tables, were shaped by an automatic learning process that digested hundreds of hours of spoken examples from thousands of Carnegie Mellon volunteers enticed by rewards of pizza and ice cream. Several practical voice-control and dictation systems are sold for personal computers today, and some heavy users are substituting larynx for wrist damage.

More computer power is needed to reach human performance, but how much? Human and animal brain sizes imply an answer, if we can relate nerve volume to computation. Structurally and functionally, one of the best understood neural assemblies is the retina of the vertebrate eye. Happily, similar operations have been developed for robot vision, handing us a rough conversion factor.

The retina is a transparent, paper-thin layer of nerve tissue at the back of the eyeball on which the eye's lens projects an image of the world. It is connected by the optic nerve, a million-fiber cable, to regions deep in the brain. It is a part of the brain convenient for study, even in living animals because of its peripheral location and because its function is straightforward compared with the brain's other mysteries. A human retina is less than a centimeter square and a half-millimeter thick. It has about 100 million neurons, of five distinct kinds. Light-sensitive cells feed wide spanning horizontal cells and narrower bipolar cells, which are interconnected by whose outgoing fibers bundle to form the optic nerve. Each of the million ganglion-cell axons carries signals from a amacrine cells, and finally ganglion cells, particular patch of image, indicating light intensity differences over space or time: a million edge and motion detections. Overall, the retina seems to process about ten one-million-point images per second.

It takes robot vision programs about 100 computer instructions to derive single edge or motion detections from comparable video images. 100 million instructions are needed to do a million detections, and 1,000 MIPS to repeat them ten times per second to match the retina.

The 1,500 cubic centimeter human brain is about 100,000 times as large as the retina, suggesting that matching overall human behavior will take about 100 million MIPS of computer power. Computer chess bolsters this yardstick. Deep Blue, the chess machine that bested world chess champion Garry Kasparov in 1997, used specialized chips to process chess moves at a the speed equivalent to a 3 million MIPS universal computer (see Figure 3-4). This is 1/30 of the estimate for total human performance. Since it is plausible that Kasparov, probably the best human player ever, can apply his brainpower to the strange problems of chess with an efficiency of 1/30, Deep Blue's near parity with Kasparov's chess skill supports the retina-based extrapolation.

The most powerful experimental supercomputers in 1998, composed of thousands or tens of thousands of the fastest microprocessors and costing tens of millions of dollars, can do a few million MIPS. They are within striking distance of being powerful enough to match human brainpower, but are unlikely to be applied to that end. Why tie up a rare twenty-million-dollar asset to develop one ersatz-human, when millions of inexpensive original-model humans are available? Such machines are needed for high-value scientific calculations, mostly physical simulations, having no cheaper substitutes. AI research must wait for the power to become more affordable.

If 100 million MIPS could do the job of the human brain's 100 billion neurons, then one neuron is worth about 1/1,000 MIPS, i.e., 1,000 instructions per second. That's probably not enough to simulate an actual neuron, which can produce 1,000 finely timed pulses per second. Our estimate is for very efficient programs that imitate the aggregate function of thousand-neuron assemblies. Almost all nervous systems contain subassemblies that big.


The small nervous systems of insects and other invertebrates seem to be hardwired from birth, each neuron having its own special predetermined links and function. The few-hundred-million-bit insect genome is enough to specify connections of each of their hundred thousand neurons. Humans, on the other hand, have

100 billion neurons, but only a few billion bits of genome. The human brain seems to consist largely of regular structures whose neurons are trimmed away as skills are learned, like featureless marble blocks chiseled into individual sculptures. Analogously, robot programs were precisely hand-coded when they occupied only a few hundred thousand bytes of memory. Now that they've grown to tens of millions of bytes, most of their content is learned from example. But there is a big practical difference between animal and robot learning. Animals learn individually, but robot learning can be copied from one machine to another. For instance, today's text and speech understanding programs were painstakingly trained over months or years, but each customer's copy of the software is "born" fully educated. Decoupling training from use will allow robots to do more with less. Big computers at the factory—maybe supercomputers with 1,000 times the power of machines that can reasonably be placed in a robot—will process large training sets under careful human supervision, and distill the results into efficient programs and arrays of settings that are then copied into myriads of individual robots with more modest processors.

Programs need memory as well as processing speed to do their work. The ratio of memory to speed has remained constant during computing history. The earliest electronic computers had a few thousand bytes of memory and could do a few thousand calculations per second. Medium computers of 1980 had a million bytes of memory and did a million calculations per second. Supercomputers in 1990 did a billion calculations per second and had a billion bytes of memory. The latest, greatest supercomputers can do a trillion calculations per second and can have a trillion bytes of memory. Dividing memory by speed defines a "time constant," roughly how long it takes the computer to run once through its memory. One megabyte per MIPS gives one second, a nice human interval. Machines with less memory for their speed, typically new models, seem fast, but unnecessarily limited to small programs. Models with more memory for their speed, often ones reaching the end of their run, can handle larger programs, but unpleasantly slowly. For instance, the original Macintosh was introduced in 1984 with 1/2 MIPS and 1/8 megabyte, and was then considered a very fast machine. The equally fast "fat Mac" with 1/2 megabyte ran larger programs at tolerable speed, but the 1 megabyte "Mac plus" verged on slow. The four megabyte "Mac classic," the last 1/2 MIPS machine in the line, was intolerably slow, and was soon supplanted by ten-times-faster processors in the same enclosure. Customers maintain the ratio by asking "would the next dollar be better spent on more speed or more memory?"

The best evidence about nervous system memory puts most of it in the synapses connecting the neurons. Molecular adjustments allow synapses to be in a number of distinguishable states, let's say one byte's worth. Then the 100-trillion-synapse brain would hold the equivalent 100 million megabytes. This agrees with our earlier estimate that it would take 100 million MIPS to mimic the brain's function. The megabyte/MIPS ratio seems to hold for nervous systems too! The contingency is the other way around: computers are configured to interact at human time scales, and robots interacting with humans seem also to be best at that ratio. On the other hand, faster machines, for instance audio and video processors and controllers of high-performance aircraft, have many MIPS for each megabyte. Very slow machines, for instance time-lapse security cameras and automatic data libraries, store many megabytes for each of their MIPS. Flying insects seem to be a few times faster than humans, so may have more MIPS than megabytes. As in animals, cells in plants signal one other electrochemically and enzymatically. Some plant cells seem specialized for communication, though apparently not as extremely as animal neurons. One day we may find that plants remember much, but process it slowly (how does a redwood tree manage to rebuff rapidly evolving pests during a 2,000 year lifespan, when it took mosquitoes only a few decades to overcome DDT?).

With our conversions, a 100-MIPS robot, for instance Navlab, has mental power similar to a 100,000-neuron housefly. The following figure rates various entities.



MIPS and Megabytes. to mimic their behavior. Note the scale. Entities rated by the computational power and memory of the smallest universal computer needed is logarithmic on both axes: each vertical division represents a thousandfold increase in processing power, and each horizontal division a thousandfold increase in memory size. Universal computers can imitate other entities at their location in the diagram, but the more specialized entities cannot. A 100-million-MIPS computer may be programmed not only to think like a human, but also to imitate other similarly-sized computers. But humans cannot imitate 100-million-MIPS computers--our general-purpose calculation ability is under a millionth of a MIPS. Deep Blue's special-purpose chess chips process moves like a 3-million-MIPS computer, but its general-purpose power is only a thousand MIPS. Most of the non-computer entities in the diagram can't function in a general-purpose way at all. Universality is an almost magical property, but it has costs. A universal machine may use ten or more times the resources of one specialized for a task. But if the task should change, as it usually does in research, the universal machine can be reprogrammed, while the specialized machine must be replaced.


Extrapolation

By our estimate, today's very biggest supercomputers are within a factor of a hundred of having the power to mimic a human mind. Their successors a decade hence will be more than powerful enough. Yet, it is unlikely that machines costing tens of millions of dollars will be wasted doing what any human can do, when they could instead be solving urgent physical and mathematical problems nothing else can touch. Machines with human-like performance will make economic sense only when they cost less than humans, say when their "brains" cost about \$1,000. When will that day arrive?

The expense of computation has fallen rapidly and persistently for a century. Steady improvements in mechanical and electromechanical calculators before World War II had increased the speed of calculation a thousandfold over hand calculation. The pace quickened with the appearance of electronic computers during the war--from 1940 to 1980 the amount of computation available at a given cost increased a millionfold.

Vacuum tubes were replaced by transistors, and transistors by integrated circuits, whose components became ever smaller and more numerous. During the 1980s microcomputers reached the consumer market, and the industry became more diverse and competitive. Powerful, inexpensive computer workstations replaced the drafting boards of circuit and computer designers, and an increasing number of design steps were automated. The time to bring a new generation of computer to market shrank from two years at the beginning of the 1980s to less than nine months. The computer and communication industries grew into the largest on earth.

Computers doubled in capacity every two years after the war, a pace that became an industry given: companies that wished to grow sought to exceed it, companies that failed to keep up lost business. In the 1980s the doubling time contracted to 18 months, and computer performance in the late 1990s seems to be doubling every 12 months.



Faster than Exponential Growth in Computing Power. The number of MIPS in \$1000 of computer from 1900 to the present. Steady improvements in mechanical and electromechanical calculators before World War II had increased the speed of calculation a thousandfold over manual methods from 1900 to 1940. The pace quickened with the appearance of electronic computers during the war, and 1940 to 1980 saw a millionfold increase. The pace has been even quicker since then, a pace which would make humanlike robots possible before the middle of the next century. The vertical scale is logarithmic, the major divisions represent thousandfold increases in computer performance. Exponential growth would show as a straight line, the upward curve indicates faster than exponential growth, or, equivalently, an accelerating rate of innovation. The reduced spread of the data in the 1990s is probably the result of intensified competition: underperforming machines are more rapidly squeezed out. The numerical data for this power curve are presented in [the appendix](#).

At the present rate, computers suitable for humanlike robots will appear in the 2020s. Can the pace be sustained for another three decades? The graph shows no sign of abatement. If anything, it hints that further contractions in time scale are in store. But, one often encounters thoughtful articles by knowledgeable people in the semiconductor industry giving detailed reasons why the decades of phenomenal growth must soon come to an end.

The keynote for advancing computation is miniaturization: smaller components have less inertia and operate more quickly with less energy, and more of them can be packed in a given space. First the moving parts shrunk, from the gears in mechanical calculators, to small contacts in electromechanical machines, to bunches of electrons in electronic computers. Next, the switches' supporting structure underwent a vanishing act, from thumb-sized vacuum tubes, to fly-sized transistors, to ever-diminishing flyspecks on integrated circuit chips. Similar to printed circuits before them, integrated circuits were made by a photographic process. The desired pattern was projected onto a silicon chip, and subtle chemistry used to add or remove the right sorts of matter in the exposed areas.

In the mid-1970s, integrated circuits, age 15, hit a crisis of adolescence. They then held ten thousand components, just enough for an entire computer, and their finest details were approaching 3 micrometers in size. Experienced engineers wrote many articles warning that the end was near. Three micrometers was barely larger than the wavelength of the light used to sculpt the chip. The number of impurity atoms defining the tiny components had grown so small that statistical scatter would soon render most components out of spec, a problem aggravated by a similar effect in the diminishing number of signaling electrons. Increasing electrical gradients across diminishing gaps caused atoms to creep through the crystal, degrading the circuit. Interactions between ever-closer wires were about to ruin the signals. Chips would soon generate too much heat to remove, and require too many external connections to fit. The smaller memory cells were suffering radiation-induced forgetfulness.

A look at the computer growth graph shows that the problems were overcome, with a vengeance. Chip progress not only continued, it sped up. Shorter-wavelength light was substituted, a more precise way of implanting impurities was devised, voltages were reduced, better insulators, shielding designs, more efficient transistor designs, better heat sinks, denser pin patterns and non-radioactive packaging materials were found. Where there is sufficient financial incentive, there is a way. In fact, solutions had been waiting in research labs for years, barely noticed by the engineers in the field, who were perfecting established processes, and worrying in print as those ran out of steam. As the need became acute, enormous resources were redirected to draft laboratory possibilities into production realities.

In the intervening years many problems were met and solved, and innovations introduced, but now, nearing a mid-life 40, the anxieties seem again to have crested. In 1996 major articles appeared in scientific magazines and major national newspapers worrying that electronics progress might be a decade from ending. The cost of building new integrated circuit plants was approaching a prohibitive billion dollars. Feature sizes were reaching 0.1 micrometers, the wavelength of the sculpting ultraviolet light. Their transistors, scaled down steadily from 1970s designs, would soon be so small that electrons would quantum "tunnel" out of them. Wiring was becoming so dense it would crowd out the components, and slow down and leak signals. Heat was increasing.

The articles didn't mention that less expensive plants could make the same integrated circuits, if less cheaply and in smaller quantities. Scale was necessary because the industry had grown so large and competitive. Rather than signaling impending doom, it indicated free-market success, a battle of titans driving down costs to the users. They also failed to mention new contenders, waiting on lab benches to step in should the leader fall.

The wave-like nature of matter at very small scales is a problem for conventional transistors, which depend on the smooth flow of masses of electrons. But, it is a property exploited by a radical new class of components known as single-electron transistors and quantum dots, which work by the interference of electron waves.

These new devices work better as they grow smaller. At the scale of today's circuits, the interference patterns are so fine that it takes only a little heat energy to bump electrons from crest to crest, scrambling their operation. Thus, these circuits have been demonstrated mostly at a few degrees above absolute zero. But, as the devices are reduced, the interference patterns widen, and it takes ever larger energy to disrupt them. Scaled to about 0.01 micrometers, quantum interference switching works at room temperature. It promises more than a thousand times higher density than today's circuits, possibly a thousand times the speed, and much lower power consumption, since it moves a few electrons across small quantum bumps, rather than pushing them in large masses through resistive material. In place of much wiring, quantum interference logic may use chains of switching devices. It could be manufactured by advanced descendants of today's chip fabrication machinery (Goldhaber–Gordon et al. 1997). Proposals abound in the research literature, and the industry has the resources to perfect the circuits and their manufacture, when the time comes.

Wilder possibilities are brewing. Switches and memory cells made of single molecules have been demonstrated, which might enable a volume to hold a billion times more circuitry than today. Potentially blowing everything else away are "quantum computers," in which a whole computer, not just individual signals, acts in a wavelike manner. Like a conventional computer, a quantum computer consists of a number of memory cells whose contents are modified in a sequence of logical transformations. Unlike a conventional computer, whose memory cells are either 1 or 0, each cell in a quantum computer is started in a quantum superposition of both 1 and 0. The whole machine is a superposition of all possible combinations of memory states. As the computation proceeds, each component of the superposition individually undergoes the logic operations. It is as if an exponential number of computers, each starting with a different pattern in memory, were working on the problem simultaneously. When the computation is finished, the memory cells are examined, and an answer emerges from the wavelike interference of all the possibilities. The trick is to devise the computation so that the desired answers reinforce, while the others cancel. In the last several years, quantum algorithms have been devised that factor numbers and search for encryption keys much faster than any classical computer. Toy quantum computers, with three or four "qubits" stored as states of single atoms or photons, have been demonstrated, but they can do only short computations before their delicate superpositions are scrambled by outside interactions. More promising are computers using nuclear magnetic resonance, as in hospital scanners. There, quantum bits are encoded as the spins of atomic nuclei, and gently nudged by external magnetic and radio fields into magnetic interactions with neighboring nuclei. The heavy nuclei, swaddled in diffuse orbiting electron clouds, can maintain their quantum coherence for hours or longer. A quantum computer with a thousand or more qubits could tackle problems astronomically beyond the reach of any conceivable classical computer.

Molecular and quantum computers will be important sooner or later, but humanlike robots are likely to arrive without their help. Research within semiconductor companies, including working prototype chips, makes it quite clear that existing techniques can be nursed along for another decade, to chip features below 0.1 micrometers, memory chips with tens of billions of bits and multiprocessor chips with over 100,000 MIPS. Towards the end of that period, the circuitry will probably incorporate a growing number of quantum interference components. As production techniques for those tiny components are perfected, they will begin to take over the chips, and the pace of computer progress may steepen further. The 100 million MIPS to match human brain power will then arrive in home computers before 2030.

False Start

It may seem rash to expect fully intelligent machines in a few decades, when the computers have barely matched insect mentality in a half-century of development. Indeed, for that reason, many long-time artificial intelligence researchers scoff at the suggestion, and offer a few centuries as a more believable period. But there are very good reasons why things will go much faster in the next fifty years than they have in the last fifty.

The stupendous growth and competitiveness of the computer industry is one reason. A less appreciated one is

that intelligent machine research did not make steady progress in its first fifty years, it marked time for thirty of them! Though general computer power grew a hundred thousand fold from 1960 to 1990, the computer power available to AI programs barely budged from 1 MIPS during those three decades.


In the 1950s, the pioneers of AI viewed computers as locomotives of thought, which might outperform humans in higher mental work as prodigiously as they outperformed them in arithmetic, if they were harnessed to the right programs. Success in the endeavor would bring enormous benefits to national defense, commerce and government. The promise warranted significant public and private investment. For instance, there was a large project to develop machines to automatically translate scientific and other literature from Russian to English. There were only a few AI centers, but those had the largest computers of the day, comparable in cost to today's supercomputers. A common one was the IBM 704, which provided a good fraction of a MIPS.

By 1960 the unspectacular performance of the first reasoning and translation programs had taken the bloom off the rose, but the unexpected launching by the Soviet Union of Sputnik, the first satellite in 1957, had substituted a paranoia. Artificial Intelligence may not have delivered on its first promise, but what if it were to suddenly succeed after all? To avoid another nasty technological surprise from the enemy, it behooved the US to support the work, moderately, just in case. Moderation paid for medium scale machines costing a few million dollars, no longer supercomputers. In the 1960s that price provided a good fraction of a MIPS in thrifty machines like Digital Equipment Corp's innovative PDP-1 and PDP-6.

The field looked even less promising by 1970, and support for military-related research declined sharply with the end of the Vietnam war. Artificial Intelligence research was forced to tighten its belt and beg for unaccustomed small grants and contracts from science agencies and industry. The major research centers survived, but became a little shabby as they made do with aging equipment. For almost the entire decade AI research was done with PDP-10 computers, that provided just under 1 MIPS. Because it had contributed to the design, the Stanford AI Lab received a 1.5 MIPS KL-10 in the late 1970s from Digital, as a gift.

Funding improved somewhat in the early 1980s, but the number of research groups had grown, and the amount available for computers was modest. Many groups purchased Digital's new Vax computers, costing \$100,000 and providing 1 MIPS. By mid-decade, personal computer workstations had appeared. Individual researchers reveled in the luxury of having their own computers, avoiding the delays of time-shared machines. A typical workstation was a Sun-3, costing about \$10,000, and providing about 1 MIPS.

By 1990, entire careers had passed in the frozen winter of 1-MIPS computers, mainly from necessity, but partly from habit and a lingering opinion that the early machines really should have been powerful enough. In 1990, 1 MIPS cost \$1,000 in a low-end personal computer. There was no need to go any lower. Finally spring thaw has come. Since 1990, the power available to individual AI and robotics programs has doubled yearly, to 30 MIPS by 1994 and 500 MIPS by 1998. Seeds long ago alleged barren are suddenly sprouting. Machines read text, recognize speech, even translate languages. Robots drive cross-country, crawl across Mars, and trundle down office corridors. In 1996 a theorem-proving program called EQP running five weeks on a 50 MIPS computer at Argonne National Laboratory found a proof of a boolean algebra conjecture by Herbert Robbins that had eluded mathematicians for sixty years. And it is still only spring. Wait until summer.



The big freeze. From 1960 to 1990 the cost of computers used in AI research declined, as their numbers dilution absorbed computer-efficiency gains during the period, and the power available to individual AI programs remained almost unchanged at 1 MIPS, barely insect power. AI computer cost bottomed in 1990, and since then power has doubled yearly, to several hundred MIPS by 1998. The major visible exception is computer chess (shown by a progression of knights), whose prestige lured the resources of major computer companies and the talents of programmers and machine designers. Exceptions also exist in less public competitions, like petroleum exploration and intelligence gathering, whose high return on investment gave them regular access to the largest computers.

The Game's Afoot

A summerlike air already pervades the few applications of artificial intelligence that retained access to the largest computers. Some of these, like pattern analysis for satellite images and other kinds of spying, and in seismic oil exploration, are closely held secrets. Another, though, basks in the limelight. The best chess-playing computers are so interesting they generate millions of dollars of free advertising for the winners, and consequently have enticed a series of computer companies to donate time on their best machines and other resources to the cause. Since 1960 IBM, Control Data, AT&T, Cray, Intel and now again IBM have been sponsors of computer chess. The "knights" in the AI power graph show the effect of this largesse, relative to mainstream AI research. The top chess programs have competed in tournaments powered by supercomputers, or specialized machines whose chess power is comparable. In 1958 IBM had both the first checker program, by Arthur Samuel, and the first full chess program, by Alex Bernstein. They ran on an IBM 704, the biggest and last vacuum-tube computer. The Bernstein program played atrociously, but Samuel's program, which automatically learned its board scoring parameters, was able to beat Connecticut checkers champion Robert Nealey. Since 1994, Chinook, a program written by Jonathan Schaeffer of the University of

Alberta, has consistently bested the world's human checker champion. But checkers isn't very glamorous, and this portent received little notice.

By contrast, it was nearly impossible to overlook the epic battles between world chess champion Garry Kasparov and IBM's Deep Blue in 1996 and 1997. Deep Blue is a scaled-up version of a machine called Deep Thought, built by Carnegie Mellon University students ten years earlier. Deep Thought, in turn, depended on special-purpose chips, each wired like the Belle chess computer built by Ken Thompson at AT&T Bell Labs in the 1970s. Belle, organized like a chessboard, circuitry on the squares, wires running like chess moves, could evaluate and find all legal moves from a position in one electronic flash. In 1997 Deep Blue had 256 such chips, orchestrated by a 32 processor mini-supercomputer. It examined 200 million chess positions a second. Chess programs, on unaided general-purpose computers, average about 16,000 instructions per position examined. Deep Blue, when playing chess (and only then), was thus worth about 3 million MIPS, 1/30 of our estimate for human intelligence.

Deep Blue, in a first for machinekind, won the first game of the 1996 match. But, Kasparov quickly found the machine's weaknesses, and drew two and won three of the remaining games.

In May 1997 he met an improved version of the machine. That February, Kasparov had triumphed over a field of grandmasters in a prestigious tournament in Linares, Spain, reinforcing his reputation as the best player ever, and boosting his chess rating past 2800, uncharted territory. He prepared for the computer match in the intervening months, in part by playing against other machines. Kasparov won a long first game against Deep Blue, but lost next day to masterly moves by the machine. Then came three grueling draws, and a final game, in which a visibly shaken and angry Kasparov resigned early, with a weak position. It was the first competition match he had ever lost.

The event was notable for many reasons, but one especially is of interest here. Several times during both matches, Kasparov reported signs of mind in the machine. At times in the second tournament, he worried there might be humans behind the scenes, feeding Deep Blue strategic insights!

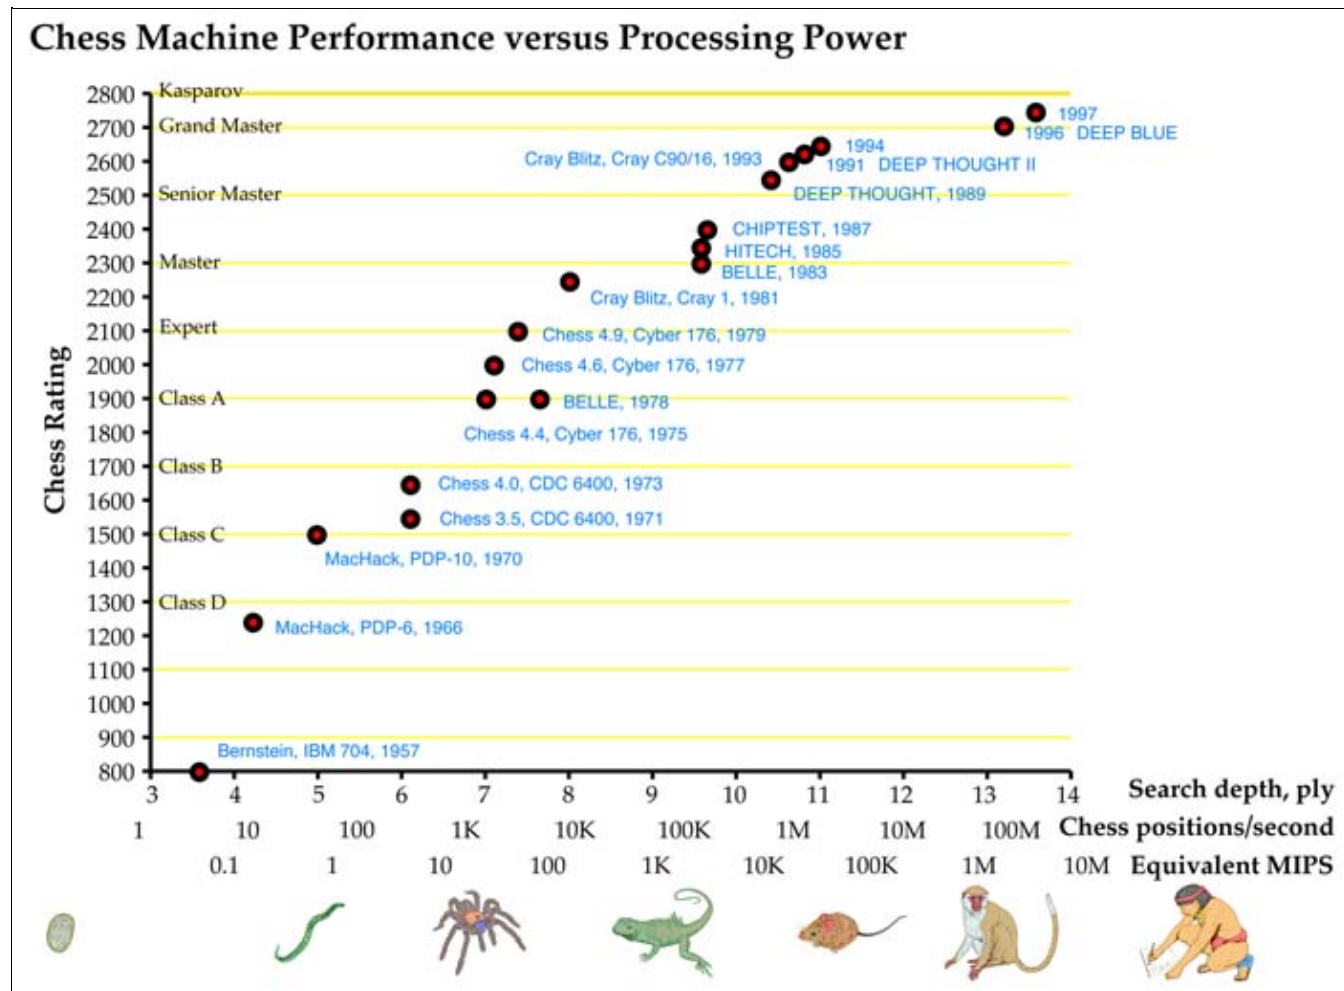
Bobby Fischer, the US chess great of the 1970s, is reputed to have played each game as if against God, simply making the best moves. Kasparov, on the other hand, claims to see into opponents' minds during play, intuiting and exploiting their plans, insights and oversights. In all other chess computers, he reports a mechanical predictability stemming from their indiscriminating but limited lookahead, and absence of long-term strategy. In Deep Blue, to his consternation, he saw instead an "alien intelligence."

In this paper-thin slice of mentality, a computer seems to have not only outperformed the best human, but to have transcended its machinehood. Who better to judge than Garry Kasparov? Mathematicians who examined EQP's proof of the Robbins conjecture, mentioned earlier, report a similar impression of creativity and intelligence. In both cases, the evidence for an intelligent mind lies in the machine's performance, not its makeup.

Now, the team that built Deep Blue claim no "intelligence" in it, only a large database of opening and end games, scoring and deepening functions tuned with consulting grandmasters, and, especially, raw speed that allows the machine to look ahead an average of fourteen half-moves per turn. Unlike some earlier, less successful, chess programs, Deep Blue was not designed to think like a human, to form abstract strategies or see patterns as it races through the move/countermove tree as fast as possible.

Deep Blue's creators know its *quantitative* superiority over other chess machines intimately, but lack the chess understanding to share Kasparov's deep appreciation of the difference in the *quality* of its play. I think this dichotomy will show up increasingly in coming years. Engineers who know the mechanism of advanced robots most intimately will be the last to admit they have real minds. From the inside, robots will indisputably be machines, acting according to mechanical principles, however elaborately layered. Only on the outside, where they can be appreciated as a whole, will the impression of intelligence emerge. A human brain, too, does not exhibit the intelligence under a neurobiologist's microscope that it does participating in a lively

conversation.



Agony to ecstasy. In forty years, computer chess progressed from the lowest depth to the highest peak of human chess performance. It took a handful of good ideas, culled by trial and error from a larger number of possibilities, an accumulation of previously evaluated game openings and endings, good adjustment of position scores, and especially a ten-million-fold increase in the number of alternative move sequences the machines can explore. Note that chess machines reached world champion performance as their (specialized) processing power reached about 1/30 human, by our brain to computer measure. Since it is plausible that Garry Kasparov (but hardly anyone else) can apply his brainpower to the problems of chess with an efficiency of 1/30, the result supports that retina-based extrapolation. In coming decades, as general-purpose computer power grows beyond Deep Blue's specialized strength, machines will begin to match humans in more common skills.

The Great Flood

Computers are universal machines, their potential extends uniformly over a boundless expanse of tasks. Human potentials, on the other hand, are strong in areas long important for survival, but weak in things far removed. Imagine a "landscape of human competence," having lowlands with labels like "arithmetic" and "rote memorization", foothills like "theorem proving" and "chess playing," and high mountain peaks labeled "locomotion," "hand-eye coordination" and "social interaction." We all live in the solid mountaintops, but it takes great effort to reach the rest of the terrain, and only a few of us work each patch.

Advancing computer performance is like water slowly flooding the landscape. A half century ago it began to drown the lowlands, driving out human calculators and record clerks, but leaving most of us dry. Now the flood has reached the foothills, and our outposts there are contemplating retreat. We feel safe on our peaks, but, at the present rate, those too will be submerged within another half century. I propose (Moravec 1998)

that we build Arks as that day nears, and adopt a seafaring life! For now, though, we must rely on our representatives in the lowlands to tell us what water is really like.

Our representatives on the foothills of chess and theorem-proving report signs of intelligence. Why didn't we get similar reports decades before, from the lowlands, as computers surpassed humans in arithmetic and rote memorization? Actually, we did, at the time. Computers that calculated like thousands of mathematicians were hailed as "giant brains," and inspired the first generation of AI research. After all, the machines were doing something beyond any animal, that needed human intelligence, concentration and years of training. But it is hard to recapture that magic now. One reason is that computers' demonstrated stupidity in other areas biases our judgment. Another relates to our own ineptitude. We do arithmetic or keep records so painstakingly and externally, that the small mechanical steps in a long calculation are obvious, while the big picture often escapes us. Like Deep Blue's builders, we see the process too much from the inside to appreciate the subtlety that it may have on the outside. But there is a non-obviousness in snowstorms or tornadoes that emerge from the repetitive arithmetic of weather simulations, or in rippling tyrannosaur skin from movie animation calculations. We rarely call it intelligence, but "artificial reality" may be an even more profound concept than artificial intelligence (Moravec 1998).

The mental steps underlying good human chess playing and theorem proving are complex and hidden, putting a mechanical interpretation out of reach. Those who can follow the play naturally describe it instead in mentalistic language, using terms like strategy, understanding and creativity. When a machine manages to be simultaneously meaningful and surprising in the same rich way, it too compels a mentalistic interpretation. Of course, somewhere behind the scenes, there are programmers who, in principle, have a mechanical interpretation. But even for them, that interpretation loses its grip as the working program fills its memory with details too voluminous for them to grasp.

As the rising flood reaches more populated heights, machines will begin to do well in areas a greater number can appreciate. The visceral sense of a thinking presence in machinery will become increasingly widespread. When the highest peaks are covered, there will be machines than can interact as intelligently as any human on any subject. The presence of minds in machines will then become self-evident.

REFERENCES

Goldhaber-Gordon, D. J. et al. (1997) "Overview of Nanoelectronic Devices", *Proceedings of the IEEE*, April 1997.

Moravec, H. (1998) *Robot, Being: mere machine to transcendent mind*. (forthcoming) Oxford University Press.

[Appendix](#)

[Open peer commentary](#)

[Journal of Transhumanism](#)