# The Senses Have no Future

*Hans Moravec*
*Robotics Institute*
*Carnegie Mellon University*
*Pittsburgh, PA 15213*
*hpm@cmu.edu*

February 1997

**Abstract**

*Senses evolved to when the world was wild, enabling our ancestors to detect subtle passing opportunities and dangers. Senses are less useful in a tamer world, where our interactions become more and more simple information exchanges. Senses, and the instincts using them, are increasingly liabilities, demanding entertainment rather than providing useful services. The anachronism will become more apparent as virtual realities, prosthetic sense organs and brain to computer interfaces become common. Imagine reading a computer screen if your eyes and visual cortex are artificial prostheses. It would be far better to bypass all the sensory processing, and insert the message from the computer directly into the thinking portions of your brain. In such manner all our senses will become obsolete, as our physical environment is inexorably refined from a rough physical place into a densely interconnected cyberspace.*

---

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--connected by a million-fiber cable--the optic nerve--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 image-processing functions seem straightforward compared with the brain's other mysteries. A human retina contains about 100 million neurons, of five distinct kinds. Light-sensitive cells feed *horizontal* cells and *bipolar* cells, which connect to *amacrine* cells, whose output goes to *ganglion* cells, whose outgoing fibers bundle to form the optic nerve. Each of the million ganglion-cell axons carries signals from a particular patch of image, representing the differences in light intensity between adjacent regions and from one time to the next--edge and motion detections that are useful also in robot vision. Overall, the retina seems to resolve about one million distinct regions in the visual field and to follow change up to about ten frames per second. Fed a video image with similar resolution, it takes a robot vision program about 100 computer operations to produce a single edge or motion detection, thus 100 million operations to match a whole "frame" of optic nerve output, and 1,000 MIPS--the power of a small supercomputer--to equal the retina's ten frames per second.

If the retina is worth 1,000 MIPS, what about the whole brain, whose larger neurons are 1,000 times as numerous, but occupy 100,000 times the volume? Multiplying the retina's computation by a compromise brain/retina ratio of 10,000 yields a rough brain equivalent of 10 million MIPS--like a million 1997 robot computers, or 100 of the biggest supercomputers. Conversely a 10 MIPS robot--like most still in use--has the mental power of a million-neuron bee. An advanced experimental robot, with 100 MIPS, matches the brain of a very small fish. The Figure 2 rates other entities. In fact, the narrowly competent performace of advanced industrial robots that do intricate assembly of electronics, and of experimental robots that drive the autobahns, has the character of a small animal. Technological development has taken us from the equivalent of single neurons to this stage in about 70 years. It took natural evolution about 700 million years to go as far - evolving humans from there required a few hundred million more. By analogy it should take technology a few decades to cover the remaining distance. Computer progress supports this timescale.

Computers have doubled in capacity every two years since 1950, a pace that has becom an industry given. The curve steepened in the 1990s, to doubling each year, as seen in Figure 2. The universal factor in improving computation has been miniaturization: smaller components have less inertia and operate more quickly with less power, and more of them can exist in a given space. Microprocessors in 1997 contain about 10 million components, but manufacturers have exhibited memory chips with a billion devices. As components shrink to atomic scales, it is possible to imagine two-dimensional chips with a trillion components, and three-dimensional arrays with a million trillion. Such numbers take us far beyond the paltry 10 million MIPS required for a human-capable robot. The--probably conservative--assumption that computer power will continue to grow at its historical rate predicts that 10 million MIPS personal computers will arrive by 2030. Giving the robotics industry a few years to get its software into shape, suggests the advent of human-like robots soon after.

**FIGURE 1:** Natural and Artificial Thinkers

**FIGURE 2:** Faster-than-exponential Growth in Computing Power

As intelligent robots design successive generations of successors, technical evolution will go into overdrive. Biological humans can either adopt the fabulous mechanisms of robots, thus becoming robots themselves, or they can retire into obscurity. A robot ecology will colonize space with intelligent machines optimized to live there. Yet, viewed from a distance, robot expansion into the cosmos will be a vigorous physical affair, a wavefront that converts raw inanimate matter into mechanisms for further expansion. It will leave in its ever-growing wake a more subtle world, with less action and more thought.

On the frontier, robots of ever increasing mental and physical ability will compete with one another in a boundless land rush. Behind the expansion wavefront, a surround of established neighbors will restrain growth, and the contest will become one of boundary pressure, infiltration and persuasion: a battle of wits. A robot with superior knowledge of matter may encroach on a neighbor's space through force, threat, or convincing promises about the benefits of merger. A robot with superior models of mind might lace attractive gifts of useful information with subtle slants that subvert others to its purposes. Almost always, the more powerful minds will have the advantage.

To stay competitive, robots will have to grow in place, repeatedly restructuring the stuff of their bounded bodies into more refined and effective forms. Inert lumps of matter, along with limbs and sense organs, will be converted into computing elements, whose components will be then miniaturized to increase their number and speed. Physical activity will gradually transform itself into a web of increasingly pure thought, where every smallest action is a meaningful computation. We cannot guess the mechanisms robots will use, since physical theory has not yet found even the exact rules underlying matter and space. Having found the rules, robots may use their prodigious minds to devise highly improbable organizations that are to familiar elementary particles as knitted sweaters are to tangled balls of yarn.

As they arrange space time and energy into forms best for computation, robots will use mathematical insights to optimize and compress the computations themselves. Every consequent increase in their mental powers will accelerate future gains, and the inhabited portions of the universe will be rapidly transformed into a cyberspace, where overt physical activity is imperceptible, but the world inside the computation is astronomically rich. Beings will cease to be defined by their physical geographic boundaries, but will establish, extend and defend identities as informational transactions in the cyberspace. The old bodies of individual robots, refined into matrices for cyberspace, will interconnect, and the minds of robots, as pure software, will migrate among them at will. As the cyberspace becomes more potent, its advantage over physical bodies will overwhelm even on the raw expansion frontier. The robot wavefront of coarse physical transformation will be overtaken by a faster wave of cyberspace conversion, the whole becoming finally a bubble of Mind expanding at near lightspeed.

## State of Mind

The cyberspace will be inhabited by transformed robots, moving and growing with a freedom impossible for physical entities. A good, or merely convincing, idea, or an entire personality, may spread to neighbors at the speed of light. Boundaries of personal identity will be very fluid, and ultimately arbitrary and subjective, as strong and weak interconnections between different regions rapidly form and dissolve. Yet some boundaries will persist, due to distance, incompatible ways of thought, and deliberate choice. The consequent competitive diversity will allow a Darwinian evolution to continue, weeding out ineffective ways of thought, and fostering a continuing novelty.

Computational speedups will extend the amount of future available to cyberspace inhabitants, because they cram more events into a given physical time, but will have only a subtle effect on immediate existence, since everything, inside and outside the individual, will be equally accelerated. Distant correspondents, however, will seem even more distant, since more thoughts will transpire in the unaltered transit time for lightspeed messages. Also, as information storage is made more efficient through both denser utilization of matter and more efficient encodings, there will be increasingly more cyber-stuff between any two points. The overall effect of improvements in computational efficiency is to increase the effective space, time and material available, that is, to expand the universe.

Because it uses resources more efficiently, a mature cyberspace will be effectively *much* bigger and longer lasting than the raw spacetime it displaces. Only an infinitesimal fraction of normal matter does work of interest to thinking beings, but in a well-developed cyberspace every bit will be part of a relevant computation or storing a significant datum. The advantage will grow as more compact and faster ways of using space and matter are invented. Today we take pride in storing information as densely as one bit per atom, but it is possible to do much better by converting an atom's mass into many low-energy photons, each storing a separate bit. As the photons' energies are reduced, more of them can be created, but their wavelength, and thus the space they occupy and the time to access them rises, while the temperature they can tolerate drops. A very general quantum mechanical calculation in this spirit by Bekenstein concludes that the maximum amount of information stored in (or fully describing) a sphere of matter is proportional to the mass of the sphere times its radius, hugely scaled. The "Bekenstein bound" leaves room for a million bits in a hydrogen atom, 10^16 in a virus, 10^45 in a human being, 10^75 for the earth, 10^86 in the solar system, 10^106 for the galaxy, and 10^122 in the visible universe.

The computer to brain comparison above suggests that a human brain could be encoded in less than 10^15 bits. If it takes a thousand times more storage to encode a body and surrounding environment, a human with living space might consume 10^18 bits, and a large city of a million human-scale inhabitants might be efficiently stored in 10^24 bits, and the entire existing world population would fit in 10^28. Thus, in an ultimate cyberspace, the 10^45 bits of a single human body could contain the efficiently-encoded biospheres of a thousand galaxies--or a quadrillion individuals each with a quadrillion times the capacity of a human mind.

Because it will be so more capacious than the conventional space it displaces, the expanding bubble of cyberspace can easily recreate internally everything of interest it encounters, memorizing the old universe as it consumes it. Traveling as fast as any warning message, it will absorb astronomical oddities, geologic wonders, ancient Voyager spacecraft, early robots in outbound starships and entire alien biospheres. Those entities may continue to live and grow as if nothing had happened, oblivious of their new status as simulations in the cyberspace--living memories in unimaginably powerful minds, more secure in their existence, and with more future than ever before, because they have become valued parts of such powerful patrons.

Earth, at the center of the expansion, can hardly escape the transformation. The conservative, somewhat backward, robots defending Earth from unpredictable robots will be helpless against a wave that subverts their very substance. Perhaps they will continue, as simulations defending a simulated Earth of simulated biological humans--in one of many, many different stories that plays itself out in the vast and fertile minds of our ethereal grandchildren.

The scenarios absorbed in the cyberspace expansion will provide not only starting points for unimaginably many tales about possible futures, but an astronomically voluminous archeological record from which to infer the past. Minds somewhere intermediate between Sherlock Holmes and God will process clues in solar-system quantities to deduce and recreate the most microscopic details of the preceding eras. Entire world histories, with all their living, feeling inhabitants, will be resurrected in cyberspace. Geologic ages, historical periods and individual lifetimes will recur again and again as parts of larger mental efforts, in faithful renditions, in artistic variations, and in completely fictionalized forms.

The Minds will be so vast and enduring, that rare infinitesimal flickers of interest by them in the human past will ensure that our entire history is replayed astronomically many times, in many places and many, many variations. Single original events will be very rare compared to the indefinitely many cyberspace replays. Most things that are experienced--this very moment, for instance, or your entire life--are far more likely to be a Mind's musings than the physical processes they seem to be. There is no way to tell for sure, and the suspicion that we are someone else's thought does not free us from the burdens of life: to a simulated entity, the simulation is reality, and must be lived by its internal rules.

## Pigs in Cyberspace?

Might an adventurous human mind escape from a bit role in a cyber deity's thoughts, to eke out an independent life among the mental behemoths of a mature cyberspace? We approach the question by extrapolating existing possibilities.

Telepresence and virtual reality are in the news. Today's pioneering systems give crude peeks into remote and simulated worlds, but maturing technology will improve the fidelity. Imagine a well-developed version of the near future: you are cocooned in a harness that, with optical, acoustical, mechanical, chemical and electrical devices drives all your senses, and measures all of your actions. The machinery presents pictures to your eyes, sounds to your ears, pressures and temperatures to your skin, forces to your muscles and even smells and tastes to your nose and mouth. Telepresence results when these inputs and outputs are relayed to a distant humanoid robot. Images from the robot's two camera eyes appear on your eyeglass viewscreens, sound from its microphones is heard in your earphones, contacts on your skin allow you to feel through its instrumented surface and smell and taste through its chemical sensors. Motions of your body cause the robot to move in exact synchrony. When you reach for something in the viewscreens, the robot grasps it, and relays to your muscles and skin the resulting weight, shape, texture and temperature, creating the perfect illusion that you inhabit the robot's body. Your sense of consciousness seems to have migrated to the robot's location, in a true "out of body" experience.

Virtual reality uses a telepresence harness, but substitutes a computer simulation for the remote robot. When connected to a virtual reality, where you are and what you see and touch do not exist in the usual physical sense, but are a kind of computer-generated dream. Like human dreams, virtual realities may contain elements from the outside world, for instance representations of other physical people connected via their own harnesses, or even real views, perhaps through simulated windows. Imagine a hybrid travel system, where a virtual "central station" is surrounded by portals with views of various physical locations. While in the station one inhabits a simulated body, but as one steps through a portal, the harness link switches seamlessly to a physical telepresence robot waiting at that location.

Linked realities are crude toys today, but driven by rapidly advancing computer and communications technologies. In a few decades people may spend more time linked than experiencing their dull immediate surroundings, just as today most of us spend more time in artificial indoor settings than in the uncomfortable outdoors. Linked realities will routinely transcend the physical and sensory limitations of the "home" body. As those limitations become more severe with age, we might compensate by turning up a kind of volume control, as with a hearing aid. When hearing aids at any volume are insufficient, it is now possible to install electronic cochlear implants that stimulate auditory nerves directly. Similarly, on a grander scale, aging users of remote bodies may opt to bypass atrophied muscles and dimmed senses, and connect sensory and motor nerves directly to electronic interfaces. Direct neural interfaces would make most of the harness hardware unnecessary, along with sense organs and muscles, and indeed the bulk of the body. The home body might be lost, but remote and virtual experiences could become more real than ever.

Picture a "brain in a vat," sustained by life-support machinery, connected by wonderful electronic links to a series of artificial rent-a-bodies in remote locations, and to simulated bodies in virtual realities. Though it may be nudged far beyond its natural lifespan by an optimal physical environment, a biological brain built to operate for a human lifetime is unlikely to function effectively forever. Why not use advanced neurological electronics like that which links it with the external world, to replace the gray matter as it begins to fail? Bit by bit our failing brain may be replaced by superior electronic equivalents, leaving our personality and thoughts clearer than ever, though, in time, no vestige of our original body or brain remains. The vat, like the harness before it, will have been rendered obsolete, while our thoughts and awareness continue. Our mind will have been transplanted from our original biological brain into artificial hardware. Transplantation to yet other hardware should be trivial in comparison. Like programs and data that can be transferred between computers without disrupting the processes they represent, our essences will become patterns that can migrate the information networks at will. Time and space will be more flexible--when our mind resides in very fast hardware, one second of real time may provide a subjective year of thinking time, while a thousand years spent on a passive storage medium will seem like no time at all. The very components of our minds will follow our sense of awareness in shifting from place to place at the speed of communication. We might find ourselves distributed over many locations, one piece of our mind here, another piece there, and our sense of awareness yet elsewhere , in what can no longer be called an out-of-body experience, for lack of a body to be out of. And, yet, we will not be truly disembodied minds.

Humans need a sense of body. After twelve hours in a sensory-deprivation tank, floating in a totally dark, quiet, contactless, odorless, tasteless, body-temperature saline solution, a person begins to hallucinate, as the mind, like a television displaying snow on an empty channel, turns up the amplification in search of a signal, becoming ever less discriminating in the interpretations it makes of random sensory hiss. To remain sane, a transplanted mind will require a consistent sensory and motor image, derived from a body, or from a simulation. Transplanted human minds will often be without physical bodies, but hardly ever without the illusion of having them.

Computers already contain many non-human entities that resemble truly bodiless minds. A typical computer chess program knows nothing about physical chess pieces or chessboards, or about the staring eyes of its opponent or the bright lights of a tournament, nor does it work with an internal simulation of those physical attributes. It reasons, instead, with a very efficient and compact mathematical representation of chess positions and moves. For the benefit of human players, this internal representation may be interpreted into a graphic on a computer screen, but such images mean nothing to the program that actually chooses the chess moves. The chess program's thoughts and sensations--its consciousness--is pure chess, uncomplicated by physical considerations. Unlike a transplanted human mind requiring a simulated body, a chess program is pure mind.

Minds in a mature, teeming, competitive cyberspace will be optimally configured to make their living there. Only successful enterprises will be able to afford the storage and computational essentials of life. Some may do the equivalent of construction, converting undeveloped parts of the universe into cyberspace, or improving the performance of existing patches, thus creating new wealth. Others may devise mathematical, physical or engineering solutions that give the developers new and better ways to construct computing capacity. Some may create programs that others can incorporate into mental repertoire. There will be niches for agents, who collect commissions for locating opportunities and negotiating deals for clients, and for banks, storing and redistributing resources, buying and selling computing space, time and information. Some mental creations will be like art, having value only because of changeable idiosyncrasies in their customers. Entities who fail to support their operating costs will eventually shrink and disappear, or merge with other ventures. Those who succeed will grow. The closest present-day parallel is the growth, evolution, fragmentation and consolidation of corporations, who plan their future, but whose options are shaped primarily by the marketplace.

A human would likely fare poorly in such a cyberspace. Unlike the streamlined artificial intelligences that zip about, making discoveries and deals, rapidly reconfiguring themselves to efficiently handle changing data, a human mind would lumber about in a massively inappropriate body simulation, like a deep-sea diver plodding through a troupe of acrobatic dolphins. Every interaction with the world would first be analogized into a recognizable quasi-physical form: other programs might be presented as animals, plants or demons, data items as books or treasure chests, accounting entries as coins or gold. Maintaining the fictions will increase the cost of doing business and decrease responsiveness, as will operating the mind machinery that reduces the physical simulations into mental abstractions in the human mind. Though a few humans may find momentary niches exploiting their baroque construction to produce human-flavored art, most will be compelled to streamline their interface to the cyberspace.

The streamlining could begin by merging processes that analogize the world with those that reduce the resulting simulated sense impressions. The cyber world would still appear as location, color, smell, faces, and so on, but only noticed details would be represented. Since physical intuitions are probably not the best way to deal with most information, humans would still be at a disadvantage to optimized artificial intelligences. Viability might be further increased by replacing some innermost mental processes with cyberspace-appropriate programs purchased from the AIs. By a large number of such substitutions, our thinking procedures might be totally liberated from any traces of our original body. But the bodiless mind that results, wonderful though it may be in its clarity of thought and breadth of understanding, would be hardly human: it will have become an AI.

So, one way or another, the immensities of cyberspace will be teeming with unhuman superminds, engaged in affairs that are to human concerns as ours are to those of bacteria. Memories of the human past will occasionally flash through their minds, as humans once in a long while think of bacteria, and by their thoughts they will recreate us. They could interface us to their realities, making us something like pets, though we would probably be overwhelmed by the experience. More likely, the re-creations would be in the original historical settings, fictional variations, or total fantasies, which would to us seem just like our present existence. Reality or re-creation, there is no way to sort it out from our perspective: we can only wallow in the scenery provided.

---

**Biography for Hans Moravec**

*Hans Moravec has been thinking about machines thinking since he was a child in the 1950s, builing his first robot, a construct of tin cans, batteries, lights and a motor, at age ten. In high school he won two science fair prizes for a light-following electronic turtle and a tape-controlled robot hand. As an undergraduate he designed a computer to control fancier robots, and experimented with learning and* *automatic programming on commercial machines. During his master's work he built a small robot with whiskers and photoelectric eyes controlled by a minicomputer, and wrote a thesis on a computer language for artificial intelligence. He received a PhD from Stanford in 1980 for a TV-equipped robot, remote controlled by a large computer, that negotiated cluttered obstacle courses. Since 1980 he has been director of the Carnegie Mellon University Mobile Robot Laboratory, birthplace of mobile robots deriving 3D spatial awareness from cameras, sonars, and other sensors. His 1988 book, "Mind Children: the future of robot and human intelligence", and forthcomg "Mind Age: transcendence through robots", consider the future prospects for humans, robots and intelligence. He has published many articles in robotics, computer graphics, multiprocessors, space travel and other speculative areas.*