B. EXPANDING HUMAN COGNITION AND COMMUNICATION

THEME B SUMMARY

Panel: W.S. Bainbridge, R. Burger, J. Canton, R. Golledge, R.E. Horn, P. Kuekes, J. Loomis, C.A. Murray, P. Penz, B.M. Pierce, J. Pollack, W. Robinett, J. Spohrer, S. Turkle, L.T. Wilson

In order to chart the most profitable future directions for societal transformation and corresponding scientific research, five multidisciplinary themes focused on major goals have been identified to fulfill the overall motivating vision of convergence described in the previous pages. The first, "Expanding Human Cognition and Communication," is devoted to technological breakthroughs that have the potential to enhance individuals' mental and interaction abilities. Throughout the twentieth century, a number of purely psychological techniques were offered for strengthening human character and personality, but evaluation research has generally failed to confirm the alleged benefits of these methods (Druckman and Bjork 1992; 1994). Today, there is good reason to believe that a combination of methods, drawing upon varied branches of converging science and technology, would be more effective than attempts that rely upon mental training alone.

The convergence of nanotechnology, biotechnology, information technology, and cognitive science could create new scientific methodologies, engineering paradigms, and industrial products that would enhance human mental and interactive abilities. By uniting these disciplines, science would become ready to succeed in a rapid program to understand the structure and functions of the human mind, The Human Cognome Project. Truly, the mind is the final frontier, and unraveling its mysteries will have tremendous practical benefits. Among the most valuable spin-offs will be a host of devices that enhance human sensory capabilities. We will be able to build a vast variety of humane machines that adapt to and reflect the communication styles, social context, and personal needs of the people who use them. We will literally learn how to learn in new and more effective ways, revolutionizing education across the life span. New tools will greatly enhance creativity, industrial design, and personal productivity. Failure to invest in the necessary multidisciplinary research would delay or even prevent these benefits to the economy, to national security, and to individual well-being.

Rapid recent progress in cognitive science and related fields has brought us to the point where we could achieve several breakthroughs that would be of great value to mankind. However, we will need to make a significant special effort to bring together the often widely dispersed scientific and technical disciplines that must contribute. For example, progress in the cognitive neuroscience of the human brain has been achieved through new research methodologies, based in both biology and information science, such as functional magnetic resonance imagining (fMRI) and infrared sensors. However, we are reaching the resolution limits of current instrumentation, for example, because of concerns about the safety of human research subjects (Food and Drug Administration 1998), so progress will stall quickly unless breakthroughs in NBIC can give us research tools with much greater resolution, sensitivity, and capacity to analyze data. Many other examples could be cited in which scientific, technological, and economic progress is approaching a barrier that can be surmounted only by a vigorous program of multidisciplinary research.

The panel identified five main areas in which integration of the NBIC sciences can enhance the cognitive and communicative aspects of human performance. Each of these is a challenging field for multidisciplinary research that will lead to many beneficial applications.

1. The Human Cognome Project

It is time to launch a Human Cognome Project, comparable to the successful Human Genome Project, to chart the structure and functions of the human mind. No project would be more fundamental to progress throughout science and engineering or would require a more complete unification of NBIC sciences. Success in the Human Cognome Project would allow human beings to understand themselves far better than before and therefore would enhance performance in all areas of human life.

While the research would include a complete mapping of the connections in the human brain, it would be far more extensive than neuroscience. The archaeological record indicates that anatomically modern humans existed tens of thousands of years before the earliest examples of art, a fact that suggests that the human mind was not merely the result of brain evolution but also required substantial evolution in culture and personality. Central to the Human Cognome Project would be wholly new kinds of rigorous research on the nature of both culture and personality, in addition to fundamental advances in cognitive science.

The results would revolutionize many fields of human endeavor, including education, mental health, communications, and most of the domains of human activity covered by the social and behavioral sciences. Some participants in the human cognition and communication working group were impressed by the long-term potential for uploading aspects of individual personality to computers and robots, thereby expanding the scope of human experience, action, and longevity. But at the very least, greater understanding of the human mind would allow engineers to design technologies that are well suited to human control and able to accomplish desired goals most effectively and efficiently. Success in the Human Cognome Project would greatly facilitate success in the other four areas identified by this working group.

2. Personal Sensory Device Interfaces

Fundamental scientific and engineering work needs to be done to permit development of an array of personal sensory device interfaces to enhance human abilities to perceive and communicate. Human senses are notoriously limited. Whereas we can hear ten octaves of musical tones, we can see only one octave of the colors of light, and our ears have a poor ability to form detailed “images” from sound the way our eyes can with light. Today’s communication technology has revolutionized the ability of people to communicate across large distances, but little has been done to help with small area communication, for example, between individuals in a conference room. These are only two of many areas where NBIC sensor efforts can increase human performance.

Research can develop high bandwidth interfaces between devices and the human nervous system, sensory substitution techniques that transform one type of input (visual, aural, tactile) into another, effective means for storing memory external to the brain, knowledge-based information architectures that facilitate exploration and understanding, and new kinds of sensors that can provide people with valuable data about their social and physical environments. For example, increased awareness of the chemical composition of things in our immediate environment will improve human productivity, health, and security. Artificial agents based in microelectronics, nanotechnology, and bioengineering may endow people with entirely new senses or existing senses operating in new ways, in some cases employing neural interfaces to deliver complex information directly into the human mind.

3. Enriched Community

Enlightened exploitation of discoveries in the NBIC sciences will humanize technology rather than dehumanize society. Robots, intelligent agents, and information systems need to be sensitive to human needs, which is another way of saying that they must to some extent embody human personality. Over the next two decades, as nanotechnology facilitates rapid improvement of microelectronics, personal digital assistants (PDAs) are likely to evolve into smart portals to a whole world of information sources, acting as context aware personal brokers interacting with other systems maintained by corporations, governments, educational institutions, and individuals. Today's email and conference call systems could evolve into multi-media telepresence communication environments. Global Positioning System (GPS) units could become comprehensive guides to the individual's surroundings, telling the person his or her location and also locating everything of interest in the immediate locale.

To accomplish these practical human goals, we must invest in fundamental research on how to translate human needs, feelings, beliefs, attitudes, and values into forms that can guide the myriad devices and embedded systems that will be our artificial servants of the future. We must understand how interacting with and through machines will affect our own sense of personhood as we create ever more personable machines. As they become subtle reflections of ourselves, these technologies will translate information between people who are separated by perspective, interests, and even language. Without the guidance provided by the combined NBIC sciences, technology will fail to achieve its potential for human benefit. Multidisciplinary research to humanize computing and communications technology will expand the social competence of individuals and increase the practical effectiveness of groups, social networks, and organizations.

4. Learning How to Learn

We need to explore fresh instructional approaches, based in the NBIC sciences, to help us learn how to learn. Such educational tools as interactive multimedia, graphical simulations, and game-like virtual reality will enhance learning not merely from kindergarten through graduate school but also throughout the entire life course in school, in corporations, and at home. The results of past efforts have often been disappointing, because they failed to draw upon a sufficiently broad and deep scientific base. For example, instructional software typically lacked a firm grounding in the findings of cognitive science about how people actually think and learn (Bransford, Brown, and Cocking 1999).

In the future, everyone will need to learn new skills and fundamental knowledge throughout life, often in fields connected to mathematics, engineering, and the sciences. Thus we will need new kinds of curricula, such as interactive virtual reality simulations run over the Internet that will allow a student anywhere to experience the metabolic processes that take place within a living cell, as if seeing them from a nanoscale perspective. New, dynamic ways to represent mathematical logic could be developed based on a correct understanding of how the human mind processes concepts like quantity and implication, allowing more people to learn mathematics more quickly, thoroughly, and insightfully. The social interaction resulting from multiuser video games can be harnessed as a strong learning motivator, if they are designed for the user's demographic and cultural background and can infuse the learning with mystery, action, and drama. The goal would be to revolutionize science, mathematics, and engineering education through experiences that are emotionally exciting, substantively realistic, and based on accurate cognitive science knowledge about how and why people learn.

5. Enhanced Tools for Creativity

As technology becomes ever more complex, engineering design becomes an increasingly difficult challenge. For example, it is extremely costly to create large software systems, and the major bottlenecks reducing their effectiveness are unreliability and inefficiency. Similar problems beset systems for large-scale organization administration, supply chain management, industrial design, mass media, and government policy making. We can anticipate that future industries in biotechnology and nanotechnology will present unprecedented design challenges.

Investment in research and development of wholly new industrial design methods will pay great dividends. Among these, biologically inspired techniques, such as evolutionary design methods analogous to genetic algorithms, are especially promising. Terascale and petascale computer simulations are excellent approaches for many design problems, but for the foreseeable future the cost of creating a facility to do such work would be prohibitive for universities and most companies. Therefore, a national center should be established for high-end engineering design simulations. This facility could be linked to a network of users and specialized facilities, providing a distributed design environment for advanced research in engineering. Good models for creating the National Center for Engineering Design would be the supercomputer networks established by the National Science Foundation: the National Computational Science Alliance, the National Partnership for Advanced Computational Infrastructure, and the new Terascale Computing System.

At the same time, radically new methods would enhance small-scale design activities by a wide range of individuals and teams in such fields as commercial art, entertainment, architecture, and product innovation. New developments in such areas as visual language, personalized design, designing around defects, and the cognitive science of engineering could be extremely valuable. Breakthroughs in design could become self-reinforcing, as they energize the economic and technical feedback loops that produce rapid scientific and technological progress.

Statements and Visions

Participants in the human cognition and communication panel contributed a number of *statements*, describing the current situation and suggesting strategies for building upon it, as well as transformative *visions* of what could be accomplished in 10 or 20 years through a concentrated effort. The contributions include statements about societal opportunities and challenges, sensory systems, networking architecture, spatial cognition, visual language, and "companion" computers, as well as visions on predicting social behavior, design complexity, enhancing personal area sensing, understanding the brain, stimulating innovation and accelerating technological convergence.

References

Bransford, J.D., A.L. Brown, and R.R. Cocking, eds. 1999. *How people learn: Brain, mind, experience, and school.* Washington D.C.: National Research Council.

Druckman, D. and R.A. Bjork, eds. 1992. *In the mind's eye: Enhancing human performance*. Washington, D.C.: National Research Council.

_____. 1994. *Learning, remembering, believing: Enhancing human oerformance*. Washington, D.C.: National Research Council.

Food and Drug Administration. 1998. Guidance for the submission of premarket notifications for magnetic resonance diagnostic devices. U.S. Food and Drug Administration, November 14 1998: http://www.fda.gov/cdrh/ode/mri340.pdf.

Kurzweil, R. 1999. *The age of spiritual machines: When computers exceed human intelligence*. New York: Viking.

STATEMENTS

NBICS (NANO-BIO-INFO-COGNO-SOCIO) CONVERGENCE TO IMPROVE HUMAN PERFORMANCE: OPPORTUNITIES AND CHALLENGES

Jim Spohrer, IBM, CTO Venture Capital Relations, spohrer@us.ibm.com

This paper is an exploration of new opportunities and challenges for improving human performance from the perspective of rapid technological change and convergence. In the past two million years, human performance has primarily been improved in two ways: evolution (physical-cognitive-social changes to people) and technology (human-made artifacts and other changes to the environment). For example, approximately one hundred thousand generations ago, physical-cognitive-social evolution resulted in widespread spoken language communication among our ancestors. About 500 generations ago, early evidence of written language existed. Then the pace of technological progress picked up: 400 generations ago, libraries existed; 40 generations ago, universities appeared; and 24 generations ago, printing of language began to spread. Again, the pace of technological advancements picked up: 16 generations ago, accurate clocks appeared that were suitable for accurate global navigation; five generations ago, telephones were in use; four, radios; three, television; two, computers; and one generation ago, the Internet.

In the next century (or in about five more generations), breakthroughs in nanotechnology (blurring the boundaries between natural and human-made molecular systems), information sciences (leading to more autonomous, intelligent machines), biosciences or life sciences (extending human life with genomics and proteomics), cognitive and neural sciences (creating artificial neural nets and decoding the human cognome), and social sciences (understanding “memes“ and harnessing collective IQ) are poised to further pick up the pace of technological progress and perhaps change our species again in as profound a way as the first spoken language learning did some one hundred thousand generations ago. NBICS (nano-bio-info-cogno-socio) technology convergence has the potential to be the driver of great change for humankind. Whether or not this is in fact desirable, reasoned speculation as to how this may come to pass and the threats posed by allowing it to come to pass are increasingly available from futurists. Currently, this technology road of human performance augmentations is at the stage of macroscopic external human-computer interfaces tied into large social networking systems that exist today. Recently, there are the tantalizing first experiments of microscopic internal interfaces to assist the elderly or others with special needs; and then there is the further speculative road, with potentially insurmountable obstacles by today’s standards, that leads to the interfaces of the future.

After setting the stage with longer term visions and imaginings, this paper will focus on the nearer term opportunities and challenges afforded by NBICS research and development (R&D) over the next half a generation or so. In conclusion, while futurists may be overestimating the desirability and feasibility of achieving many of their visions, we are probably collectively underestimating the impact of many of the smaller technological steps along the way.

Introduction: Motivations and Goals

At the beginning of the NBIC workshop, the participants were challenged by Newt Gingrich to think outside the box and to ambitiously consider the possible implications of the nano-info-bio-cogno convergence over the coming decades. We were also instructed to consider human dignity as an important issue, which tempered some of the cyborg speculations and other visions of humans with technology implants and augments that might seem unappealing to most people today. Thus, while social norms can shift significantly over several generations, we were primarily concerned with the world of our children and our own old-age years. We were also treated to a number of presentations describing state-of-the-art results in areas such as nanotechnology; learning technology; social acceptance of technology; designer drugs to combat diseases and other degenerative conditions; neurological implants; advanced aircraft designs highlighting smart, polymorphic (shape-shifting) materials; reports on aging, blindness, and other challenges; evolutionary software and robots; the needs of the defense department for the military of the future; augmented reality and virtual reality; and other useful perspectives on the topic of augmenting human performance. While it would be well beyond the scope of this paper to try to summarize all of these perspectives, I have tried to integrate ideas from these presentations into my own thinking about nano-info-bio-cogno convergence. Additionally, my perspective has been shaped by interactions with Doug Engelbart, whose pioneering work in the area of human augmentation systems stresses the importance of the co-evolution of technological and social systems. Because the social sciences will strongly influence which paths humans will ultimately explore as well as help us understand why, we are really concerned here with nano-bio-info-cogno-socio convergence.

Nano-bio-info-cogno-socio convergence assumes tremendous advances in each of the component science and technology areas:

1. Nanoscience advances in the coming decade will likely set the stage for a new generation of material science, biochemistry, and molecular electronics, as well as of new tools for measuring and manipulating the world at the level of individual atoms and molecules. Nanotechnology advances are poised to give humans the capabilities that bacteria have had for billions of years, the ability to create molecular machines that solve a wide range of problems on a global scale. Ultimately, these advancements will blur the distinction between natural and human-made objects.
2. Bioscience or life sciences will expand the mapping of the human genome to the human proteome, leveraging both to create new drugs and therapies to address a host of maladies of the past and new threats on the horizon.
3. Information science advances will find many applications in the ongoing e-business transformation already underway, as well as pervasive communication and knowledge management tools to empower individuals. More importantly, information science will provide both the interlingua to knit the other technologies together and the raw computational power needed to store and manipulate mountains of new knowledge.
4. Cognitive science and neuroscience will continue to advance our understanding of the human information processing system and the way our brains work.
5. Social science advances (obtained from studies of real systems as well as simulations of complex adaptive systems composed of many interacting individuals) will provide fresh insights into the collective IQ of humans, as well as interspecies collective IQ and the spread of memes. A meme, which is a term coined by the author and zoologist Richard Dawkins, is “a habit, a technique, a twist of feeling, a sense of things, which easily flips from one brain to another.” It is no coincidence that meme rhymes with gene, for one is about replicating ideas (from one brain to another brain) and the other is about replicating molecules (from one cell to another cell).
6. Thus, the central question of this paper is “how might the convergence of nano-bio-info-cogno-socio technologies be accomplished and used to improve human performance” or, in the words of one workshop participant, Sandia National Laboratory scientist Gerry Yonas, to “make us all healthier, wealthier, and wiser”?
7. To gain some traction on this question, a framework, here termed simply the Outside-Inside Framework, is proposed in the next section. This framework makes explicit four of the key ways that new technologies might be used to augment human performance: (a) outside the body (environmental); (b) outside the body (personal); (c) inside the body (temporary); (d) inside the body (permanent). This framework will be shown to be largely about how and where information is encoded and exchanged: (i) info: bits and the digital

environment, (ii) cogno-socio: brains and memes and the social environments, (iii) nano-bio: bacteria and genes and the bioenvironment, (iv) nano-cogno: bulk atoms, designed artifacts, and the physical environments. In conclusion, near-term implications of NBICS technology convergence will be discussed.

The Outside-Inside Framework and Future Imaginings

The Outside-Inside framework consists of four categories of human performance-enhancing technologies:

- Outside the body and environmental
- Outside the body and personal
- Inside the body and temporary
- Inside the body and permanent

In this section, while briefly describing the categories and subcategories, some extremely speculative visions of the future will be discussed to help stretch our imaginations before “coming back to earth” in the last section to discuss more practical and near term possibilities. Readers are encouraged to view this section as a number of imagination challenges and to create their own answers to questions like what new materials, agents, places, mediators, ingestibles, senses, and species might come to be in the next few decades. In the true spirit of brainstorming, anything goes in this section. Also, it is worth noting that while futurists may be overestimating the desirability and feasibility of how quickly, if ever, we can achieve many of their visions, we are probably collectively underestimating the impact of many of the smaller technological steps along the way. Finally, as an example of improving human performance, the task of learning will be considered, focusing on the way existing and imaginary technologies may improve our ability to learn and/or perform more intelligently.

Outside the Body and Environmental.

People perform tasks in a variety of environmental contexts or places, such as homes, offices, farms, factories, hotels, banks, schools, churches, restaurants, amusement parks, cars, submarines, aircraft, space stations, and a host of other environments that have been augmented by what is termed here environmental technologies. From the materials that are used to construct the buildings and artifacts at these locations to the agents (people, domesticated animals) that provide services in these locations to the very nature of the places themselves, environmental technologies account for most of the advances in human performance that have occurred in the past 500 generations of recorded history (most of us overlap and therefore experience only about five generations of perspectives from grandparents to grandchildren). For the task of learning, consider the important roles that the three innovations — paper (material), teachers (agents), and schools (places) — have had on education. NBICS convergence will surely lead to new materials, new agents, and new places.

Outside the body and environmental: Materials. We expect that the progression from rocks, wood, bricks, cloth, ceramics, glass, bronze, iron, cement, paper, steel, rubber, plastic, semiconductors, and so on, will be augmented with new materials, such as smart, chromatically active (change color), polymorphic (change shape) materials such as those NASA is already experimenting with. For a thought-provoking vision of where new materials could lead, the reader is directed to the currently infeasible but intriguing notion of "utility fog" developed by Rutgers computer science professor J. Storrs Hall in the early 1990s. Smaller than dust, "foglets" are speculative tiny interlocking machines that can run "programs" that make collections of billions of them work together to assume any shape, color, and texture, from flowing, clear water to fancy seat belt body suits that appear only when an accident has occurred. If utility fog were a reality, most artifacts could be made invisible until needed, making them quite portable. There would be no need to carry luggage on trips; one could simply create clothes out of utility fog. Materializing objects out of thin air (or fog), while wildly infeasible today, nevertheless provides an interesting springboard for imagining some of the ultimate human-computer interfaces (such as a second skin covering human bodies, eyes, ears, mouth, nose, and skin) that may someday exist. Perhaps these ultimate interfaces might connect us to telerobotic versions of ourselves assembled out of utility fog in distance places.

There are many reasons to be skeptical about utility fog (the Energy budget, for one), but notions like utility fog help us understand the potential of NBICS. For example, multi-cellular organisms provide a vast library of examples of the ways cells can be interlinked and grouped to produce shapes, textures, and macroscopic mechanical structures. Social insects like ants have been observed to interlink to solve problems in their environments. And while I'm unaware of any types of airborne bacteria that can spontaneously cluster into large groups, I suspect that mechanisms that bacteria and slime molds use for connecting in various arrangements may one day allow us to create new kinds of smart materials. Hopefully the notion of utility fog has served its brainstorming purpose of imagination stretching, and there are a number of related but nearer term investigations underway. For example, U.C.-Berkeley professor and microroboticist Kris Pister's Smart Dust and Micromechanical Flying Insect projects are good examples of the state-of-the-art in building microrobots, and as these microrobots get smaller, they may very well pave the way to many exciting new materials.

Outside the body and environmental: Agents. Interacting with intelligent agents, such as other people and other species (e.g., guide dogs), has clear advantages for augmenting human performance. Some of the most important agents we interact with daily are role-specialized people and businesses (organization as agent). The legal process of incorporating a business or nonprofit organization is essentially equivalent to setting up a fictitious person with specialized rights, responsibilities, and capabilities. The notion of new agents was an active area of discussion among the workshop participants: from the implications of digital personae (assumed identities on-line) to artificial intelligence and robotics, as well as the evolution of new types of organizations. The successful entrepreneur and futurist Ray Kurzweil has a website *kurzweilai.net* (see Top KurzweilAI News of 2001) that explores these and other futures and interestingly includes Kurzweil's alter-ego, Ramona!, that has been interviewed by the press to obtain Kurzweil's views on a variety of subjects. Undoubtedly, as technology evolves, more digital cloning of aspects of human interactions will occur. An army of trusted agents that can interact on our behalf has the potential to be very empowering as well as the potential to be quite difficult to update and maintain synchrony with the real you. What happens when a learning agent that is an extension of you becomes more knowledgeable about a subject than you? This is the kind of dilemma that many parents and professors have already faced.

Outside the body and environmental: Places. New places create new opportunities for people. The exploration of the physical world (trade connecting ancient civilization, New World, the Wild West, Antarctica, the oceans, the moon, etc.) and the discovery of new places allows new types of human activities and some previously constrained activities to flourish. For example, the New World enhanced the Puritans' abilities to create the kind of communities they wanted for themselves and their children. Moving beyond the physical world, science fiction writer William Gibson first defined the term *cyberspace*. The free thinking artist and futurist Jaron Lanier, who coined the term *virtual reality*, and many other people have worked to transform the science fiction notion of cyberspace into working virtual reality technologies. Undoubtedly, the digital world will be a place of many possibilities and affordances that can enhance human performance on a wide variety of tasks, including both old, constrained activities as well as new activities. The increasing demand for home game machines and combinatorial design tools used by engineers to explore design possibilities is resulting in rapid advances in the state-of-the-art creation of simulated worlds and places. Furthermore, in the context of learning, inventor and researcher Warren Robinett, who was one of the workshop participants, co-created a project that allows learners to "feel" interactions with simulated molecules and other nanostructures via virtual realities with haptic interfaces. In addition, Brandeis University professor Jordan Pollack, who was also one of the workshop participants, described his team's work in the area of combinatorial design for robot evolution, using new places (simulated worlds) to evolve new agents (robots) and then semi-automatically manifest them as real robots in the real world. Also, it is worth noting that in simulated worlds, new materials, such as utility fog, become much easier to implement or, more accurately, at least emulate.

Outside the Body and Personal

The second major category, personal technologies, are technologies that are outside of the body, but unlike environmental technologies are typically carried or worn by a person to be constantly available. Two of the earliest examples of personal technologies were of course clothing and jewelry, which both arose thousands of generations ago. For hunter gatherers as well as cowboys in the American West, weapons were another form of early personal technology. Also included in this category are money, credit cards, eyeglasses, watches, pens, cell phones, handheld game machines, and PDAs (personal digital assistants). For learners, a number of portable computing and communication devices are available, such as leapfrog, which allows students to prepare for quizzes on chapters from their school textbooks, and graphing calculators from Texas Instruments. Recently, a number of wearable biometric devices have also appeared on the market.

Outside the body and personal: Mediators. Mediators are personal technologies that include cellphones; PDAs; and handheld game machines that connect their users to people, information, and organizations and support a wide range of interactions that enhance human performance. WorldBoard is a vision of an information infrastructure and companion mediator devices for associating information with places. WorldBoard, as originally conceived in my papers in the mid-1990s, can be thought of either as a planetary augmented reality system or a sensory augment that would allow people to perceive information objects associated with locations (e.g., virtual signs and billboards). For example, on a nature walk in a national park a person could use either heads-up display glasses or a cell phone equipped with a display, camera, and GPS (Global Positioning System) to show the names of mountains, trees, and buildings virtually spliced into the scenes displayed on the glasses or cell phone. WorldBoard mediators might be able to provide a pseudo X-ray vision, allowing construction equipment operators to see below the surface to determine the location of underground buried pipes and cables rather than consulting blueprints that might not be available or might be cumbersome to properly orient and align with reality. The slogan of WorldBoard is "putting information in its place" as a first step to contextualizing and making useful the mountains of data being created by the modern day information explosion.

Human-made tools and artifacts are termed mediators, in this paper, because they help externalize knowledge in the environment and mediate the communication of information between people. Two final points are worth making before moving inside the body. First, the author and cognitive scientist Don Norman, in his book *Things that Make Us Smart* provides an excellent, in-depth discussion of the way human-made tools and artifacts augment human performance and intelligence. Furthermore, Norman's website includes a useful article on the seeming inevitability of implants and indeed cyborgs in our future, and why implants will become increasingly accepted over time for a wider and wider range of uses. A second point worth making in the context of mediators is that human performance could be significantly enhanced if people had more will power to achieve the goals that they set for themselves. Willpower enforcers can be achieved in many ways, ranging from the help of other people (e.g., mothers for children) to mediator devices that remove intentionality from the equation and allow multitasking (e.g., FastAbs electric stimulation workout devices).

Inside the Body and Temporary

The third major category, inside the body temporary technologies, includes most medicines (pills) as well as new devices such as the camera that can be swallowed to transmit pictures of a journey through a person's intestines. A number of basic natural human processes seem to align with this category, including inhaling and exhaling air; ingesting food and excreting waste; spreading infections that eventually overcome the body's immune system; as well as altered states of awareness such as sleep, reproduction, pregnancy, and childbirth.

Inside the body and temporary: Ingestibles. Researchers at Lawrence Livermore National Laboratories have used mass spectrometry equipment to help study the way that metabolisms of different people vary in their uptake of certain chemical components in various parts of the body. Eventually, this line of investigation may lead to precisely calibrating the amount of a drug that an individual should take to achieve an optimal benefit from ingesting it. For example, a number of studies show positive effects of mild stimulants, such as coffee, used by subjects who were studying material to be learned, as well as positive effects from being in the appropriate mental and physical states when performing particular tasks. However, equally clear from the data in these studies are indications that too much or too little of a good thing can result in no enhancement or detrimental side effects instead of enhanced performance.

With the exception of an Air Force 2025 study done by the Air University, I have not yet found a reference (besides jokes, science fiction plots, and graduate school quiz questions), to what I suspect is someone‘s ultimate vision of this ingestible enhancements subcategory, namely a learning pill or knowledge pill. Imagine that some day we are able to decode how different brains store information, and one can simply take a custom designed learning pill before going to sleep at night to induce specific learning dreams, and when morning arrives the person’s wetware will have been conditioned or primed with memories of the new information. Staggeringly improbable, I know.

Nevertheless, what if someone could take a pill before falling asleep at night, and awaken in the morning knowing or being conditioned to more rapidly learn how to play, for example, a game like chess? If learning could be accelerated in this manner, every night before going to bed, people would have a “learning nightcap.” Imagine an industry developing around this new learning pill technology. The process at first might require someone spending the time to actually learn something new, and monitoring and measuring specific neurological changes that occur as a result of the learning experience, and then re-encoding that information in molecular machines custom-designed for an individual to attach himself or herself to locations in the brain and interact with the brain to create dream-like patterns of activation that induce time-released learning. Businesses might then assign learning pills to their employees, schools might assign learning pills to their students, soldiers might take learning pills before being sent out on missions (per the Air Force 2025 study that mentioned a “selective knowledge pill”), and families might all take learning pills before heading out on vacations. However, perhaps like steroids, unanticipated side effects could cause more than the intended changes.

What makes the learning pill scenario seem so far-fetched and improbable? Well, first of all, we do not understand much about the way that specific bits of information are encoded in the brain. For example, what changes in my brain (short-term and then long-term memory) occur when I learn that there is a new kind of oak tree called a Live Oak that does not lose its leaves in the winter? Second, we do not know how to monitor the process of encoding information in the brain. Third, different people probably have idiosyncratic variations in the ways their brains encode information, so that one person’s encoding of an event or skill is probably considerably different from another person’s. So how would the sharing work, even if we did know how it was encoded in one person’s brain? Fourth, how do we design so many different molecular machines, and what is the process of interaction for time-released learning? Fifth, exactly how do the molecular machines attach to the right parts of the brain? And how are they powered? We could go on and on, convincing ourselves that this fantasy is about as improbable as any that could possibly be conceived. Nevertheless, imagination-stretching warmups like these are useful to help identify subproblems that may have nearer term partial solutions with significant impacts of their own.

Inside the Body and Permanent

The fourth major category, inside the body permanent technologies, raises the human dignity flag for many people, as negative images of cyborgs from various science fiction fare leap immediately to mind. The science fact and e-life writer Chris O'Malley recently wrote a short overview of this area. Excerpts follow:

> Largely lost in the effort to downsize our digital hardware is the fact that every step forward brings us closer to an era in which computers will routinely reside within us. Fantasy? Hardly. We already implant electronics into the human body. But today's pacemakers, cochlear implants, and the like will seem crude — not to mention huge — in the coming years. And these few instances of electronic intervention will multiply dramatically… The most pervasive, if least exciting, use of inner-body computing is likely to be for monitoring our vital stats (heart rate, blood pressure, and so on) and communicating the same, wirelessly, to a home healthcare station, physician's office, or hospital. But with its ability to warn of imminent heart attacks or maybe even detect early-stage cancers, onboard monitoring will make up in saved lives what it lacks in sex appeal... More sensational will be the use of internal computers to remedy deficiencies of the senses. Blindness will, it seems reasonable to speculate, be cured through the use of electronic sensors — a technology that's already been developed. So, too, will deafness. Someday, computers may be able to mimic precisely the signal that our muscles send to our brain and vice versa, giving new mobility to paralysis victims. Indeed, tiny computers near or inside our central processing unit, the human brain, could prove a cure for conditions such as Alzheimer's, depression, schizophrenia, and mental retardation... Ethical dilemmas will follow, as always...

Inside the body and permanent: New organs (senses and effectors). This subcategory includes replacement organs, such as cochlear implants, retinal implants, and pacemakers, as well as entirely new senses. People come equipped with at least five basic senses: sight, hearing, touch, taste, and smell. Imagine if we were all blind but had the other four senses. We'd design a world optimized for our sightless species, and probably do quite well. If we asked members of that species to design a new sense, what might they suggest? How would they even begin to describe vision and sight? Perhaps they might describe a new sense in terms of echolocation, like a species of bats, that would provide a realtime multipoint model of space in the brain of the individual that could be reasoned to be capable of preventing tripping on things in hostile environments.

In our own case, because of the information explosion our species has created, I suggest that the most valuable sixth sense for our species would be a sense that would allow us to quickly understand, in one big sensory gulp, vast quantities of written information (or even better, information encoded in other people's neural nets). Author Robert Lucky has estimated that all senses give us only about 50 bits per second of information, in the Shannon sense. A new high bandwidth sense might be called a Giant UpLoad Process or the GULP Sense. Imagine a sixth sense that would allow us to take a book and gulp it down, so that the information in the book was suddenly part of our wetware, ready for inferencing, reference, etc., with some residual sense of the whole, as part of the sensory gulp experience. Just as some AI programs load ontologies and rules, the GULP sense would allow for rapid knowledge uptake. A GULP sense would have a result not unlike the imaginary learning pill above. What makes the information-gulping sixth sense and the learning pill seem so fantastic has to do in part with how difficult it is for us to transform information encoded in one format for one set of processes into information encoded in another format for a different set of processes — especially when one of those formats is idiosyncratic human encoding of information in our brains. Perhaps the closest analogy today to the complexity of transforming information in one encoding to another is the ongoing transformation of businesses into e-businesses, which requires linking idiosyncratic legacy systems in one company to state-of-the-art information systems in another company.

The process of creating new sensory organs that work in tandem with our own brains is truly in a nascent state, though the cochlear implant and retinal implant directions seem promising. University of Texas researcher Larry Cauller, who was one of the workshop participants, grabbed the bull by the horns and discussed ways to attack the problem of building an artificial brain as well as recent technology improvements in the area of direct neural interfaces. As neural interface chips get smaller, with finer and more numerous pins, and leveraging RF ID tag technology advances, the day is rapidly approaching where these types of implants can be done in a way that does minimal damage to a brain receiving a modern neural interface implant chip. Improved neural interface chips are apparently already paying dividends in deepening the understanding of the so-called mirror neurons that are tied in with the “monkey see, monkey do” behaviors familiar in higher primates. One final point on this somewhat uncomfortable topic, MIT researcher and author Sherry Turkle, who was also a workshop participant, presented a wealth of information on the topic of sociable technologies as well as empirical data concerning people’s attitudes about different technologies. While much of the discussion centered on the human acceptance of new agents such as household entertainment robots (e.g., Sony’s AIBO dog), there was unanimous agreement among all the participants that as certain NBICS technologies find their way into more universally available products, attitudes will be shaped, positively as well as negatively, and evolve rapidly, often in unexpected ways for unexpected reasons.

Tokyo University’s Professor Isao Shimoyama has created a robo-roach or cyborg roach that can be controlled with the same kind of remote that children use to control radio-controlled cars. Neural interfaces to insects are still crude, as can be seen by going to Google and searching for images of “robo-roach.” Nevertheless, projecting the miniaturization of devices that will be possible over the next decade, one can imagine tools that will help us understand the behaviors of other species at a fine level of detail. Ultimately, as our ability to rapidly map genes improves, neural interface tools may even be valuable for studying the relationship between genes and behaviors in various species. NBICS convergence will accelerate as the linkages between genes, cellular development, nervous systems, and behavior are mapped.

Inside the body and permanent: New skills (new uses of old sensors and effectors). Senses allow us to extract information from the world, exchange information between individuals, and encode and remember relevant aspects of the information in our brains (neural networks, wetware). Sometimes physical, cognitive, and social evolution of a species allows an old sense to be used in a new way. Take, for example, verbal language communication or speech. Long before our ancestors could effectively listen to and understand spoken language, they could hear. A lion crashing through the jungle at them registered a sound pattern in their prehuman brains and caused action. However, over time, a set of sound associations with meaning and abstractions, as well as an ability to create sounds, along with increased brain capacity for creating associations with symbols and stringing them together via grammars to create complex spoken languages, occurred. Over time, large groups of people shared and evolved language to include more sounds and more symbolic, abstract representations of things, events, and feelings in their world. An important point about acquiring new skills, such as sounds in a language, is that infants and young children have certain advantages. Evidence indicates that brains come prewired at the neural level for many more possibilities than actually get used, and if those connections are not needed, they go away. Once the connections go away, learning can still occur, but the infant brain advantage is no longer available. Essentially, the infant brain comes prewired to facilitate the development of new uses of old sensors and effectors.

Entrepreneur and author Bob Horn, who was also a participant at the workshop, argues that visual languages have already evolved and can be further evolved — perhaps, dramatically so for certain important categories of complex information, and thus progressing towards the information gulp-like sense alluded to above. In addition, researchers at IBM‘s Knowledge Management Institute and elsewhere offer stories and story languages as a highly evolved, and yet mostly untapped — except for entertainment purposes — way to rapidly convey large volumes of information. For example, when I mention the names of two television shows, The Honeymooners and The Flintstones, many TV-literate Americans in their forties and fifties will understand that these have in fact the same basic story formula, and will immediately draw on a wealth of abstractions and experience to interpret new data in terms of these stories. They may even be reminded of a Honeymooner episode when watching a Flintstone cartoon — this is powerful stuff for conveying information. The generation of television and videogame enthusiasts have a wealth of new cognitive constructs that can be leveraged in the evolution of a new sense for rapid, high volume information communication. Certainly, new notations and languages (e.g., musical notation, programming languages, and mathematics) offer many opportunities for empowering people and enhancing their performance on particular tasks. All these approaches to new uses of old senses are primarily limited by our learning abilities, both individually and collectively. Like the evolution of speech, perhaps new portions of the brain with particular capabilities could accelerate our ability to learn to use old senses in new ways. An ability to assimilate large amounts of information more rapidly could be an important next step in human evolution, potentially as important as the evolution of the first language spoken between our ancestors.

Inside the body and permanent: New genes. If the notion of “computers inside” or cyborgs raise certain ethical dilemmas, then tinkering with our own genetic code is certain to raise eyebrows as well. After all, this is shocking and frightening stuff to contemplate, especially in light of our inability to fully foresee the consequences of our actions. Nevertheless, for several reasons, including, for the sake of completeness in describing the Outside-Inside Framework, this is an area worth mentioning. While selective breeding of crops, animals, and people (as in ancient Sparta) is many hundreds of generations old, only recently have gene therapies become possible as the inner working of the billion-year-old molecular tools of bacteria for slicing and splicing DNA have been harnessed by the medical and research communities. Just as better understanding of the inner working of memory of rodents and its genetic underpinnings have allowed researchers to boost the IQs of rodents on certain maze running tasks, soon we can expect other researchers building on these results to suggest ways of increasing the IQs of humans.

University of Washington researcher and medical doctor Jeffrey Bonadio (Bonadio 2002), who was a workshop participant, discussed emerging technologies in the area of gene therapies. Gene therapy is the use of recombinant DNA as a biologic substance for therapeutic purposes, using viruses and other means to modify cellular DNA and proteins for a desired purpose.

In sum, the Outside-Inside Framework includes four main categories and a few subcategories for the ways that technology might be used to enhance human performance:

- Outside the body and environmental
  - new materials
  - new agents
  - new places
  - new mediators (tools and artifacts)
- Outside the body, personal
  - new mediators (tools and artifacts)
- Inside the body, temporary
  - new ingestibles
- Inside the body, permanent
  - new organs (new sensors and effectors)
  - new skills (new uses of old sensors and effectors)
  - new genes

The four categories progress from external to internal changes, and span a range of acceptable versus questionable changes. In the next section, we'll consider these categories from the perspective of information encoding and exchange processes in complex dynamic systems.

Information Encoding and Exchange: Towards a Unified Information Theoretic Underpinning

The Outside-Inside Framework provides one way to organize several of the key issues and ideas surrounding the use of NBICS technology advances to enhance human performance ("make us all healthier, wealthier, and wiser"). This simple framework can be shown to be largely about understanding and controlling how, where, and what information is encoded and exchanged. For example, consider the following four loosely defined systems and the way information is encoded differently, and interdependently, in each: (a) bits and the digital environment (information), (b) brains and memes and the social environment (cogno-socio), (c) bacteria and genes and the bioenvironment (nano-bio), (d) bulk atoms, raw materials, designed artifacts, and the physical environment (nano-based).

At this point, a brief digression is in order to appreciate the scale of successful information encoding and evolution in each of these loosely defined systems. People have existed in one form or another for about 2 million years, which is a few hundred thousand generations (to an order of magnitude). Today, there are about six billion people on Earth. The human body is made up of about 10^{13} cells, the human brain about 10^{10} cells (10^{27} atoms), and the human genome is about 10^{9} base pairs. Humans have been good problem solvers over the generations, creating successful civilizations and businesses as well as creating a growing body of knowledge to draw on to solve increasingly complex and urgent problems. However, in some ways, even more impressive than humans are bacteria, according to author Howard Bloom (2001). Bacteria have existed on Earth for about 3.5 billion years, which is an estimated 10^{14} bacteria generations ago. Today, there are an estimated 10^{30} bacteria (or about one hundred million bacteria for every human cell) on Earth living inside people, insects, soil, deep below the surface of the Earth, in geothermal hot springs in the depths of the ocean, and in nearly every other imaginable place. Bacteria have been successful "problem-solvers," as is evidenced by their diversity and ever-growing bag of genetic tricks for solving new problems. People have made use of bacteria for thousands of generations (though electronic digital computers only recently) in producing bread, wine, and cheese, but only in the past couple of generations have bacteria become both a tool kit and a road map for purposeful gene manipulation. Bacteria and viruses are both an ally and a threat to humans. For example, bacterial or viral plagues like the influenza outbreak of 1917 are still a major threat today. Among our best new allies in this fight are the advances in life sciences technologies enabled by more powerful digital technology. Most recently, electronic transistors have been around for less than a century, and at best, we have only a few dozen generations of manufacturing technology. Today, there are more than 10^{18} transistors on Earth, and very roughly 10 million transistors per microprocessor, 100 million PCs manufactured per year, and 10 billion embedded processors.

Returning to the issue of understanding and controlling how, where, and what information is encoded and exchanged, consider the following milestones in human history (where GA is human generations ago), as seen through the lens of the Outside-Inside Framework:

- Speech (100,000 GA): A new skill (new use of old sensors and effectors, requires learning a new audible language), encoding information in sounds, for exchanging information among people. Probably coincides with the evolution of new brain centers, new organs.
- Writing (500 GA): A new mediator and new skill (new use of old sensors and effectors, requires learning a new visual language), encoding information in visual symbols on materials from the environment for recording, storing, and

exchanging information between people. Did not require new brain centers beyond those required for spoken language.

- Libraries (400 GA): A new place and agent (organization) for collecting, storing and distributing written information.
- Universities (40 GA): A new place and agent (organization) for collecting, storing, and distributing information as social capital.
- Printing (14 GA): A new mediator (tool) for distributing information by making many physical copies of written and pictorial information.
- Accurate clocks (16 GA): A new mediator (tool) for temporal information and spatial information (accurate global navigation).
- Telephone (5 GA): A new mediator (tool) for exchanging audio information encoded electrically and transported via wires over great distances.
- Radio (4 GA): A new mediator (tool) for distributing audio information encoded electromagnetically and transported wirelessly over great distances.
- Television (3 GA): A new mediator (tool) for distributing audiovisual information encoded electromagnetically, transported wirelessly over great distances.
- Computers (2 GA): A new mediator and agent for storing, processing, creating, and manipulating information encodable in a binary language.
- Internet (1 GA): A new mediator for distributing information encodable in a binary language.
- Global Positioning System or GPS (0 GA): A new mediator for spatial and temporal (atomic clock accuracy) information.

Stepping back even further for a moment (per Bloom 2001), we can identify six fundamental systems for encoding and accumulating information: matter, genes, brains, memes, language, and bits:

- Big Bang (12 billion years ago): New place and new material - the Universe and matter
- Earth (4.5 billion years ago): New place and new materials - the Earth and its natural resources
- Bacteria (3.5 billion years ago): New species and agent, encoding information in primitive genome (DNA) in cells
- Multicellular (2.5 billion years ago): New species with multicellular chains and films
- Clams (720 million years ago): New species with multiple internal organs with primitive nervous systems
- Trilobites (500 million years ago): New species with simple brains for storing information (memes possible)
- Bees (220 million years ago): New species and agent; social insect with memes, collective IQ

- Humans and Speech (2 million years ago): New species and agent, with primitive spoken language and tools, extensive memes, collective IQ
- Writing (about 10 thousand years ago): New mediator, recordable natural language and symbolic representations
- Computers (about 50 years ago): New mediator and agent, binary language and predictable improvement curve through miniaturization

Of course, all these dates are very approximate. The important point is simply this: if the past is the best predictor of the future, then we can expect NBICS convergence to shed light on all of these key systems for encoding, exchanging, and evolving information. If (and this is a big if) we can (1) truly understand (from an information processing standpoint) the working of material interactions, genes and proteins, nervous systems and brains, memes and social systems, and natural language, and translate all this into appropriate computational models, and (2) use this deep model-based understanding to control and directly manipulate their inner workings to short-cut the normal processes of evolution, then perhaps we can create improvements (solve complex urgent problems) even faster. Of course, accelerating evolution in this way is both staggeringly difficult to do in reality as well as potentially very empowering and dangerous if we should succeed.

Again, the point here is simply that NBICS convergence has zeroed in on the key, few separate information systems that drive enhancements not only to human performance, but to the universe as we know it: matter, genes, brains, memes, language, and bits. Does this mean that we have bitten off too much? Perhaps, but it does seem to be time to ask these kinds of convergence questions, much as physicists in the late 1800s began a quest to unify the known forces. In essence, the quest for NBICS convergence is looking for the Maxwell's equations or, better yet, the "unified field theory" for complex dynamic systems that evolve, but in terms of models of information encoding, and exchange instead of models of particle and energy exchange. Author and scientist Richard Dawkins in his book *The Selfish Gene* foreshadows some of this thinking with his notion of a computational zoology to better understand why certain animal behaviors and not others make sense from a selfish gene perspective. Author and scientist Stuart Kaufman in his book *At Home in the Universe: The Search for the Laws of Self-Organization and Complexity* is searching for additional mechanisms beyond evolution's natural selection mechanism that could be at work in nature. Testing and applying these theories will ultimately require enormous computational resources.

It is interesting to note that computational power may become the limiting factor to enhancing human performance in many of the scenarios described above. What happens when Moore's Law runs out of steam? To throw one more highly speculative claim into the hopper, perhaps quantum computing will be the answer. Recently, IBM researchers and collaborators controlled a vial of a billion-billion (10^{18}) molecules designed to possess seven nuclear spins. This seven qubit quantum computer correctly factored the number 15 via Shor's algorithm and had its input programmed by radio frequency pulses and output detected by a nuclear magnetic resonance instrument. Certainly, there is no shortage of candidates for the next big thing in the world of more computing power.

Concluding Remarks: Near-Term Opportunities for e-Business Infrastructure

So what are the near-term opportunities? The R&D community is engaged. From an R&D perspective, the five innovation ecosystems (university labs, government labs, corporate labs, venture capital backed start-ups, and nonprofit/nongovernment organizations) have already geared up initiatives in all the separate NBICS (nano-bio-info-cogno-socio) areas, somewhat less in socio, and cogno is perhaps secondary to neuro. However, what about real products and services coming to market and the converged NBICS as opposed to separate threads?

From a business perspective, a number of existing technology trends generally align with and are supportive of NBICS directions. One of the major forces driving the economy these days is the transformation of businesses into e-businesses. The e-business evolution (new agent) is really about leveraging technology to enhance all of the connections that make businesses run: connections to customers, connections to suppliers, connections between employees and the different organizations inside a business, and connections to government agencies, for example. Some aspects of the NBICS convergence can not only make people healthier, wealthier, and wiser, but can make e-businesses healthier, wealthier, and wiser as well. I suspect that while many futurists are describing the big impact of NBICS convergence on augmenting human performance, they are overlooking the potentially larger and nearer term impacts of NBICS convergence on transforming businesses into more complete e-businesses. The area of overlap between what is good for business and what is good for people is in my mind one of the first big, near term areas of opportunity for NBICS convergence. Improving human performance, like improving business performance will increasingly involve new interfaces to new infrastructures.

a) *Communication infrastructure*: The shift from circuits to packets and electronics to photonics and the roll out of broadband and wireless will benefit both businesses and individuals.

b) *Knowledge infrastructure*: Knowledge management, semantic search, and natural language tools will make businesses and people act smarter.

c) *Sensor infrastructure*: Realtime access to vital information about the health of a person or business will be provided.

d) *Simulation infrastructure*: There will be a shift from *in vitro* to *in silico* biology for the design and screening of new drugs for people and new products for businesses.

e) *Intellectual property, valuations and pricing, human capital infrastructure*: Inefficiencies in these areas are a major drag on the economy overall.

f) *Miniaturization, micromanipulation, microsensing infrastructure*: Shrinking scales drive chip businesses and open new medical applications.

g) *Computing infrastructure (grid - social)*: This is still emerging, but ultimately, computer utility grids will be an enormous source of computing power for NBICS efforts.

h) *Computing infrastructure (autonomic - biological)*: The cost of managing complex technology is high; the autonomic borrows ideas from biological systems.

Already, IBM Research has begun to articulate some of the challenges and the promise of autonomic computing (http://www.research.ibm.com/autonomic/), which seeks to build a new generation of self-managing, self-regulating, and self-repairing information technology that has some of the advantages of living systems. As NBICS convergence happens, our information technology infrastructure will benefit, making many businesses more efficient and more viable.

Ultimately, NBIC convergence will lead to complete computational models of materials, genes, brains, and populations and how they evolve, forever improving and adapting to the demands of changing environments. A first step is to understand the way information is encoded and exchanged in each of these complex dynamic systems and to apply that new understanding to enhance each system. While this is an exciting undertaking, especially in light of recent advances in mapping the human genome, nanotechnology advances, and 30 some years of unabated miniaturization (Moore's Law) driving up computational capabilities, it is also a time to admit that this is still a multi-decade undertaking with lots of twists and turns in the road ahead. Better frameworks that help us inventory and organize the possibilities, as well as glimpse the ultimate goal of NBICS convergence, are still needed.

References

Bloom, H. 2001. Global brain: The evolution of mass mind from the big bang to the 21st century. John Wiley & Sons.

Bonadio, J. 2002. Gene therapy: Reinventing the wheel or useful adjunct to existing paradigms? In this volume.

SENSOR SYSTEM ENGINEERING INSIGHTS ON IMPROVING HUMAN COGNITION AND COMMUNICATION

Brian M. Pierce, Raytheon Company

The improvement of human cognition and communication can benefit from insights provided by top-down systems engineering used by Raytheon and other aerospace and defense companies to design and develop their products. Systems engineering is fundamental to the successful realization of complex systems such as multifunction radar sensors for high performance aircraft or the Army's Objective Force Warrior concept for the dismounted soldier. Systems engineering is very adept at exploring a wide trade space with many solutions that involve a multitude of technologies. Thus, when challenged by the theme of the workshop to evaluate and explore convergent technologies (nanoscience and nanotechnology, biotechnology and biomedicine, information technology, and cognitive science) for improving human cognition and communication, it was natural to start with a top-down systems engineering approach.

One of the first questions to be asked is what is meant by improvement. In sensor systems engineering, improvement covers a wide range of issues such as performance, cost, power and cooling, weight and volume, reliability, and supportability. The ranking of these issues depends on the mission for the system in question and on the sensor platform. For example, a surveillance radar system on an aircraft has much more emphasis on the power and cooling issues than does a ground-based radar system.

Improvement has many facets in the context of the Army's Objective Force Warrior system for the dismounted soldier: enhanced fightability without impeding movement or action; minimal weight; efficient, reliable, and safe power; integratability; graceful degradation; trainability; minimal and easy maintenance (ultra-reliability); minimal logistics footprint; interoperability; and affordability. The prioritization of these requirements could change depending on whether the warrior is based on space, airborne, surface ship, or undersea platforms. Ideally, the adaptability of the system is high enough to cover a wide range of missions and platforms, but issues like cost can constrain this goal.

Improvement is a relative term, and improvement objectives in the case of human cognition depend on the definition of the baseline system to be improved, e.g., healthy versus injured brain. Furthermore, does one focus solely on cognition in the waking conscious state, or is the Rapid Eye Movement (REM) sleeping conscious state also included? Although recent memory, attention, orientation, self-reflective awareness, insight, and judgment are impaired in the REM sleep state, J. Allen Hobson suggests that this state may be the most creative one, in which the chaotic, spontaneous recombination of cognitive elements produces novel configurations of new information resulting in new ideas (Hobson 1999).

Improvement objectives for human communication include enhancements in the following:

a) communication equipment *external* to the individual, e.g., smaller, lighter cell phones operable over more frequencies at lower power
b) information transfer *between* equipment and individual, i.e., through human-machine interface
c) communication and cognitive capabilities *internal* to the individual, e.g., communication outside the normal frequency bands for human vision and hearing.

If one reviews the evolution of cognitive and communication enhancement for the dismounted soldier during the last several decades, improvements in equipment *external* to the soldier and the human-machine interface predominate. For example, Raytheon is developing uncooled infrared imagers for enhanced night vision, a tactical visualization module to enable the visualization of a tactical situation by providing realtime video, imagery, maps, floor plans, and "fly-through" video on demand, and GPS and antenna systems integrated with the helmet or body armor. Other external improvements being developed by the Department of Defense include wearable computers, ballistic and laser eye protection, sensors for detection of chemical and biological warfare agents, and smaller, lighter, and more efficient power sources. Improvements that would be inside the individual have been investigated as well, including a study to enhance night vision by replacing the visual chromophores of the human eye with ones that absorb in the infrared, as well as the use of various drugs to achieve particular states of consciousness.

The convergent technologies of nanoscience and nanotechnology, biotechnology and biomedicine, information technology, and cognitive science have the potential to accelerate evolutionary improvements in cognition and communication *external* to the individual and the human-machine interface, as well as enable revolutionary improvements *internal* to the individual. The recent workshop on nanoscience for the soldier identified several potential internal improvements to enhance soldier performance and to increase soldier survivability: molecular internal computer, sensory, and mechanical enhancement; active water reclamation; short-term metabolic enhancement; and regeneration/self-healing (Army Research Laboratory 2001).

The trend in sensor systems is towards the integrated, wide band, multifunction sensor suite, in which processor/computer functions are extended into the sensing elements so that digitization occurs as early as possible in the sensing process. This type of sensor architecture enables a very high degree of adaptability and performance. However, one still has to trade the pros and cons of handling the increasing torrent of bits that results from digitizing closer to the sensor's front-end. For example, the power consumption associated with digitization can be an important consideration for a given platform and mission.

Improvements in human cognition and communication will also follow a path of higher integration and increased functionality. The exciting prospect is that the convergent technologies encompass the three major improvement paths: external, human-machine interface, and internal. This breadth should make it possible to pursue a more complete system solution to a particular need. If better night vision is desired, the convergent technologies could make it possible to trade a biological/chemical approach of modifying the photoreceptors in the eye, a micro/nano-optoelectronic imager external to the eye, or a hybrid of the two. Memory enhancement is an important element of improving human cognition, and perhaps convergent technologies could be used to build on work that reports using external electrical stimulation (Jiang, Racine, and Turnbull 1997) or infusion of nerve growth factor (Frick et al. 1997) to improve/restore memory in aged rats.

Sensor systems have benefited enormously from architectures inspired by the understanding of human cognition and communication. The possibility exists for sensor system engineering to return the favor by working in concert with the convergent technologies of nanoscience and nanotechnology, biotechnology and biomedicine, information technology, and cognitive science.

References

Army Research Laboratory/Army Research Office. 2001. Workshop on Nanoscience for the Soldier. February 8-9.

Frick, K.M., D.L. Price, V.E. Koliatsos, and A.L. Markowska. 1997. The effects of nerve growth factor on spatial recent memory in aged rats persist after discontinuation of treatment. *Journal of Neuroscience*. 17(7):2543-50 (Apr 1).

Hobson, J.A. 1999. *Consciousness*. NY: Scientific American Library. pg. 45.

Jiang, F., R. Racine, and J. Turnbull. 1997. Electrical stimulation of the septal region of aged rats improves performance in an open-field maze. *Physiology and Behavior* 62(6):1279-82 (Dec.).

CAN NANOTECHNOLOGY DRAMATICALLY AFFECT THE ARCHITECTURE OF FUTURE COMMUNICATIONS NETWORKS?

Cherry A. Murray, Lucent Technologies

We live in an era of astounding technological transformation — the Information Revolution — that is as profound as the two great technological revolutions of the past — the Agricultural and Industrial Revolutions. All around us are now familiar technologies whose very existence would have seemed extraordinary just a generation ago: cellular telephones, the optical fiber telecommunications backbone, the Internet, and the World Wide Web. All of the underlying technologies of the Information Age are experiencing exponential growth in functionality due to decreasing size and cost of physical components — similar to Moore's Law in silicon-integrated electronics technology. In the next decade, the size scale of many communications and computing devices — such as individual transistors — is predicted to decrease to the dimension of nanometers; where fundamental limits may slow down, single device functionality will increase. Before these fundamental limits are even attained, however, we must address the difficult assembly and interconnection problems with a network of millions of small devices tied together to provide the increased functionality at lower cost. If the interconnection problem is solved and if the cost of physical elements is dramatically reduced, the architectures of future communications networks — and the Internet itself — can be dramatically changed. In order for this to happen, however, we must have a breakthrough in our ability to deal with the statistical nature of devices in the simulation and design of complex networks on several levels.

Fundamental Limits to Individual Devices

Communications and computing rely ultimately on individual devices such as transistors, optical switching elements, memory elements, and detectors of electrical, optical, and radio signals. These devices are linked into physical modules like integrated circuits that perform the necessary functions or computations needed for the communications network or computer. For the last two decades, the trends of technology have dramatically decreased the size and power requirements of individual elements so that they can be integrated into a single complex package, thus reducing parts needed, space, and cost of functional modules such as communications receivers. I expect that these trends will continue, using what we have learned from nanotechnology research, until either fundamental physical limits to the individual devices are reached, or which is more likely, until we hit a new bottleneck of how to design and achieve the interconnection of these devices. When devices approach the nanometer scale, they will no longer be identical but will have a statistical distribution of characteristics: for example, in a 10 nm channel length transistor, the number of dopant atoms will be in the tens or hundreds and vary from transistor to transistor produced in an identical way, due to the random nature of the dopant diffusion process. This means that there will necessarily be a statistical variation of turn-on voltages, break-down voltages, channel conductivity, and so forth.

The Interconnection Problem

The engineering research of the next decade will most likely bring to fruition the integration of different functionalities on a single small platform, such as compact electro-photonic modules using engineered photonic bandgap structures to reduce the size of optical modulation-demodulation elements. I expect that the newest integration architectures will necessarily include fully three-dimensional circuits and up to a million or so “zero cost” single devices. These integrated modules will, in themselves, be complex networks of individual elements, each element type described by a statistical distribution of characteristics. We will need a breakthrough in the methods of integrated circuit simulation and design in order to deal with the complexity of designing these modules and to deal with the latency of signals travelling across long paths or through many connections. Right now, the simulation and design software for merely pure electronic integrated circuits, assuming that each element type is identical, is a major bottleneck in the production of application-specific integrated circuits (ASICS). One possibility in the far future is to harness the methods of directed self-assembly of the network of devices, much as our brain manages to learn from its environment how to assemble the synapses between neurons. We are not even close today.

Future Communications Network Architectures

As extremely small and low-cost communications modules are developed, certainly personal access networks — the equipment used by an individual to communicate with his or her near surroundings and to gain access to larger area local area networks and ultimately to the global wide area communications networks of the future — will become ubiquitous. These will mostly be wireless ad hoc networks, since people are mobile. Local area networks, for example, campus or in-building networks with range below one km, will be ubiquitous as well, whether wireless or wireline, depending on deployment costs. But how will the dramatic reduction of cost of the physical infrastructure for communications equipment affect the major communication long haul or wide area networks? Currently, the architectures of cross-continental or undersea or satellite communications systems are determined not only by the cost of components but by the costs associated with deployment, provisioning, reconfiguration, protection, security, and maintenance. The simulation and design tools used for complex wide area networks are in their infancy, as are the simulation and design tools for the integrated modules of which they are comprised. We need a breakthrough in simulation and design techniques. As the costs of the physical hardware components for wide area networks come down, the deployment costs will not fall as much, due to the power requirements needed in wide area systems, and this and the complexity of network management will probably determine network architectures. For example, the complexity of managing security and quality of service in a nationwide ad hoc wireless network comprised of billions of only small, low power base stations is enormous. Thus it is much more likely to have hierarchies of scale in networks, first personal, then local, and then medium range, culminating in a backbone network similar to what we have today. Again, we may be able to learn much from how biological networks configure themselves as we develop self-configuring, self-protecting, and self-monitoring networks.

SPATIAL COGNITION AND CONVERGING TECHNOLOGIES

Reginald G. Golledge, University of California at Santa Barbara

This paper explores aspects of spatial cognition and converging technologies following five themes:

1. Nano-Bio-Info-Cognitive technology (NBIC) and improving learning
2. Enhancing sensory and cognitive capabilities in the spatial domain
3. NBIC and improving human-machine interfaces
4. Suggestions about what should be done
5. Expected outcomes

NBIC and Improving Learning

What will NBIC allow us to achieve in the learning domain that we cannot achieve now?

The effects of NBIC may be

- improved knowledge of brain functioning and capabilities
- new learning domains such as immersive virtual environments
- more widespread use of nonvisual experiences for solving spatial problems
- examining sensory substitution as a way to enhance learning .

Let us briefly examine how these might occur.

Improving Knowledge of Brain Functioning and Capabilities: Place Cell Analysis.

Advances in Magnetic Resonance Imagery (MRI) have given some promise for tracking what parts of the brain are used for what functions. Opinions differ regarding the value of this technology, but much of the negative criticism is directed towards identifying which parts of the brain appear to be used for emotions such as love or hate, or for aesthetic reactions to concepts of beauty, danger, and fear. Somewhat less controversy is present in the spatial domain, where the 25-year-old hypothesis of O'Keefe and Nadel (1978) that the hippocampus is one's "cognitive map" (or place where spatial information is stored) is being actively investigated. Neurobiologists may be able to determine which neurons "fire" (or are excited) when spatial information relating to objects and their locations are sensed and stored. If NBIC can develop reliable place cell analysis, the process of mapping the human brain could be transformed into examining the geography of the brain. To do this in a thorough manner, we need to know more about spatial cognition, including understanding spatial concepts, spatial relations, spatial thinking, and spatial reasoning.

Within the domains of spatial thinking and reasoning — domains that span all scales of science and technology from the nano scale to a universe-wide scale — there is enormous potential for improving our understanding of all facets of the spatial domain. Spatial thinking and reasoning are dominated by perceptualizations, which are the multisensory expansion of visualization. The major processes of information processing include encoding of sensed experiences, the internal manipulation of sensed information in working memory, the decoding of manipulated information, and the use of the results in the decision-making and choice processes involved in problem-solving and spatial behavior. According to Golledge (2002), thinking and reasoning spatially involves

- Understanding the effects of scale
- Competently mentally transforming perceptions and representations among different geometric dimensions (e.g., mentally expanding 1-dimensional traverses or profiles to 2-D or 3-D configurations similar to that involved in geological mapping, or reducing 3-D or 4-D static or dynamic observations to 2-D formats for purposes of simplification or generalization (as when creating graphs, maps, or images)
- Comprehending different frames of reference for location, distance estimation, determining density gradients, calculating direction and orientation, and other referencing purposes (e.g., defining coordinates, vectors, rasters, grids, and topologies)
- Being capable of distinguishing spatial associations among point, line, area, and surface distributions or configurations
- Exercising the ability to perform spatial classification (e.g., regionalization)
- Discerning patterns in processes of change or spread (e.g., recognizing patterns in observations of the spatial spread of AIDS or city growth over time)
- Revealing the presence of spatial and nonspatial hierarchies

Each of the above involves sensing of phenomena and cognitive processing to unpack embedded detail. It should also be obvious that these perceptual and cognitive processes have their equivalents in information technology (IT), particularly with respect to creating, managing, and analyzing datasets. While we are creating multiple terabytes of data each day from satellites, from LIght Detection And Ranging (LIDAR), from cameras, and from visualizations, our technology for dealing with this data — particularly for dynamic updating and realtime analysis — lags somewhat, even in the most advanced systems currently invented. Even in the case of the most efficient data collector and analyzer ever developed, the human mind, there is still a need to simplify, summarize, generalize, and represent information to make it legible. The activities required to undertake this knowledge acquisition process are called education, and the knowledge accumulation resulting from this exposure is called learning. Thus, if NBIC can empower spatial thinking and reasoning, it will promote learning and knowledge accumulation among individuals and societies, and the results will have impact the entire spatial domain. (Note, there is a National Research Council committee on spatial thinking whose report is due at the end of 2002.)

To summarize, spatial thinking is an important part of the process of acquiring knowledge. In particular, spatial knowledge, defined as the product of spatial thinking and reasoning (i.e., defined as cognitive processes) can be characterized as follows:

- Spatial thinking and reasoning do not require perfect information because of the closure power of cognitive processes such as imaging, imagining, interpolating, generalizing, perceptual closure, gestalt integration, and learning

- Spatial metaphors are being used — particularly in IT related database development and operation — but it is uncertain whether they may or may not be in congruence with equivalent cognitive functioning.
- Spatial thinking has become an important component of IT. IT has focused on visualization as a dominant theme in information representation but has paid less attention to other sensory modalities for its input and output architectures; more emphasis needs to be given to sound, touch, smell, gaze, gesture, emotion, etc. (i.e., changing emphasis from visualizations to perceptualizations).

New Learning Domains

One specific way that NBIC developments may promote learning is by enhancement of virtual systems. In geography and other spatial sciences, learning about places other than one's immediate environment is achieved by accessing secondary information, as in books, maps, images, and tables. In the future, one may conceive of the possibility that all place knowledge could be learned by primary experience in immersive virtual environments. In fact, within 20 years, much geospatial knowledge could be taught in immersive virtual environments (VE) labs. This will require

- solution of the space sickness or motion sickness problems sometimes associated with immersion in VE
- quick and immediate access to huge volumes of data — as in terabytes of data on a chip — so that suitably real environments can be created
- adoption of the educational practice of "learning by doing"
- major new development of hardware and virtual reality language (VRL) software
- conviction of teachers that use of VE labs would be a natural consequence of the educational premise that humans learn to think and reason best in the spatial domain by directly experiencing environments.
- Investigation of which types of learning experiences are best facilitated by use of VE.

Using More Nonvisual Methods

Because of the absence of geography in many school curricula in the United States, many people have severely restricted access to (and understanding of) representations of the environment (for example, maps and images) and more abstract concepts (including spatial concepts of hierarchy and association or adjacency displayed by maps or data represented only in tables and graphs) that are fundamental in education and daily life. Representations of the geographic world (maps, charts, models, graphs, images, tables, and pictures) have the potential to provide a rich array of information about the modern world. Learning from spatialized representations provides insights into layout, association, adjacency, and other characteristics that are not provided by other learning modes. But, electronic spatial representations (maps and images) are not accessible to many groups who lack sight, training, or experience with computerized visualizations, thus contributing to an ever-widening digital divide. With new technological developments, such as the evolution from textual interfaces to graphically based Windows environments, and the increasing tendencies for website information to be restricted to those who can access visualizations and images, many people are being frustrated in their attempts to access necessary information — even that relevant to daily life, such as weather forecasts.

When viewing representations of the geographic world, such as a map on a computer screen, sight provides a gestalt-like view of information, allowing the perception of the synoptic whole and almost simultaneously recognizing and integrating its constituent parts. However, interacting with a natural environment is in fact a multi-modal experience. Humans engage nearly all of their sensory modalities when traversing space. Jacobson, Rice, Golledge and Hegarty (2002) summarize recent literature relating to non-visual interfaces. They suggest that, in order to attend to some of this multisensory experience and to provide access to information for individuals with restricted senses, several research threads can be identified for exploring the presentation of information multimodally. For example, information in science and mathematics (such as formulae, equations, and graphs) has been presented through auditory display (e.g., hearing a sine wave) and through audio-guided keyboard input (Gardner et al. 1998; Stevens et al. 1997). Mynatt (1977) has developed a tonal interface that allows users without vision to access Windows-style graphical user interfaces. Multimodal interfaces are usually developed for specialist situations where external vision is not necessarily available, such as for piloting and operating military aircraft (Cohen and Wenzel 1995; Cohen and Oviatt 1995; Rhyne and Wolf 1993).

Jacobson et al. also point out that abstract sound variables have been used successfully for the presentation of complex multivariate data. Parkes and Dear (1990) incorporated "sound painting" into their tactual-auditory information system (NOMAD) to identify gradients in slope, temperature, and rainfall. Yeung (1980) showed that seven chemistry variables could be presented through abstract sound and reported a 90% correct classification rate prior to training and a 98% correct response rate after training. Lunney and Morrison (1981) have shown that sound graphs can convey scientific data to visually impaired students. Sound graphs have also been compared to equivalent tactual graphs; for example, Mansur et al. (1985) found comparable information communication capabilities between the two media, with the auditory displays having the added benefit of being easier to create and quicker to read. Recent research has represented graphs by combining sound and brailled images with the mathematical formula for each graph being verbally presented while a user reads the brailled shape. Researchers have investigated navigating the Internet World Wide Web through audio (Albers 1996; Metois and Back 1996) and as a tool to access the structure of a document (Portigal and Carey 1994). Data sonification has been used to investigate the structure of multivariate and geometric data (Axen and Choi 1994; Axen and Choi 1996; Flowers et al. 1996), and auditory interfaces have been used in aircraft cockpits and to aid satellite ground control stations (Albers 1994; Ballas and Kieras 1996; Begault and Wenzel 1996). But while hardware and software developments have shown "proof of concept," *there appear to be few successful implementations of the results for* *general use (except for some gaming contexts) and no conclusive behavioral experiments to evaluate the ability of the general public or specialty groups (e.g., the vision-impaired) to use these innovations to interpret on screen maps, graphics, and images*.

Thus, while Jacobson et al. (2002) have illustrated that multimodal interfaces have been explored within computer science and related disciplines (e.g., Delclos and Hartman 1993; Haga and Nishino 1995; Ladewski 1996; Mayer and Anderson 1992; Merlet, Nadalin, Soutou, Lapujade, and Ravat 1993; Morozov 1996; Phillips 1994; Stemler 1997; Hui et al. 1995; and others), and a number of researchers have looked at innovative interface mediums such as gesture, speech, sketching, and eye tracking (e.g., Ballas and Kieras 1996; Briffault and Denis 1996; Dufresne et al. 1995; Schomaker et al. 1995; Taylor et al. 1991), they also claim that only recently are such findings beginning to have an impact upon technology for general education, a view shared by Hardwick et al. (1996; 1997).

In summary, extrapolating from this example, one can assume that developments in NBIC will impact the learning activities of many disciplines by providing new environments for experience, by providing dynamic realtime data to explore with innovative teaching methods, and (if biotechnology continues to unpack the secrets of the brain and how it stores information as in place cell theory), the possibility of direct human-computer interaction for learning purposes may all be possible. Such developments could

- enhance the process of spatial learning by earlier development of the ability to reason abstractly or to more readily comprehend metric and nonmetric relations in simple and complex environments
- assist learning by discovering the biotechnological signatures of phenomena and discovering the place cells where different kinds of information are stored, and in this way enhance the encoding and storage of sensed information
- where functional loss in the brain occurs (e.g., if loss of sight leaves parts of the brain relatively inactive), to find ways to use the cells allocated to sight to be reallocated to other sensory organs, thus improving their functioning capabilities.
- Representations of the geographic world (maps, charts, models, graphs, images, tables, and pictures) have the potential to provide a rich array of information about the modern world.
- Learning from spatialized representations provides insights into layout, association, adjacency, and other spatial characteristics that are not provided by other learning modes.
- However, interacting with a natural environment is in fact a multimodal experience. Humans engage nearly all of their sensory modalities when traversing or experiencing space.

Given the dominance of computer platforms for representing information and the overwhelming use of flat screens to display such information, there is reason to believe that multimodal representations may not be possible until alternatives to 2-D screen surfaces have been developed for everyday use. The reasons for moving beyond visualization on flat screens are compelling and are elaborated on later in this chapter.

Enhancing Sensory and Cognitive Capabilities in the Spatial Domain

How can we exploit developments in NBIC to enhance perceptual and cognitive capabilities across the life span, and what will be the types of developments needed to achieve this goal?

To enhance sensory and cognitive capabilities, a functional change in the way we encode information, store it, decode it, represent it, and use it may be needed. Much of the effort in Information Technology has been directed towards developing bigger and bigger databases that can be used on smaller and smaller computers. From satellites above we get terabytes of data (digitized records of the occurrence of phenomena), and we have perhaps outgrown our ability to examine this data. As nanotechnology and IT come into congruence, the terabytes of data being stored in boxes will be stored on chips and made accessible in real time via wearable and mobile computers, and even may be fed into smart fabrics woven into the clothes we wear. But just how well can we absorb, access, or use this data? How much do we need to access? And how best *can* we access it and use it? The question arises as to how we can exploit human perception and cognition to best help in this process, and the answer is to find out more about these processes so that they can be enhanced. Examples of questions to be pursued include the following:

- How can we enhance the sensory and cognitive aspects of human wayfinding for use in navigating in cyberspace?
- What particular sensory and cognitive capabilities are used in the field, and how do we enhance them for more effective fieldwork with wearable and mobile computers (e.g., for disaster responses)?
- How do we solve problems of filtering information for purposes of representation and analysis (e.g., enhance visualizations)?
- How do we solve the problem of resolution, particularly on the tiny screens typical of wearable and field computers?
- What alternatives to visualization may be needed to promote ease of access, representation, and use of information?
- What is the best mode for data retrieval in field settings (e.g., how do we get the information we need now)?
- How can we build technology to handle realtime dynamic input from several sources, as is done by human sensory organs and the human brain?
- Will we need a totally new approach to computer design and interface architecture (e.g., abandon keyboards and mice) that will allow use of the full range of sensory and cognitive capabilities, such as audition, touch, gaze, and gesture (e.g., the use of Talking Signs® and Internet connections to access websites tied to specific locations)?

Visualization is the dominant form of human-IT interaction. This is partly because the visual sense is so dominant, particularly in the spatial domain. It is also the dominant mode for representation of analyzed data (on-screen). But visualization is but a subset of spatialization, which goes beyond the visual domain by using everyday multimodal situations (from desktops and file cabinets to overlay and digital worlds) to organize and facilitate access to stored information. These establish a linking by analogy and metaphor between an information domain and familiar elements of everyday experience. Spatial (and specifically geographic) metaphors have been used as database organizing systems. But even everyday geospatial experiences are biased, and to enhance our sensory and cognitive abilities we need to recognize those biases and mediate them if successful initiation of everyday knowledge and experience (including natural languages) are to be used to increase human-IT interactions.

The main problem arising from these usages is simply that an assumption of general geospatial awareness is false. Basic geographic knowledge (at least in the United States) is minimal, and knowledge of even rudimentary spatial concepts like distance, orientation, adjacency, and hierarchy is flawed. Recent research in spatial cognition has revealed a series of biases that permeate naïve spatial thinking. Partly because of a result of cognitive filtering of sensed information and partly because of inevitable technical errors in data capture and representation, biases occur. Golledge (2002) has suggested that these include the following:

- conceptual bias due to improper thinking and reasoning (e.g., applying metric principles to nonmetric situations)
- perceptual biases, including misunderstandings and misconceptions of notions of symmetry, alignment, clustering, classification, closure, and so on (e.g., assuming Miami, Florida, MUST be east of Santiago, Chile, because Miami is on the east cost of North America and Santiago is on the west coast of South America) (Fig. B.1)

Figure B.1. Cognitive East/ West alignment effects.

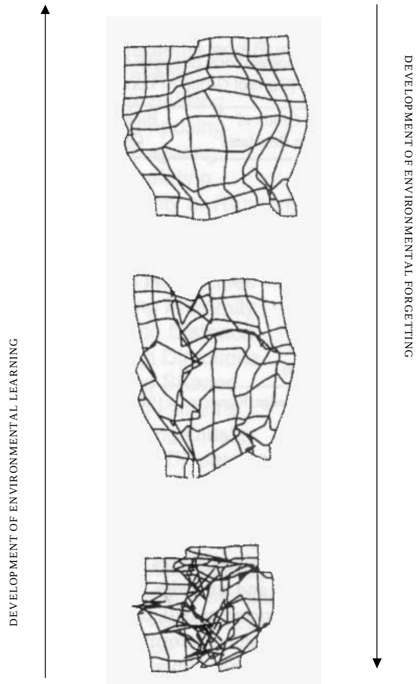

Figure B.2. Three examples of cognitive maps, of long-term residents (top), mid-term residents, (middle), and newcomers (bottom), recovered using non-metric multidimensional scaling of cognitive interpoint distances. (The exact parallel reversals for memory loss is speculative.)

- violating topological features of inclusion and exclusion when grouping (spatial) data
- assuming distance asymmetry when distance symmetry actually exists, and vice versa (e.g., different perceptions of trips to and from work)
- inappropriate use of cognitive concepts of rotation and alignment (e.g., misreading map orientation)
- cognitively overestimating shorter distances and under-estimating longer distances (Stevens' Law or regression towards the mean)
- distortions in externalized spatial products (e.g., distorted cognitive maps) (Liben 1982; Fig. B.2)
- bias that results from using imprecise natural language (e.g., fuzzy spatial prepositions like "near" and "behind" that are perspective dependent). (Landau and Jackendoff 1993)

Golledge has argued that these (and other storage, processing, and externalizing biases) result in perceptual and cognitive errors in encoding, internally manipulating, decoding, and using stored information. The following are examples of the accommodations humans make to deal with these biases (incidentally developing new ones):

- making naturally occurring irregular shapes and areas regular for purposes of simplification, representation, and generalization
- mentally rotating features or distributions to fit preconceptions (e.g., vertically aligning North and South America, as shown in Figure B.1)
- imposing hierarchical orderings to clarify distributions (e.g., systems of landmarks)
- making effective rational decisions without perfect information
- cognitively generalizing from one scale to another without appropriate empirical evidence (e.g., from laboratory to real world scales)
- realizing that data collected for machine use has to be more perfect than data collected for human use.

NBIC and Improving Human-Computer Interfaces and Interactions

A key question is why won't existing interface architecture be appropriate for human-computer interaction in the future?

Existing interface architecture is still being modeled on dated technology — the typewriter keyboard and the cursor driven mouse — and not for ease of human-computer interaction. The interface concern is the most pressing problem of HCI and is its most critical part. It is the medium through which information is accessed, questions are posed, and solution paths are laid out and monitored. It is the tool with which the user manipulates and interacts with data. Interface architectures like the desktop, filing cabinet, and digital world are implemented (still) via keyboards and mice. Today's interfaces are cursor dependent and contribute significantly to creating a digital divide that impedes 8 million sight-impaired and 82 million low-vision (potential) users from freely interacting with the dominant IT of this age.

Communicating involves transferring information; to do so requires compatibility between sender and receiver. The interface architecture that controls human-computer information exchange, according to Norman (1988), must accomplish the following:

- facilitate the exchange of knowledge in the environment and knowledge in the head
- keep the interaction task simple
- ensure that operations are easy to do
- ensure correct transfer among information domains
- understand real and artificial restraints on interaction
- acknowledge existence of error and bias due to modal difficulties
- eventually standardize procedures

Thus, the interface must maximize the needs of both human user and computer.

These needs raise the question of what cutting edge hardware (e.g., rendering engines, motion tracking by head mounted display units, gaze tracking, holographic images, avatars complete with gestures, and auditory, tactual, and kinesthetic interface devices), adds to information processing? Besides the emphasis on historic input devices (keyboard and mouse), there is a similar emphasis on a dated output device, the limited domain of the flat computer screen (inherited from the TV screen of the 1930s), which is suited primarily for visualization procedures for output representation. While there is little doubt that the visual senses are the most versatile mode for the display of geospatial data and data analysis (e.g., in graph, table, map, and image mode), it is also argued that multiple modality interfaces could enrich the type, scale, and immediacy of displayed information. One of the most critical interface problems relates to the size and resolution of data displays. This will be of increasing importance as micro-scale mobile and wearable computers have to find alternatives to 2-inch square LED displays for output presentation. The reasons for moving beyond visualization on flat screens are compelling. Examples include the following:

- multimodal access to data and representations provide a cognitively and perceptually rich form of interaction

- multimodal input and output interfaces allow HC interaction when sight is not available (e.g., for blind or sight-impaired users) or when sight is an inappropriate medium (e.g., accessing onscreen computer information when driving a vehicle at high speeds)
- when absence of light or low precludes the use of sight
- when visual information needs to be augmented
- when a sense other than vision may be necessary (e.g., for recording and identifying bird calls in the field)

Nonvisual technology allows people with little or no sight to interact (e.g., using sound, touch, and force-feedback) with computers. Not only is there a need for text to speech conversion, but there is also a need to investigate the potential use of nonvisual modalities for accessing cursor-driven information displays, icons, graphs, tables, maps, images, photos, windows, menus, or other common data representations. Without such access, sight-disabled and low-sight populations are at an immense disadvantage, particularly when trying to access spatial data. This need is paramount today as home pages on the World Wide Web encapsulate so much important information in graphic format, and as digital libraries (including the Alexandria Digital Map and Image Library at the University of California, Santa Barbara) become the major storage places for multidimensional representations of spatial information.

In the near future, one can imagine a variety of new interfaces, some of which exist in part now but which need significant experimentation to evaluate human usability in different circumstances before being widely adopted. Examples of underutilized and underinvestigated technologies include the following:

- a force-feedback mouse that requires building virtual walls around on-screen features, including windows, icons, objects, maps, diagrams, charts, and graphs. The pressure-sensitive mouse allows users to trace the shape of objects or features and uses the concept of a gravity well to slip inside a virtual wall (e.g., a building entrance) to explore the information contained therein (Jacobson et al. 2002).
- vibrotactile devices (mice) that allow sensing of different surfaces (dots, lines, grates, and hachures) to explore flat, on-screen features (e.g., density shading maps and meteorological or isoline temperature maps) (O'Modhrain and Gillespie 1995; Jacobson, et al. 2002)
- use of real, digitized, or virtual sounds including speech to identify on-screen phenomena (e.g., Loomis, Golledge, and Klatzky 2001)
- avatars to express emotions or give directions by gesturing or gazing
- smart clothing that can process nearby spatial information and provide information on nearby objects or give details of ambient temperature, humidity, pollution levels, UV levels, etc.

Currently, the use of abstract sound appears to have significant potential, although problems of spatial localization of sound appear to offer a significant barrier to further immediate use. Some uses (e.g., combinations of sound and touch — NOMAD — and sound and Braille lettering — GPS Talk — are examples of useful multimodal interfaces (e.g., Parkes and Dear 1990; Brabyn and Brabyn 1983; Sendero Group 2002). Some maps (e.g., isotherms/density shading) have proven amenable to sound painting, and researchers in several countries have been trying to equate sound and color. At present, much of the experimentation with multimodal interfaces is concentrated in the areas of video games and cartoon-like movies. Researchers such as Krygier (1994) and Golledge, Loomis, and Klatzky (1994) have argued that auditory maps may be more useful than tactual maps and may, in circumstances such as navigating in vision-obstructed environments, even prove more useful than visual maps because they don‘t require map-reading ability but rely on normal sensory experiences to indicate spatial information such as direction.

What Needs to be Done to Help NBIC Make Contributions in the Spatial Domain?

- If space is to be used as a metaphor for database construction and management, and if human wayfinding/navigation practices are to be used as models for Internet search engines, there are a host of spatial cognition research activities that need to be pursued. First there is a need for a concept-based common vocabulary. There must be a sound ontology, an understanding of spatial primitives and their derivatives, and a meaningful way to communicate with a computer using natural language and its fuzzy spatial prepositions (i.e., a common base of spatial linguistics, including a grammar).
- We need to find matches between information types and the best sensory modalities for representing and using each type of information.
- We need an educated and IT-enlightened science and engineering community that understands spatial thinking and reasoning processes.
- We need to change educational and learning practices to produce an NBIC-enlightened public and an IT-enlightened set of decision makers. Part of this need can be achieved by producing spatially aware professionals who understand and use actual or enhanced sensory and cognitive capabilities to understand and react to different situations and settings.
- We need to explore the cognitive processes used in risky decision making and use innovative IT practices to develop databases, management systems, and analytical techniques that are cognitively compatible with these processes (Montello 2001).
- We need to develop new realtime dynamic human-computer interfaces (both input and output) that facilitate collaborative decision making. This may involve building virtual environments suited for realtime collaborative image exchange and simultaneous use, analysis, modification, and representation of data, even when researchers are continents apart.
- We need to determine what dimensions of cyberspace are compatible with perceptualization and visualization, particularly in the spatial domain.
- We need to define the impacts of selecting specific scales and levels of resolution for visual or perceptual representation of information.

- We need to explore the value of changing network representations and displays of information in cyberspace to grid layout or configurational displays — the expansion from 1- to 2- or 3-dimensional information representations would facilitate a higher level of abstract thinking and reasoning to be implemented in analyzing configurational displays.
- The explosion of interfaces built upon visualization has produced too many graphic interfaces that do not maximize cognitive capabilities of users and have further disadvantaged disabled groups such as the blind or sight-impaired. This latter fact is continuing the computer alienation of aged populations, where over 70% have low vision or other sight problems. There are, according to census estimates, over 52 million disabled people in the United States. Approximately 3-4 million of these are blind, legally blind, or severely vision-impaired. A further 80+ million people have low vision. We cannot ignore these groups or exclude them from use of future technology.
- We need to determine optimal output interfaces for wearable computers that do not limit the user to visually reading complex displays (e.g., maps) on tiny screens. This carries with it the various cartographic representation problems of choosing scale, resolution, degree of simplification, generalization, and accuracy. This is not just a computer graphics problem, but a problem for cartographic theorists, empirical researchers, and researchers in spatial perception and spatial cognition, and it may involve innovative nanotechnology to build “fold-out“ or “expandable“ screens.
- There is a need to explore interfaces that can meaningfully display dynamic data at various scales and degrees of resolution.
- There is a need to examine whether nano- or biotechnology can alter the senses and cognitive capabilities of humans to enhance HCI. In particular, can nano-biotechnology enhance our tactual and auditory capabilities (e.g., sensing gloves and ear implants) to ensure that information processing becomes perceptually and cognitively less biased and error ridden?
- There is a need for distributed national learning and research networks to be developed to encourage timely transfer of information from the research to the educational domains; otherwise, the current 3-5 year lags needed for much of this transfer to take place will continue.
- As we learn more about how the mind stores data, there is a need to examine whether we can use the mind as a model to enhance efforts to build a national network of digital libraries.
- There is a need for solving problems associated with using immersive virtual environments (e.g., motion sickness) so that their real potential in research and decision making can be exploited and evaluated.
- There is a need to explore ways to increase the effectiveness of human-environment relations. This may involve
  - developing personal guidance and spatial information systems that allow people to carry with them in a wearable computer all the local

environmental information that they need to undertake daily activities (Fig. B.3)

– developing smart environments that allow people to access wireless information (e.g., infrared-based auditory signage or locally distributed servers that allow immediate access to the Internet and web pages) (Fig. B.4).

Figure B.3. Personal guidance system.

- Since environmental information is filtered through our senses and consequently is biased, individually selective, and related to stage of cognitive development, we need to know to what extent human sensing is dependent on perspective or point of view for encoding spatial relations. Attention must be paid to the roles of alignment, frames of reference, and scale or resolution (e.g., asymmetries of distance, orientation error, or locational inaccuracy), which produce information not always consistent with metric geometries and logically based algebras used to unpack information from data about the real world. Perhaps a new subjective mathematics is needed to interpret our cognitive maps.

Figure B.4. "Smart environments."

- We need to determine if knowledge of wayfinding in the real world can help us find our way in cyberspace. Spatial knowledge in humans develops from landmark route configurational understanding. Much high-order spatial knowledge in humans concerns understanding spatial relations embedded in configurational or layout knowledge, whereas much of the knowledge in IT is link- and network-based, potentially reducing its information potential by requiring human ability to integrate information obtained from specific routes in cyberspace.

There are two dominant ways for NBIC to impact the 52+ million disabled people in the United States:

1. free them from the tyranny of print and other “inaccessible” visual representations
2. help them obtain independence of travel

Enacting measures like the following will increase mobility, employability, and quality of life:

- changing computer interface architecture so that disabled groups (e.g., blind, sight impaired, dyslexic, arthritic, immobile) can access the Internet and its webpages as transparently and quickly as able-bodied people
- enabling wearable computers for use in everyday living (e.g., finding when the next bus is due or where it is now) (Fig. B.4)
- developing voice-activated personal guidance systems using GPS, GIS, and multimodal interfaces that will enable people to travel in unfamiliar environments (Fig. B.4)
- improve speech recognition for input to computers
- use infrared-based remote auditory signage systems (RASS) (e.g., talking sign technology) to facilitate wayfinding, business or object location identification, recognition of mass transit services and promotion of intermodal transfer, and to define other location-based services and information systems

Outcomes

Following are some outcomes of the integration of spatial cognition and converging NBI technologies:

- Expanding sensory and cognitive capabilities should improve learning and result in a more NBIC-enlightened public, scientists, engineers, and public policymakers.
- Developing multimodal input and output interfaces will enrich human ability to process and analyze information, covering all types of spatial information required for microscopic, global, or extraterrestrial research. It will also help to remove the rapidly growing effects of the digital divide by allowing more disabled (or otherwise disadvantaged) people to join the computer-literate population, thus improving employment possibilities and improving quality of life.

Converging NBIC technology will broaden our abilities to think “outside the box” in a variety of sensory domains, such as the following examples of convergence of NBI and spatial cognition methods:

- Natural language-driven mobile and wearable computers
- Internet search engines based on human wayfinding practices
- Smart fabrics that sense the environment and warn us of pollution levels, etc.
- Smart environments (e.g., remote auditory signage systems) that talk to us as we travel through them

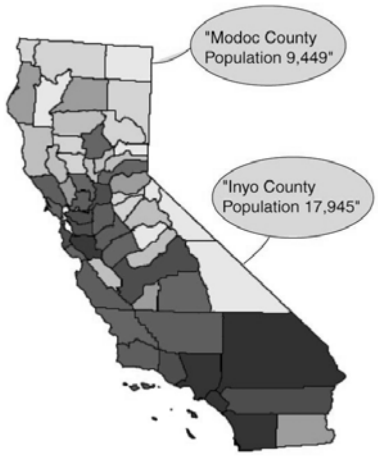

Figure B.5. Talking maps.

- GPS-based personal guidance systems that facilitate travel (e.g., tourism) in unfamiliar places
- Smart maps that explain themselves at the touch of a stylus or as a result of gaze or gesture (e.g., "You are here" maps or on-screen computer representations of data) (Fig. B.5)
- Robotic guide dogs that carry large environmental databases and can develop routes to unfamiliar places
- Smart buildings that inform about their contents and inhabitants, e.g., transit terminals (Fig. B.6).

Of particular interest are NBIC-based knowledge and devices that enhance spatial cognition used in wayfinding performance:

- Remote auditory signage (Talking Signs/Remote Infrared Auditory Signage) (at places or on vehicles, including mass transit)
- Talking fluorescent lights inside buildings such as shopping centers and transit terminals (Fig. B.7)

GPS-based guidance systems with Pointlink capabilities to locations and websites for place-based information.

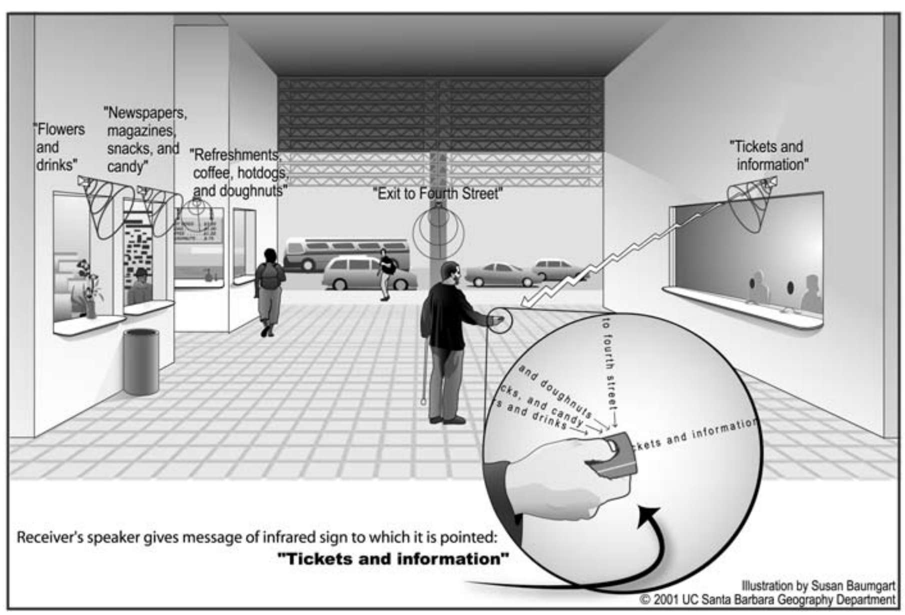

Figure B.6. Transit terminal with remote auditory signage.

Figure B.7. Talking neon lights in airport terminal.

Conclusion

The convergence of nano-, bio-, info- technology and spatial cognition research will

- broaden our ability to think outside the box
- ensure that NBI technologies are compatible with ways humans think and reason
- facilitate new product development
- help remove barriers to the natural integration of disabled and disadvantaged groups into the community, thus improving their quality of life
- provide new environments for learning
- enhance cognitive functioning by improving perceptual and cognitive capabilities
- help create less abstract and more "naturally human" computer interface architecture
- once we have learned how and where spatial information is stored in the brain (place cell analysis), this may prompt new ideas about how we think and reason

For example, eventually, the most powerful computer interface will rely on an architecture that combines geospatial metaphors with spatialization principles and multimodal input and output devices that provide access to text, maps, images, tables, and gestures.

But there is the inevitable downside, such as the thorny ethical and legal issues of defining and maintaining appropriate levels of individual privacy and security of public or business information. But developments in NBIC are the future of humankind, and these and other unrealized problems, must — in the way of humankind — be faced and solved.

Finally, if VE can be developed in an effective way, humans will have many of the capabilities of the Star Trek holodeck. They will stroll through the Amazon jungles, trek to the North or South Pole, explore an active volcano, avalanche, or hurricane, redesign cities or parts of them, change transport systems to maximize the benefits of intelligent highways, visit drought areas, explore areas of poverty or crime, all within the safety of VE. The contribution of such systems to education, research, and decision making in the policy arena could be immense. As long as we can solve the cognition and technical problems of building and using VE, these goals may be achievable.

References

Albers, M.C. 1994. The Varese System, hybrid auditory interfaces and satellite-ground control: Using auditory icons and sonification in a complex, supervisory control system. In G. Kramer and S. Smith (Eds.), *Proceedings of the Second International Conference on Auditory Display* (pp. 3-13). Sante Fe, NM.

Albers, M.C. 1996. Auditory cues for browsing, surfing, and navigating the WWW: The audible web. In S.P. Frysinger and G. Kramer (Eds.), *Proceedings of the Third International Conference on Auditory Display* (pp. 85-90). Palo Alto, Ca.

Axen, U., and I. Choi. 1994. Using additive sound synthesis to analyze simple complexes. In G. Kramer and S. Smith (Eds.), *Proceedings of the Second International Conference on Auditory Display* (pp. 31-44). Sante Fe, NM.

Axen, U., and I. Choi. 1996. Investigating geometric data with sound. In S.P. Frysinger and G. Kramer (Eds.), *Proceedings of the Third International Conference on Auditory Display* (pp. 25-28). Palo Alto, Ca.

Ballas, J., and D.E. Kieras. 1996. Computational modeling of multimodal I/O in simulated cockpits. In S.P. Frysinger and G. Kramer (Eds.), *Proceedings of the Third International Conference on Auditory Display* (pp. 135-136). Palo Alto, Ca.

Begault, D.R., and E.M. Wenzel. 1996. A virtual audio guidance and alert system for commercial aircraft operations. In S.P. Frysinger and G. Kramer (Eds.), *Proceedings of the Third International Conference on Auditory Display* (pp. 117-122). Palo Alto, Ca.

Brabyn, L.A., and J.A. Brabyn. 1983. An evaluation of "Talking Signs" for the blind. *Human Factors, 25*(1), 49-53.

Briffault, X., and M. Denis. 1996 (August). *Multimodal interactions between drivers and co-drivers: An analysis of on-board navigational dialogues.* Paper presented at the Proceedings of the Twelfth European Conference on Artificial Intelligence, Budapest.

Cohen, M., and E.M. Wenzel. 1995. The design of multidimensional sound interfaces. In T. A. F. I. I. I. E. Woodrow Barfield (Ed.), *Virtual environments and advanced interface design.* (pp. 291-346): Oxford University Press, New York, NY, US.

Cohen, P.R., and S.L. Oviatt. 1995. The role of voice input for human-machine communication. *Proceedings of the National Academy of Sciences, 92*(22), 9921-9927.

Delclos, V.R., and A. Hartman. 1993. The impact of an interactive multimedia system on the quality of learning in educational psychology: an exploratory study. *Journal of Research on Technology in Education, 26*(1), 83-93.

Dufresne, A., O. Martial, and C. Ramstein. 1995. *Multimodal user interface system for blind and "visually occupied" users: Ergonomic evaluation of the haptic and auditive dimensions*.

Flowers, J.H., D.C. Buhman, and K.D. Turnage. 1996. Data sonification from the desktop: Should sound be a part of standard data analysis software. In S.P. Frysinger and G. Kramer (Eds.), *Proceedings of the Third International Conference on Auditory Display* (pp. 1-8). Palo Alto, Ca.

Gardner, J.A., R. Lundquist, and S. Sahyun. 1998 (March). *TRIANGLE: A Tri-Modal Access Program for Reading, Writing and Doing Math.* Paper presented at the Proceedings of the CSUN International Conference on Technology and Persons with Disabilities, Los Angles.

Golledge, R.G. 2002. The nature of geographic knowledge. *Annals of the Association of American Geographers, 92*(1), In Press.

Golledge, R.G., J.M. Loomis, and R.L. Klatzky. 1994 (February 20-26). *Auditory maps as alternatives to tactual maps.* Paper presented at the 4th International Symposium on Maps and Graphics for the Visually Impaired, Sao Paulo, Brazil.

Haga, H., and M. Nishino. 1995 (17-21 June). *Guidelines for designing hypermedia teaching materials.* Paper presented at the Educational Multimedia and Hypermedia 1995. Proceedings of ED-MEDIA 95 - World Conference on Educational Multimedia and Hypermedia, Graz, Austria.

Hardwick, A., S. Furner, and J. Rush. 1996 (21 Jan). *Tactile access for blind people to virtual reality on the World Wide Web.* Paper presented at the IEE Colloquium on Developments in Tactile Displays (Digest No.1997/012) IEE Colloquium on Developments in Tactile Displays, London, UK.

Hardwick, A., J. Rush, S. Furner, and J. Seton. 1997. Feeling it as well as seeing it-haptic display within gestural HCI for multimedia telematics services. In P.A. Harling and A.D. N. Edwards (Eds.), *Progress in Gestural Interaction. Proceedings of Gesture Workshop '96, York, UK* (pp. 105-16). Berlin, Germany: Springer-Verlag.

Hui, R., A. Ouellet, A. Wang, P. Kry, S. Williams, G. Vukovich, and W. Perussini. 1995. Mechanisms for haptic feedback, *IEEE International Conference on Robotics and Automation* (pp. 2138-2143).

Jacobson, R.D., M. Rice, R.G. Golledge, and M. Hegarty. 2002. *Force feedback and auditory interfaces for interpreting on-screen graphics and maps by blind users* (Technical Paper funded by the UCSB Research Across the Disciplines (RAD) Program): Geography Departments of Florida State University and University of California Santa Barbara.

Krygier, J.B. 1994. Sound and Geographic Visualisation. In A.M. MacEachren and D.R. Fraser-Taylor (Eds.), *Visualisation in Modern Cartography* (pp. 149-166): Pergamon.

Ladewski, B.G. 1996. Interactive multimedia learning environments for teacher education: comparing and contrasting four systems. *Journal of Computers in Mathematics and Science Teaching 15*(1-2) 173-97.

Landau, B., and R. Jackendoff. 1993. "What" and "where" in spatial language and spatial cognition. *Behavioral and Brain Sciences 16*, 217-238.

Loomis, J.M., R.G. Golledge, and R.L. Klatzky. 2001. GPS-based navigation systems for the visually impaired. In W. Barfield and T. Caudell (Eds.), *Fundamentals of Wearable Computers and Augmented Reality* (pp. 429-446). Mahway, NJ: Erlbaum.

Lunney, D., and R. Morrison. 1981. High Technology Laboratory Aids for Visually Handicapped Chemistry Students. *Journal of Chemical Education, 8*(3), 228-231.

Mansur, D., M. Blattner, and K. Joy. 1985. Sound graphs: A numerical data analysis for the blind. *Journal of Medical Systems, 9*(3) 163-174.

Marston, J. 2002. *Towards an Accessible City: Empirical Measurement and Modeling of Access to Urban Opportunities for those with Vision Impairments Using Remote Infrared Audible Signage.* , UCSB, Santa Barbara. Unpublished Ph.D. Dissertation.

Mayer, R.E., and R.B. Anderson. 1992. The instructive animation: Helping students build connections between words and pictures in multimedia learning. *Journal of Educational Psychology, 84*, 444-452.

Merlet, J.F., C. Nadalin, C. Soutou, A. Lapujade, and F. Ravat. 1993 (17-20 Oct.). *Toward a design method of a multimedia information system.* Paper presented at the 1993 International Conference on Systems, Man and Cybernetics. Systems Engineering in the Service of Humans (Cat. No.93CH3242-5) Proceedings of IEEE Systems Man and Cybernetics Conference - SMC, Le Touquet, France.

Metois, E., and M. Back. 1996. BROWeb: An interactive collaborative auditory environment on the world wide web. In S.P. Frysinger and G. Kramer (Eds.), *Proceedings of the Third International Conference on Auditory Display* (pp. 105-110). Palo Alto, Ca.

Montello, D.R., ed. 2001. *Spatial Information Theory: Foundations of Geographic Information Science. Proceedings, International Conference, COSIT 2001, Morro Bay, CA, September*. New York: Springer.

Morozov, M. 1996 (14-16 Sept.). *Multimedia lecture room: a new tool for education.* Paper presented at the Multimedia, Hypermedia, and Virtual Reality. Models, Systems, and Applications. First International Conference, Moscow, Russia.

Mynatt, E.D. 1997. Transforming graphical interfaces into auditory interfaces for blind users. *Human-Computer Interaction 12*(1-2), 7-45.

Norman, D.A. (1988). *The Psychology of Everyday Things*. New York: Basic Books.

O'Keefe, J., and L. Nadel. 1978. *The Hippocampus as a Cognitive Map*. Oxford: Clarendon Press.

O'Modharain, M., and B. Gillespie. 1995. *The Moose: A Haptic User Interface for Blind Persons* (Stan-M95; CCRMA). Stanford, CA: Stanford.

Parkes, D., and R. Dear. 1990. *NOMAD: AN interacting audio-tactile graphics interpreter.* Paper presented at the Reference Manual, Version 2.0, NSW Australia: Institute of Behavior Science, University of Newcastle.

Phillips, R. 1994. Producing interactive multimedia computer-based learning projects. *Computer Graphics, 28*(1), 20-4.

Portigal, S., and T. Carey. 1994. Auralization of document structure. In G. Kramer and S. Smith (Eds.), *Proceedings of the Second International Conference on Auditory Display* (pp. 45-54). Sante Fe, NM.

Rhyne, J.A., and C. Wolf. 1993. Recognition based user interfaces. In H.R. Harston and D. Hix (Eds.), *Advances in Human-Computer Interaction: Vol. 4* (pp. 191-250). Norwood, NJ: Ablex.

Schomaker, L., J. Nijtmans, A. Camurri, F. Lavagetto, P. Morasso, C. Benoit, T. Guiard-Marginy, B. Le Goff, J. Robert-Ribes, A. Adjoudani, I. Defee, S. Munch, K. Hartnung, and J. Blauert. 1995. *A taxonomy of mulitmodal interaction in the human information processing system* : Esprit Project 8579.

Sendero Group. 2002.. Available: www.senderogroup.com.

Stemler, L.K. 1997. Educational characteristics of multimedia: a literature review. *Journal of Educational Multimedia and Hypermedia, 6*(3-4), 339-59.

Stevens, R.D., A.D.N. Edwards, and P.A. Harling. 1997. Access to mathematics for visually disabled students through multimodal interaction. *Human-Computer Interaction 12*, 47-92.

Taylor, M.M., F. Neel, and D.G. Bouwhuis, eds. 1991. *The structure of multimodal dialogue.* Amsterdam: North-Holland.

Yeung, E. 1980. Pattern recognition by audio representation of multivariate analytical data. *Analytical Chemistry, 52*(7) 1120-1123.

VISUAL LANGUAGE AND CONVERGING TECHNOLOGIES IN THE NEXT 10-15 YEARS (AND BEYOND)

Robert E. Horn, Visiting Scholar, Stanford University

Visual language is one of the more promising avenues to the improvement of human performance in the short run (the next 10 to 15 years) (Horn 2000b, 2000c). The current situation is one of considerable diversity and confusion as a new form of communication arises. But visual language also represents many great opportunities. People think visually. People think in language. When words and visual elements are closely intertwined, we create something new and we augment our communal intelligence.

Today, human beings work and think in fragmented ways, but visual language has the potential to integrate our existing skills to make them tremendously more effective. With support from developments in information technology, visual language has the potential for increasing human "bandwidth," the capacity to take in, comprehend, and more efficiently synthesize large amounts of new information. It has this capacity on the individual, group, and organizational levels. As this convergence occurs, visual language will enhance our ability to communicate, teach, and work in fields such as nanotechnology and biotechnology.

Definition

Visual language is defined as the tight integration of words and visual elements and has characteristics that distinguish it from natural languages as a separate communication tool as well as a distinctive subject of research. It has been called visual language, although it might well have been called visual-verbal language.

A preliminary syntax, semantics, and pragmatics of visual language have been described. (Horn 1998) Description of, understanding of, and research on visual language overlap with investigations of scientific visualization and multimedia.

History

The tight integration of words and visual elements has a long history (Horn 1998, Chapter 2). Only in the last 50 years, with the coming together of component visual

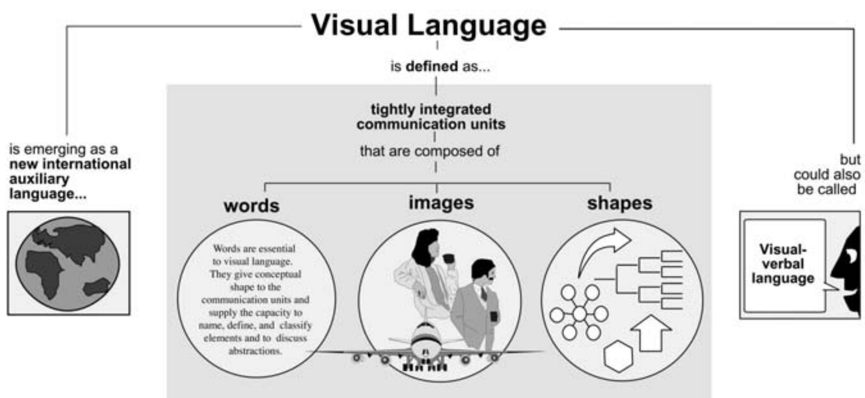

Figure B.8. Defining visual language.

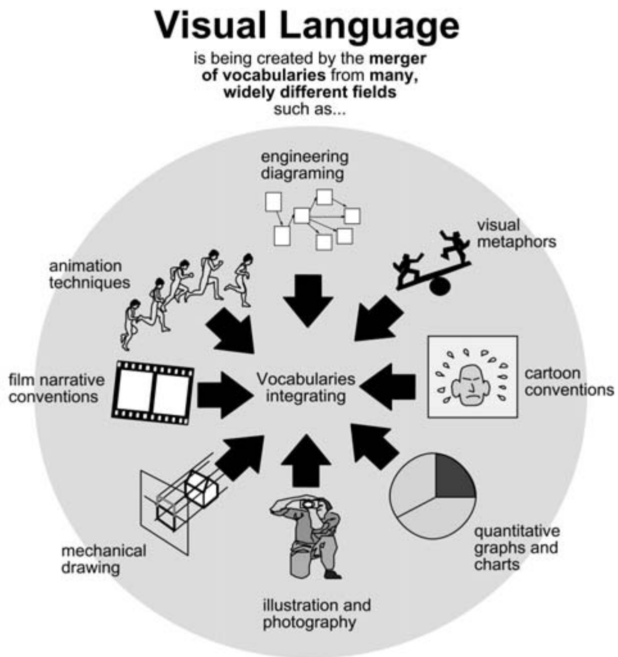

Figure B.9. Creation of visual language.

vocabularies from such widely separate domains as engineering diagramming technologies developed in medical illustration and hundreds of expressive visual conventions from the world of cartooning, has something resembling a full, robust visual-verbal language appeared (Tufte 1983, 1990).

Its evolution has been rapid in the past 10 years, especially with the confluence of scientific visualization software; widespread use of other quantitative software that permits the creation of over one hundred quantitative graphs and charts with the push of a single function key; and the profusion of multimedia presentation software, especially PowerPoint which, it is said, has several million users a day.

The Promise of *More Effective Communication*

There is widespread understanding that visual-verbal language enables forms and efficiencies of communication that heretofore have not been possible. For example, improvements in human performance from 23% to 89% have been obtained by using integrated visual-verbal stand-alone diagrams. In this case, stand-alone diagrams refer to diagrams that have all the verbal elements necessary for complete understanding without reading text elsewhere in a document (Chandler and Sweller 1991; Mayer 2001; Horton 1991).

There are several key advantages of the emerging visual-verbal language:

1. **It facilitates representation.** This new language facilitates presentation of complex, multidimensional visual-verbal thought, and — with multimedia tools — can incorporate animation, as well. Researchers and scholars are no longer constrained by the scroll-like thinking of endless paragraphs of text.

2. **It facilitates big, complex thoughts**. Human cognitive effectiveness and efficiency is constrained by the well-known limitations of working memory that George Miller identified in 1957 (Miller 1957). Large visual displays have for some time been known to help us overcome this bandwidth constraint. But only since the recent advances in visual language have we been able to imagine a major prosthesis for this human limitation. The prosthesis consists of a suite of visual language maps. This visual-verbal language (together with computer-based tools) may eliminate the major roadblocks to thinking and communicating big, complex thoughts, i.e., the problem of representing and communicating mental models of these thoughts efficiently and effectively.

This especially includes the so-called “messy” (or “wicked” or “ill-structured”) problems (Horn 2001a). Problems have straightforward solutions; messy problems do not. They are

- more than complicated and complex; they are ambiguous
- filled with considerable uncertainty — even as to what the conditions are, let alone what the appropriate actions might be
- bounded by great constraints and tightly interconnected economically, socially, politically, and technologically

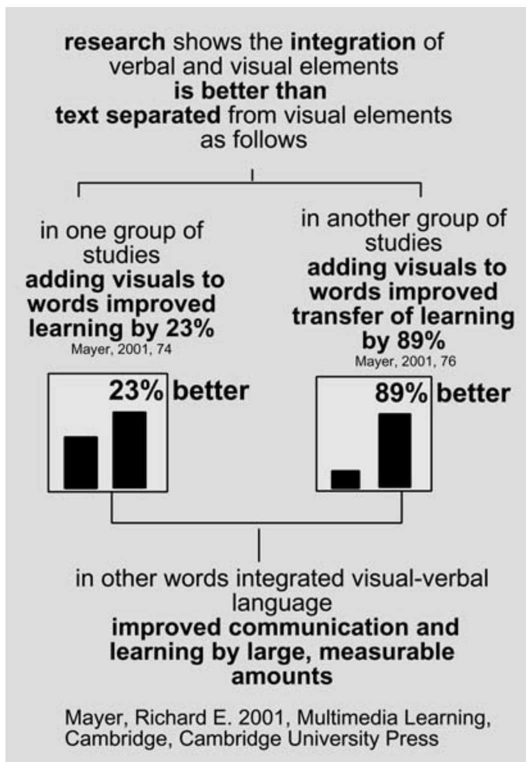

Figure B.10. Enhancing learning through visual language.

- seen differently from different points of view and quite different worldviews
- comprised of many value conflicts
- often alogical or illogical

These kinds of problems are among the most pressing for our country, for the advancement of civilization, and for humanity; hence, the promise of better representation and communication of complex ideas using visual-verbal language constructs has added significance.

Premises Regarding Visual Language

A deep understanding of the patterns of visual language will permit the following:

- more rapid, more effective interdisciplinary communication
- more complex thinking, leading to a new era of thought
- facilitation of business, government, scientific, and technical productivity
- potential breakthroughs in education and training productivity
- greater efficiency and effectiveness in all areas of knowledge production and distribution
- better cross-cultural communication

Readiness for Major Research and Development

A number of major jumping-off research platforms have already been created for the rapid future development of visual language: the Web; the ability to tag content with XML; database software; drawing software; a fully tested, widely used content-organizing and tagging system of structured writing known as Information Mapping® (Horn 1989); and a growing, systematic understanding of the patterns of visual-verbal language (Kosslyn 1989, 1994; McCloud 1993; Horton 1991; Bertin 1983).

Rationale for the Visual Language Projects

A virtual superhighway for rapid development in visual language can be opened, and the goals listed above in the premises can be accomplished, if sufficient funds over the next 15 years are applied to the creation of tools, techniques, and taxonomies, and to systematically conducting empirical research on effectiveness and efficiency of components, syntax, semantics, and pragmatics of this language. These developments, in turn, will aid the synergy produced in the convergence of biotechnology, nanotechnology, information technology, and cognitive science.

Goals of a Visual-Verbal Language Research Program

A research program requires both bold, general goals and specific landmarks along the way. A major effort to deal with the problem of increasing complexity and the limitations of our human cognitive abilities would benefit all human endeavors and could easily be focused on biotechnology and nanotechnology as prototype test beds. We can contemplate, thus, the steady, incremental achievement of the following goals as a realistic result of a major visual language program:

1. **Provide policymakers with comprehensive visual-verbal models**. The combination of the ability to represent complex mental models and the ability to collect realtime data will provide sophisticated decision-making tools for social policy. Highly visual cognitive maps will facilitate the management of and navigation through major public policy issues. These maps provide patterned abstractions of policy landscapes that permit the decisionmakers and their advisors to consider which roads to take within the wider policy context. Like the hundreds of different projections of geographic maps (e.g., polar or Mercator), they provide different ways of viewing issues and their backgrounds. They enable policymakers to drill down to the appropriate level of detail. In short, they provide an invaluable information management tool.
2. **Provide world-class, worldwide education for children.** Our children will inherit the results of this research. It is imperative that they receive the increased benefits of visual language communication research as soon as it is developed**.** The continued growth of the Internet and the convergence of intelligent visual-verbal representation of mental models and computer-enhanced tutoring programs will enable children everywhere to learn the content and skills needed to live in the 21st century. But this will take place only if these advances are incorporated into educational programs as soon as they are developed.
3. **Achieved large breakthroughs in scientific research.** The convergence of more competent computers, computer-based collaborative tools, visual representation breakthroughs, and large databases provided by sensors will enable major improvements in scientific research. Many of the advances that we can imagine will come from interdisciplinary teams of scientists, engineers, and technicians who will need to become familiar rapidly with fields that are outside of their backgrounds and competencies. Visual language resources (such as the diagram project described below) will be required at all levels to make this cross-disciplinary learning possible. This could be the single most important factor in increasing the effectiveness of nano-bio-info teams working together at their various points of convergence.
4. **Enrich the art of the 21st century.** Human beings do not live by information alone. We make meaning with our entire beings: emotional, kinesthetic, and somatic. Visual art has always fed the human spirit in this respect. And we can confidently predict that artistic communication and aesthetic enjoyment in the 21st century will be enhanced significantly by the scientific and technical developments in visual language. Dynamic visual-verbal murals and art pieces will become one of the predominant contemporary art forms of the century, as such complex, intense representation of meaning joins abstract and expressionistic art as a major artistic genre. This has already begun to happen, with artists creating the first generation of large visual language murals (Horn 2000).
5. **Develop smart, visual-verbal thought software.** The convergence of massive computing power, thorough mapping of visual-verbal language patterns, and advances in other branches of cognitive science will provide for an evolutionary leap in capacity and in multidimensionality of thought processes. Scientific visualization software in the past 15 years has led the

way in demonstrating the necessity of visualization in the scientific process. We could not have made advances in scientific understanding in many fields without software that helps us convert “firehoses of data“ (in the vivid metaphor of the 1987 National Science Foundation report on scientific visualization) into visually comprehensible depictions of *quantitative* phenomena and simulations. Similarly, every scientific field is overwhelmed with *tsunami*s of new *qualitative* concepts, procedures, techniques, and tools. Visual language offers the most immediate way to address these new, highly demanding requirements.

6. **Open wide the doors of creativity.** Visualization in scientific creativity has been frequently cited. Einstein often spoke of using visualization on his *gedanken* experiments. He saw in his imagination first and created equations later. This is a common occurrence for scientists, even those without special training. Visual-verbal expression will facilitate new ways of thinking about human problems, dilemmas, predicaments, emotions, tragedy, and comedy. “The limits of my language are the limits of my world,” said Wittgenstein. But it is in the very nature of creativity for us to be unable to specify what the limits will be. Indeed, it is not always possible to identify the limits of our worlds until some creative scientist has stepped across the limit and illuminated it from the other side.

Researchers in biotechnology and nanotechnology will not have to wait for the final achievement of these goals to begin to benefit from advances in visual language research and development. Policymakers, researchers, and scholars will be confronting many scientific, social, ethical, and organizational issues; each leap in our understanding and competence in visual language will increase our ability to deal with these kinds of complex issues. As the language advances in its ability to handle complex representation and communication, each advance can be widely disseminated because of the modular nature of the technology.

Major Objectives Towards Meeting Overall Goals of Visual-Verbal Language Research

The achievement of the six goals described above will obviously require intermediate advances on a number of fronts to achieve specific objectives:

1. **Diagram an entire branch of science with stand-alone diagrams.** In many of the newer introductory textbooks in science, up to one-third of the total space consists of diagrams and illustrations. But often, the function of scientific diagrams in synthesizing and representing scientific processes has been taken for granted. However, recent research cited above (Mayer 2001, Chandler and Sweller 1991) has shown how stand-alone diagrams can significantly enhance learning. Stand-alone diagrams do what the term indicates: everything the viewer needs to understand the subject under consideration is incorporated into one diagram or into a series of linked diagrams. The implication of the research is that the text in the other two thirds of the textbooks mentioned above should be distributed into diagrams.

   “Stand-alone” is obviously a relative term, because it depends on previous learning. One should note here that automatic prerequisite linkage is one of the easier functions to imagine being created in software packages designed to handle linked diagrams. One doesn't actually have to take too large a leap of imagination to see this as achievable, as scientists are already exchanging PowerPoint slides that contain many diagrams. However, this practice frequently does not take advantage of either the stand-alone or linked property.

Stand-alones can be done at a variety of styles and levels of illustration. They can be abstract or detailed, heavily illustrated or merely shapes, arrows, and words. They can contain photographs and icons as well as aesthetically pleasing color.

Imagine a series of interlinked diagrams for an entire field of science. Imagine zooming in and out — always having the relevant text immediately accessible. The total number of diagrams could reach into the tens of thousands. The hypothesis of this idea is that such a project could provide an extraordinary tool for cross-disciplinary learning. This prospect directly impacts the ability of interdisciplinary teams to learn enough of each other's fields in order to collaborate effectively. And collaboration is certainly the key to benefiting from converging technologies.

Imagine, further, that using and sharing these diagrams were *not* dependent on obtaining permission to reproduce them, which is one of the least computerized, most time-consuming tasks a communicator has to accomplish these days. Making permission automatic would remove one of the major roadblocks to the progress of visual language and a visual language project.

Then, imagine a scientist being able to send a group of linked, stand-alone diagrams to fellow scientists.

2. **Create "periodic" table(s) of types of stand-alone diagrams.** Once we had tens of thousands of interlinked diagrams in a branch of science, we could analyze and characterize all the components, structures, and functions of all of the types of diagrams. This would advance the understanding of "chunks of thinking" at a fine-grained level. This meta understanding of diagrams would also be a jumping-off point for building software tools to support further investigations and to support diagramming of other branches of science and the humanities.

3. **Automatically create diagrams from text.** At the present moment, we do not know how to develop software that enables the construction from text of a wide variety of kinds of elaborate diagrams. But if the stand-alone diagrams prove as useful as they appear, then an automatic process to create diagrams, or even just first drafts of diagrams, from verbal descriptions will turn out to be extremely beneficial. Imagine scientists with new ideas of how processes work speaking to their computers and the computers immediately turning the idea into the draft of a stand-alone diagram.

4. **Launch a project to map the human cognome.** In the Converging Technologies workshop I suggested that we launch a project that might be named "Mapping the Human Cognome." If properly conceived, such a project would certainly be the project of the century. If the stand-alone diagram project succeeds, then we would have a different view of human thought chunks. Since human thought-chunks can be understood as fundamental building blocks of the human cognome, the rapid achievement

of stand-alone diagrams for a branch of science could, thus, be regarded as a starting point for at least one major thrust of the Human Cognome Project (Horn 2002c).

5. **Create tools for collaborative mental models based on diagramming.** Ability to come to rapid agreement at various stages of group analysis and decision-making with support from complex, multidimensional, visual-verbal murals is becoming a central component of effective organizations. This collaborative problem-solving, perhaps first envisioned by Douglas Engelbart (1962) as augmenting human intellect, has launched a vibrant new field of computer-supported collaborative work (CSCW). The CSCW community has been facilitating virtual teams working around the globe on the same project in a 24/7 asynchronous timeframe. Integration of (1) the resources of visual language display, (2) both visual display hardware and software, and (3) the interactive potential of CSCW offers possibilities of great leaps forward in group efficiency and effectiveness.
6. **Crack the unique address dilemma with fuzzy ontologies.** The semantic web project is proceeding on the basis of creating unique addresses for individual chunks of knowledge. Researchers are struggling to create "ontologies," by which they mean hierarchical category schemes, similar to the Dewey system in libraries. But researchers haven't yet figured out really good ways to handle the fact that most words have multiple meanings. There has been quite a bit of progress in resolving such ambiguities in machine language translation, so there is hope for further incremental progress and major breakthroughs. An important goal for cognitive scientists will be to produce breakthroughs for managing the multiple and changing meanings of visual-verbal communication units on the Web in real time.
7. **Understand computerized visual-verbal linkages.** Getting computers to understand the linkage between visual and verbal thought and their integration is still a major obstacle to building computer software competent to undertake the automatic creation of diagrams. This is likely to be less of a problem as the stand-alone diagram project described above (objective #1) progresses.
8. **Crack the "context" problem.** In meeting after meeting on the subject of visual-verbal language, people remark at some point that "it all depends on the context." Researchers must conduct an interdisciplinary assault on the major problem of carrying context and meaning along with local meaning in various representation systems. This may well be accomplished to a certain degree by providing pretty good, computerized common sense. To achieve the goal of automatically creating diagrams from text, there will have to be improvements in the understanding of common sense by computers. The CYC project, the attempt to code all of human common sense knowledge into a single database — or something like it — will have to demonstrate the ability to reason with almost any subject matter from a base of 50 million or more coded facts and ideas. This common-sense database must somehow be integrally linked to visual elements.

Conclusion

It is essential to the accelerating research in the fields of nanotechnology, biotechnology, information technology, and cognitive science that we increase our understanding of visual language. In the next decade, we must develop visual language research centers, fund individual researchers, and ensure that these developments are rapidly integrated into education and into the support of the other converging technologies.

References

Bertin, J. 1983. *Semiology of graphics: Diagrams, networks, and maps.* Madison, WI: Univ. of Wisconsin Press.

Chandler, P., and J. Sweller. 1991. Cognitive load theory and the format of instruction. *Cognition and Instruction* 8(4): 293-332.

Engelbart, D.C. 1962. *Augmenting human intellect: A conceptual framework.* Stanford Research Institute, Washington, D.C.: Air Force Office Of Scientific Research, AFOSR-3233, Contract AF49(638)-1024 , SRI Project No. 3578. October.

Horn, R.E. 1989. *Mapping hypertext*, Lexington, MA: The Lexington Institute (http://www.stanford.edu/ ~rhorn/MHContents.html).

Horn, R.E. 1998a. *Mapping great debates: Can computers think?* Bainbridge Island, WA: MacroVU, Inc. (http://www.stanford.edu/~rhorn/CCTGeneralInfo.html).

Horn, R.E. 1998b. *Visual language: Global communication for the 21st century.* Bainbridge Island, WA: MacroVU, Inc. (http://www.stanford.edu/~rhorn/VLBkDescription.html).

Horn, R.E. 2000. *The representation of meaning—Information design as a practical art and a fine art.* A speech at the Stroom Center for the Visual Arts, The Hague (http://www.stanford.edu/~rhorn/ VLbkSpeechMuralsTheHague.html).

Horn, R.E. 2001a. *Knowledge mapping for complex social messes.* A speech to the Packard Foundation Conference on Knowledge Management (http://www.stanford.edu/~rhorn/ SpchPackard.html).

Horn, R.E. 2001b. *What kinds of writing have a future?* A speech prepared in connection with receiving Lifetime Achievement Award by the Association of Computing Machinery SIGDOC, October 22.

Horn, R.E. 2002a. *Think link, invent, implement, and collaborate! Think open! Think change! Think big!* Keynote Speech at Doug Engelbart Day in the State of Oregon, Oregon State University, Corvalis OR, January 24.

Horn, R.E. 2002b. *Conceptual map of a vision of the future of visual language research.* (To download PDF file http://www.stanford.edu/~rhorn/MapFutureVisualLang.html).

Horn, R.E. 2002c. *Beginning to conceptualize the human cognome project*. A paper prepared for the National Science Foundation Conference on Converging Technologies (Nano-Bio-Info-Cogno) (To download PDF file: http://www.stanford.edu/~rhorn/ArtclCognome.html).

Horton, W. 1991. *Illustrating computer documentation: The art of presenting information graphically in paper and online,* N.Y.: Wiley.

Kosslyn, S.M. 1989. Understanding charts and graphs. *Applied Cognitive Psychology* 3: 185-226.

Kosslyn, S.M. 1994. *Elements of graph design.* N.Y.: W.H. Freeman.

McCloud, S. 1993. *Understanding comics: The invisible art.* Northampton, MA: Kitchen Sink Press.

Mayer, R.E. 2001. *Multimedia learning*. Cambridge: Cambridge Univ. Press.

Tufte, E. 1983. *The visual display of quantitative information*. Cheshire, CT: Graphics Press.

Tufte, E. 1990. *Envisioning information.* Cheshire, CT: Graphics Press.

Sociable Technologies: Enhancing Human Performance When the Computer is not a Tool but a Companion

Sherry Turkle, Massachusetts Institute of Technology

> *"Replacing human contact [with a machine] is an awful idea. But some people have no contact [with caregivers] at all. If the choice is going to a nursing home or staying at home with a robot, we think people will choose the robot."* Sebastian Thrun, Assistant Professor of Computer Science, Carnegie Mellon University

> *"AIBO [Sony's household entertainment robot] is better than a real dog. It won't do dangerous things, and it won't betray you. ...Also, it won't die suddenly and make you feel very sad.'"* A 32-year -old woman on the experience of playing with AIBO

> "*Well, the Furby is alive for a Furby. And you know, something this smart should have arms. It might want to pick up something or to hug me."* Ron, age six, answering the question, "Is the Furby alive?"

Artificial intelligence has historically aimed at creating objects that might improve human performance by offering people intellectual complements. In a first stage, these objects took the form of tools, instruments to enhance human reasoning, such as programs used for medical diagnosis. In a second stage, the boundary between the machine and the person became less marked. Artificial intelligence technology functioned more as a prosthetic, an extension of human mind. In recent years, even the image of a program as prosthetic does not capture the intimacy people have with computational technology. With "wearable" computing, the machine comes closer to the body, ultimately continuous with the body, and the human person is redefined as a cyborg. In recent years, there has been an increased emphasis on a fourth model of enhancing human performance through the use of computation: technologies that would improve people by offering new forms of social relationships. The emphasis in this line of research is less on how to make machines "really" intelligent (Turkle 1984, 1995) than on how to design artifacts that would cause people to experience them as having subjectivities that are worth engaging with.

The new kind of object can be thought of as a relational artifact or as a sociable technology. It presents itself as having affective states that are influenced by the object's interactions with human beings. Today's relational artifacts include children's playthings (such as Furbies, Tamagotchis, and My Real Baby dolls); digital dolls and robots that double as health monitoring systems for the elderly (Matsushita's forthcoming Tama, Carnegie Mellon University's Flo and Pearl); and pet robots aimed at the adult (Sony's AIBO, MIT's Cog and Kismet). These objects are harbingers of a new paradigm for computer-human interaction.

In the past, I have often described the computer as a Rorschach. When I used this metaphor I was trying to present the computer as a relatively neutral screen onto which people were able to project their thoughts and feelings, a mirror of mind and self. But today's relational artifacts make the Rorschach metaphor far less useful. The computational object is no longer affectively "neutral." Relational artifacts do not so much invite projection as demand engagement. People are learning to interact with computers through conversation and gesture. People are learning that to relate successfully to a computer you do not have to know how it works but can take it "at interface value," that is, assess its emotional "state," much as you would if you were relating to another person. Through their experiences with virtual pets and digital dolls, which present themselves as loving and responsive to care, a generation of children is learning that some objects require emotional nurturing and some even promise it in return. Adults, too, are encountering technology that attempts to offer advice, care, and companionship in the guise of help-software-embedded wizards, intelligent agents, and household entertainment robots such as the AIBO "dog."

New Objects are Changing Our Minds

Winston Churchill once said, "We make our buildings and then they make us." We make our technologies, and they in turn shape us. Indeed, there is an unstated question that lies behind much of our historic preoccupation with the computer's capabilities. That question is not what can computers do or what will computers be like in the future, but instead, what will we be like? What kind of people are we becoming as we develop more and more intimate relationships with machines? The new technological genre of relational, sociable artifacts is changing the way we think. Relational artifacts are new elements in the categories people use for thinking about life, mind, consciousness, and relationship. These artifacts are well positioned to affect people's way of thinking about themselves, about identity, and about what makes people special, influencing how we understand such "human" qualities as emotion, love, and care. We will not be taking the adequate measure of these artifacts if we only consider what they do *for* us in an *instrumental* sense. We must explore what they do not just for us but to us as people, to our relationships, to the way our children develop, to the way we view our place in the world.

There has been a great deal of work on how to create relational artifacts and maximize their ability to evoke responses from people. Too little attention, however, has gone into understanding the human implications of this new computational paradigm, both in terms of how we relate to the world and in terms of how humans construct their sense of what it means to be human and alive. The language for assessing these human implications is enriched by several major traditions of thinking about the role of objects in human life.

Objects as Transitional to Relationship

Social scientists Claude Levi-Strauss (1963), Mary Douglas (1960), Donald Norman (1988), Mihaly Csikzentmihalyi (1981), and Eugene Rochberg-Halton (1981) have explored how objects carry ideas, serving as enablers of new individual and cultural meanings. In the psychoanalytic tradition Winnicott (1971) has discussed how objects mediate between the child's earliest bond with the mother, who the infant experiences as inseparable from the self, and the child's growing capacity to develop relationships with other people, who will be experienced as separate beings.

In the past, the power of objects to act in this transitional role has been tied to the ways in which they enabled the child to project meanings onto them. The doll or the teddy bear presented an unchanging and passive presence. Relational artifacts take a more active stance. With them, children's expectations that their dolls want to be hugged, dressed, or lulled to sleep don't come from the child's projection of fantasy or desire onto inert playthings, but from such things as a digital doll's crying inconsolably or even saying, "Hug me!" "It's time for me to get dressed for school!" The psychology of the playroom turns from projection to social engagement, in which data from an active and unpredictable object of affection helps to shape the nature of the relationship. On the simplest level, when a robotic creature makes eye contact, follows your gaze, and gestures towards you, what you feel is the evolutionary button being pushed to respond to that creature as a sentient and even caring other.

Objects as Transitional to Theories of Life

The Swiss psychologist Jean Piaget addressed some of the many ways in which objects carry ideas (1960). For Piaget, interacting with objects affects how the child comes to think about space, time, the concept of number, and the concept of life. While for Winnicott and the object relations school of psychoanalysis, objects bring a world of people and relationships inside the self, for Piaget objects enable the child to construct categories in order to make sense of the outer world. Piaget, studying children in the context of non-computational objects, found that as children matured, they homed in on a definition of life that centered around "moving of one's own accord." First, everything that moved was taken to be alive, then only those things that moved without an outside push or pull. Gradually, children refined the notion of "moving of one's own accord" to mean the "life motions" of breathing and metabolism.

In the past two decades, I have followed how computational objects change the ways children engage with classic developmental questions such as thinking about the property of "aliveness." From the first generation of children who met computers and electronic toys and games (the children of the late 1970s and early 1980s), I found a disruption in this classical story. Whether or not children thought their computers were alive, they were sure that how the toys moved was not at the heart of the matter. Children's discussions about the computer's aliveness came to center on what the children perceived as the computer's psychological rather than physical properties (Turkle 1984). Did the computer know things on its own or did it have to be programmed? Did it have intentions, consciousness, feelings? Did it cheat? Did it know it was cheating? Faced with intelligent machines, children took a new world of objects and imposed a new world order. To put it too simply, motion gave way to emotion, and physics gave way to psychology as criteria for aliveness.

By the 1990s, that order had been strained to the breaking point. Children spoke about computers as just machines but then described them as sentient and intentional. They talked about biology, evolution. They said things like, "the robots are in control but not alive, would be alive if they had bodies, are alive because they have bodies, would be alive if they had feelings, are alive the way insects are alive but not the way people are alive; the simulated creatures are not alive because they are just in the computer, are alive until you turn off the computer, are not alive because nothing in the computer is real; the Sim creatures are not alive but almost-alive, they would be alive if they spoke, they would be alive if they traveled, they're not alive because they don't have bodies, they are alive because they can have babies and would be alive if they could get out of the game and onto America Online."

There was a striking heterogeneity of theory. Children cycled through different theories to far more fluid ways of thinking about life and reality, to the point that my daughter upon seeing a jellyfish in the Mediterranean said, “Look, Mommy, a jellyfish; it looks so realistic!” Likewise, visitors to Disney’s Animal Kingdom in Orlando have complained that the biological animals that populated the theme park were not “realistic” compared to the animatronic creatures across the way at Disneyworld.

By the 1990s, children were playing with computational objects that demonstrated properties of evolution. In the presence of these objects, children’s discussions of the aliveness question became more complex. Now, children talked about computers as “just machines” but described them as sentient and intentional as well. Faced with ever more sophisticated computational objects, children were in the position of theoretical tinkerers, “making do” with whatever materials were at hand, “making do” with whatever theory could be made to fit a prevailing circumstance (Turkle 1995).

Relational artifacts provide children with a new challenge for classification. As an example, consider the very simple relational artifact, the “Furby.” The Furby is an owl-like interactive doll, activated by sensors and a pre-programmed computer chip, which engages and responds to their owners with sounds and movement. Children playing with Furbies are inspired to compare and contrast their understanding of how the Furby works to how they “work.” In the process, the line between artifact and biology softens. Consider this response to the question, “Is the Furby alive?”

> Jen (age 9): I really like to take care of it. So, I guess it is alive, but it doesn’t need to really eat, so it is as alive as you can be if you don’t eat. A Furby is like an owl. But it is more alive than an owl because it knows more and you can talk to it. But it needs batteries so it is not an animal. It’s not like an animal kind of alive.

Jen’s response, like many others provoked by playing with Furbies, suggests that today’s children are learning to distinguish between an “animal kind of alive” and a “Furby kind of alive.” In my conversations with a wide range of people who have interacted with relational artifacts — from five year olds to educated adults — an emergent common denominator has been the increasingly frequent use of “sort of alive” as a way of dealing with the category confusion posed by relational artifacts. It is a category shared by the robots’ designers, who have questions about the ways in which their objects are moving toward a kind of consciousness that might grant them a new moral status.

Human-Computer Interaction

The tendency for people to attribute personality, intelligence, and emotion to computational objects has been widely documented in the field of human-computer interaction (HCI) (Weizenbaum 1976; Nass, Moon, et al. 1997, Kiesler and Sproull 1997; Reeves and Nass 1999). In most HCI work, however, this “attribution effect” is considered in the context of trying to build “better” technology.

In *Computers are Social Actors: A Review of Current Research*, Clifford Nass, Youngme Moon, and their coauthors (1997) review a set of laboratory experiments in which “individuals engage in social behavior towards technologies even when such behavior is entirely inconsistent with their beliefs about machines" (p. 138). Even when computer-based tasks contained only a few human-like characteristics, the authors found that subjects attributed personality traits and gender to computers and adjusted their responses to avoid hurting the machines' "feelings." The authors suggest that "when we are confronted with an entity that [behaves in human-like ways, such as using language and responding based on prior inputs] our brains' default response is to unconsciously treat the entity as human" (p. 158). From this, they suggest design criteria: technologies should be made more "likeable":

> … "liking" leads to various secondary consequences in interpersonal relationships (e.g., trust, sustained friendship, etc.), we suspect that it also leads to various consequences in human-computer interactions (e.g., increased likelihood of purchase, use, productivity, etc.) (p. 138).

Nass et al. prescribe "likeability" for computational design. Several researchers are pursuing this direction. At the MIT Media Lab, for example, Rosalind Picard's Affective Computing research group develops technologies that are programmed to assess their users'' emotional states and respond with emotional states of their own. This research has dual agendas. On the one hand, affective software is supposed to be compelling to users — "friendlier," easier to use. On the other hand, there is an increasing scientific commitment to the idea that objects need affect in order to be intelligent. As Rosalind Picard writes in *Affective Computing* (1997, x),

> I have come to the conclusion that if we want computers to be genuinely intelligent, to adapt to us, and to interact naturally with us, then they will need the ability to recognize and express emotions, to have emotions, and to have what has come to be called "emotional intelligence."

Similarly, at MIT's Artificial Intelligence Lab, Cynthia Breazeal has incorporated both the "attribution effect" and a sort of "emotional intelligence" in Kismet. Kismet is a disembodied robotic head with behavior and capabilities modeled on those of a pre-verbal infant (see, for example, Breazeal and Scassellati 2000). Like Cog, a humanoid robot torso in the same lab, Kismet learns through interaction with its environment, especially contact with human caretakers. Kismet uses facial expressions and vocal cues to engage caretakers in behaviors that satisfy its "drives" and its "emotional" needs. The robot "wants" to be happy, and people are motivated to help it achieve this goal. Its evocative design seems to help, Breazeal reports: "When people see Cog they tend to say, 'That's interesting.' But with Kismet they tend to say, 'It smiled at me!' or 'I made it happy!'" (Whynott 1999). I have seen similar reactions between children and simpler digital pets (both on the screen, such as neopets and in robotic form, such as Furbies and AIBOs).

When children play with Furbies, they want to know the objects' "state," not to get something "right," but to make the Furbies happy. Children want to understand Furby language, not to "win" in a game over the Furbies, but to have a feeling of mutual recognition. When I asked her if her Furby was alive, Katherine, age five, answered in a way that typifies this response:

> "Is it alive? Well, I love it. It's more alive than a Tamagotchi because it sleeps with me. It likes to sleep with me."

Children do not ask how the Furbies "work" in terms of underlying process; they take the affectively charged toys "at interface value."

With the advent of relational artifacts and their uses of emotion, we are in a different world from the old AI debates of the 1960s to 1980s, in which researchers argued about whether machines could be "really" intelligent. The old debate was essentialist; these new objects allow researchers and their public to sidestep such arguments about what is inherent in the computer. Instead, they focus attention on what the objects evoke in us. When we are asked to care for an object (the robot Kismet or the plaything Furby), and when the cared-for object thrives and offers us its attention and concern, we experience that object as intelligent. Beyond this, we feel a connection to it. So the issue here is not whether objects "really" have emotions, but what is happening when relational artifacts evoke emotional responses in the users.

People's relationships with relational artifacts have implications for technological design (i.e., how to make the objects better, more compelling), and they have implications that are the focus of this research: they complicate people's ways of thinking about themselves, as individuals, as learners and in relationships, and within communities. To augment human potential, any discussion of how to make "better" relational artifacts must be in terms of how they can best enhance people in their human purposes. It cannot be discussed in terms of any absolute notions defined solely in terms of the objects.

The questions raised by relational artifacts speak to people's longstanding fears and hopes about technology, and to the question of what is special about being human, what is the nature of "personhood." In the case of relational technology, there is a need for examination of these questions, beginning with how these objects are experienced in the everyday lives of the individuals and groups who are closest to them.

Human Performance

When people learn that AIBO, the Sony robot dog, is being introduced into nursing homes as companions to the elderly, the first question asked is usually, "Does it work?" By this the person means, "Are the old people happier when they have a robot pet? Are they easier to take care of?" My vision of the future is that we are going to have increasingly intimate relationships with sociable technologies, and we are going to need to ask increasingly complex questions about the kinds of relationships we form with them. The gold standard cannot be whether these objects keep babies and/or the elderly "amused" or "quiet" or "easier to care for." Human performance needs to be defined in a much more complex way, beginning with a set of new questions that take the new genre of objects seriously. Taking them seriously means addressing them as new social interlocutors that will bring together biology, information science, and nanoscience. Human performance needs to take into account the way we feel about ourselves as people, in our relationships and in our social groups. From this point of view, the question for the future is not going to be whether children love their robots more than their parents, but what loving itself comes to mean. From this perspective on human enhancement, some of the questions are

- How are children adapting ideas about aliveness, intentionality, and emotion to accommodate relational artifacts?
- How are designers and early adopters adapting ideas about personhood, intentionality, and relationship to accommodate relational artifacts? How do these artifacts influence the way people think about human minds?
- How are people thinking about the ethical issues raised by relational artifacts? Is a moral code for the treatment of this new type of artifacts being developed?
- How are people using relational artifacts to address needs traditionally met by other humans and animal pets, such as companionship and nurturing?

A Vision Statement

Computational objects are "evocative objects." They raise new questions and provoke new discourse about the nature of mind, about what it means to be alive, about what is special about being a person, about free will and intentionality. Computation brings philosophy into everyday life. Objects as simple as computer toys and games raise such questions as "What is intelligence? What does it mean to be alive? Or to die? What is the nature of the self? What is special about being a person?" In the next 10 to 20 years, research that will marry biology, information science, cognitive science, and nanoscience is going to produce increasingly sophisticated relational, sociable artifacts that will have the potential to profoundly influence how people think about learning, human development, intelligence, and relationships.

- As research on relational and sociable technology progresses, there will be parallel investigations of how these objects affect the people who use them, how they influence psychological development, human relationships, and additionally, how they enter into people's thinking about themselves, including about such questions as the nature of intention, the self, and the soul.
- The development of sociable technologies will require a renaissance in the sciences that study human development and personality. There will be an increasing virtuous cycle of research to understand human personality and to create person-enhancing machines. Indeed, the notion of personable machines will come to mean person-enhancing machines.
- In the past, it has been argued that technology dehumanized life, but as we become committed to person-enhancing objects, this argument will need to be revisited. Making technology personable will entail learning about ourselves. In order to make technology enhance humans, we will humanize technology.
- Historically, when technology has been designed without human fulfillment in mind, but purely in terms of the instrumental capabilities of the machine, there has been a great deal of resistance to technology. This resistance needs to be taken seriously, because it points to the ways in which people associate technology with human loss. The development of sociable technology will require that there be a flourishing of research that takes resistance to technology as a symptom of something important that needs to be studied

rather than a problem that needs to be overcome. An understanding of human psychology is essential for the development of sociable technologies. This latter will proceed with vigilance and with the participation of humanists and scientists. Sociable technology will enhance human emotional as well as cognitive performance, not only giving us more satisfactory relationships with our machines but also potentially vitalizing our relationships with each other, because in order to build better sociable objects we will have learned more about what makes us social with each other.

References

Baltus, G., D. Fox, F. Gemperle, J. Goetz, T. Hirsch, D. Magaritis, M. Montemerlo, J. Pineau, N. Roy, J. Schulte, and S. Thrun. 2000. Towards personal service robots for the elderly. Published online at http://www.cs.cmu.edu/~thrun/papers/thrun.nursebot-early.html.

Breazeal (Ferrell), C., and B. Scassellati. 2000. Infant-like social interactions between a robot and a human caretaker. In *Adaptive behavior on simulation models of social agents,* Kerstin Dautenhahn, ed. (Special issue.)

Brooks, R., C. Breazeal, M. Marjanovic, B. Scassellati, and M. Williamson. 1998. The Cog project: Building a humanoid robot. In *Computation for metaphors, analogy and agents*, C. Nehaniv, ed. *Springer Lecture Notes in Artificial Intelligence*, Vol. 1562.

Brooks, R. 2000. Human-robot interaction and the future of robots in society. Presentation at the STS Colloquium, February 7, Massachusetts Institute of Technology, Cambridge, MA.

Csikzentmihalyi, M., and E. Rochberg-Halton. 1981. *The meaning of things: Domestic symbols and the self*. Cambridge: Cambridge University Press.

Dennett, D. [1987] 1998. *The intentional stance*. Reprint. Cambridge, MA: MIT Press.

Douglas, M. [1960] 1993. *Purity and danger: An analysis of the concepts of pollution and taboo*. Reprint. London: Routledge.

Drexler, M. 1999. Pet robots considered therapy for the elderly. *CNN Online*, March 12. http://www.cnn.com/TECH/ptech/9903/25/robocat.idg/.

Ito, M. 1997. Inhabiting multiple worlds: Making sense of SimCity 2000 in the fifth dimension. In *Cyborg Babies,* R.D. Floyd and J. Dumit, eds. New York: Routledge.

Kiesler, S. and L. Sproull. 1997. 'Social' human-computer interaction? In *Human Values and the Design of Computer Technology*, B. Friedman, ed. Stanford, CA: CSLI Publications.

Knorr-Cetina, K. 1997. Sociality with objects: Social relations in postsocial knowledge societies. *Theory, Culture and Society* 14(4): 1-30.

Latour, B. 1992. Where are the missing masses? The sociology of a few mundane artifacts. In *Shaping Technology/Building Society: Studies in Sociotechnical Change*, W. Bijker and J. Law, eds. Cambridge MA: MIT Press.

Levi-Strauss, C. 1963. *Structural anthropology*. New York: Basic Books.

Nass, C., Y. Moon, J. Morkes, E. Kim, and B.J. Fogg. 1997. Computers are social actors: A review of current research. In *Human values and the design of computer technology*, B. Friedman, ed. Stanford, CA: CSLI Publ..

Norman, D. 1988. *The design of everyday things*. New York: Currency/Double Day.

Ornstein, P.H., ed. 1978. *The search for the self: Selected writings of Heinz Kohut: 1950-1978, Vol. 2.* New York: International Universities Press, Inc.

Piaget, J. 1960. *The child's conception of the world.* Joan and Andrew Tomlinson, trans. Totowa, N.J.: Littlefield, Adams.

Picard, R.W. 1997. *Affective computing.* Cambridge, MA: MIT Press.

Reeves, B. and C. Nass. 1999. *The media equation: How people treat computers, television, and new media like real people and places*. Cambridge: Cambridge University Press.

Resnick, M. 1998. Technologies for lifelong kindergarten. *Educational technology research and development* 46:4.

Turkle, S. 1995. *Life on the screen: Identity in the age of the Internet*. New York: Simon and Schuster.
_____. 1984. *The second self: Computers and the human spirit*. New York: Simon and Schuster.
Weizenbaum, J. 1976. *Computer power and human reason: From judgement to calculation*. San Francisco: W.H. Freeman.
Whynott, D. 1999. The robot that loves people. *Discover,* October, online edition.
Winner, L. 1986. *The whale and the reactor: A search for limits in an age of high technology*. Chicago: University of Chicago Press.
Winnicott, D.W. [1971] 1989. *Playing and Reality*. Reprint. London: Routledge.
Zambrowski, J. 2000. CMU, Pitt developing "nursebot." *Tribune Review*, October 27, online edition.

VISIONARY PROJECTS

SOCIO-TECH…THE PREDICTIVE SCIENCE OF SOCIETAL BEHAVIOR

Gerold Yonas, Sandia National Laboratories,[1] *and Jessica Glicken Turnley, J., Galisteo Consulting Group, Inc.*

Socio-tech is the predictive — not descriptive — science of the behavior of societies. It is the convergence of information from the life sciences, the behavioral sciences (including psychology and the study of cognition), and the social sciences. Its data gathering and analysis approaches come from these fields and are significantly augmented by new tools from fields such as nanotechnology, engineering, and the information sciences. Agent-based simulations, models incorporating genetic algorithms, evolutionary computing techniques, and brain-machine interfaces provide new ways to gather data and to analyze the results.

Why Do We Care?

Most immediately, socio-tech can help us win the war on terrorism. It can help us to understand the motivations of the terrorists and so eliminate them. It also can help us to manage ourselves, to orchestrate our own country's response to a potential or real attack. In the longer term, as a predictive science, socio-tech can help us identify possible drivers for a wide range of socially disruptive events and allow us to put mitigating or preventative strategies in place before the fact.

What Is New?

The multiple drivers of human behavior have long been known. What have been missing are the theoretical paradigm and associated tools to integrate what we know about these drivers into an overarching understanding of human activity.

Currently, most of the data related to understanding human behavior has remained field-specific. The life sciences focus on the biological impacts of humans

[1] Sandia National Laboratories is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under Contract DE-AC04-94AL85000.

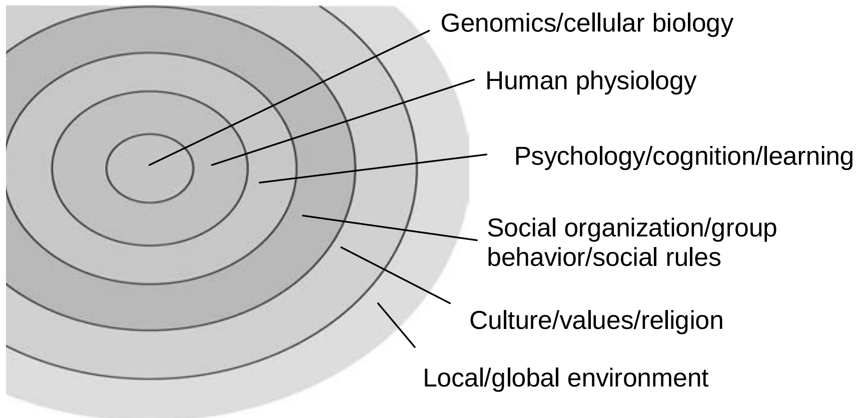

Figure B.11. Integrated studies of human behavior: Socio-tech.

functioning in physical spaces. The social sciences focus on the organizing principles of groups (rule of law, social hierarchies) and the different values groups place on behaviors (e.g., through culture or religion). The behavioral sciences are concerned with the functioning of the brain and the impact of individual experience on decision-making. The tools of science, engineering, and the information and computational sciences generally are not well integrated into these fields. C.P. Snow's 1959 Rede lecture captured this divide between the sciences on one hand and the arts and humanities on the other by the term "the two cultures."

There is little dialogue among practitioners from these different areas. They are separated by barriers of jargon, by conceptual frameworks that are difficult to translate from one field to another, and by traditional institutional compartmentalization of intellectual disciplines. Efforts such as Lewis Mumford's *Techniques and Human Development* (1989) to socially contextualize technology or E.O. Wilson's more recent and ambitious *Concilience* (1999) are the exceptions rather than the rule. We thus have no true study of human behavior, for there is no field or discipline with the interest or the tools to integrate data from these different fields. The challenge before us is to devise a way to understand data and information from each field in the context of all others. If genomics can be practiced with an awareness of human physiology, behavior, values, and environment, and, conversely, if information from genomics can be incorporated in a meaningful way into studies in these other fields, we will have made a significant leap in our understanding of human behavior (Figure B.11).

Why Now?

The time is ripe to begin such integration — to use the tremendous computing power we now have to integrate data across these fields to create new models and hence new understanding of the behavior of individuals. The ultimate goal is acquiring the ability to predict the behavior of an individual and, by extension, of groups. Recent advances in brain imaging, neuropsychology, and other sciences of the brain have significantly contributed to our knowledge of brain functioning. Genomics, molecular biology, and contributions from other areas in the life sciences

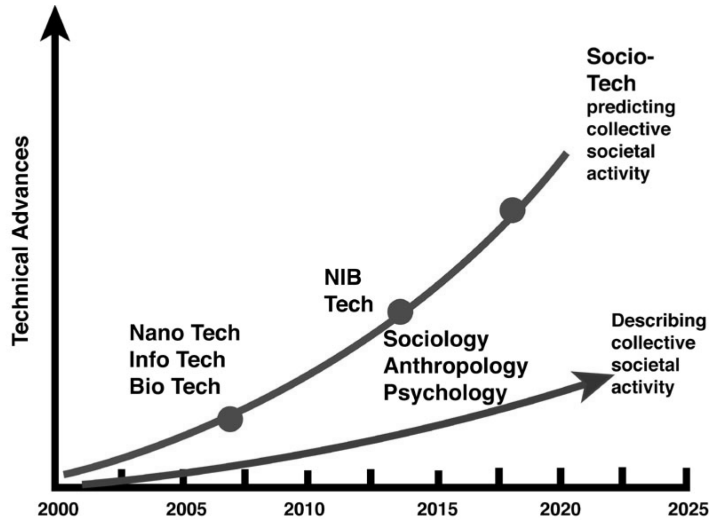

Figure B.12. Socio-tech: A qualitatively new science.

have greatly advanced our knowledge of the human body, its genetic core, and its response to various environmental stimuli. The increasing body of knowledge in the social sciences, combined with the tremendous computing (analysis) power available at affordable prices and new tools for communication and expression, have given us new ways of looking at social relationships such as social network theory, and new ways of understanding different ways of life. Incorporating these advances in a wide range of fields of study into overarching and integrating conceptual models should give us significant insights into human behavior.

Figure B.12 shows two possible trajectories for the development of knowledge. The upper trajectory combines the "two cultures," using technology to leverage the behavioral and social sciences and leads to a predictive science of behavior. The lower trajectory illustrates improvements in the behavioral and social sciences, with little incorporation of theory and tools from science and technology. It leads to greater descriptive but no predictive capabilities.

Socio-tech — the accumulation, manipulation, and integration of data from the life, social, and behavioral sciences, using tools and approaches provided by science and technology — will raise our ability to predict behaviors. It will allow us to interdict undesirable behaviors before they cause significant harm to others and to support and encourage behaviors leading to greater social goods.

References

Mumford, L. 1999. *Techniques and human developmen*t. 2 vols. New York, NY: Harcourt Brace Jovanovich.

Snow, C.P. 1959. *The two cultures.* Introduction by Stefan Collini. New York, NY: Cambridge University Press.

Wilson, E.O. 1999. *Concilience: The unity of knowledge*, New York, NY: Random House.

Breaking the Limits on Design Complexity

Jordan Pollack, Brandeis University

As we contemplate microelectromechanical systems (MEMS) and nanotechnologies (nano), we must study the history of the design of circuits and software, especially software that is supposed to have cognitive function, or artificial intelligence (AI). Having been working in the field of AI for 25 years, I can say with some authority that nanotechnology will not solve the AI problem. In fact, the repeated failures of AI artifacts to live up to claims made by their proponents can shed light on human expectations of nano and on the capacity of human teams to design complex objects.

We think that in order to design products "of biological complexity" that could make use of the fantastic fabrication abilities of new nano and MEMS factories, we must first liberate design by discovering and exploiting the principles of automatic self-organization that are seen in nature. A brain has 10^{11} connections. Chemistry often works with 10^{23} molecules. Advanced software is the most complex (and profitable) of all of human artifacts, yet each application only comprises between 10 million and 100 million lines of code, or a maximum of around 10^{8} moving parts. Suppose an animal brain, rather than requiring the specifying over time of the bonds for every molecule, ONLY required the equivalent of 10^{10} uniquely programmed parts. Why can't we engineer that?

In circuits, achieving even a function as lowly as the "bit" of memory creates a means to replication. Now we have 32 million bits on a single chip, and that is an achievement. Building blocks that can be replicated via manufacturing in hardware make things like memory chips and CPUs faster and more capable. This replication capacity and speedup of hardware enables Moore's law, a doubling of computer power, and even disk space, every 18 months. However, this periodic doubling of computer power has not led to equivalent doubling of human capacity to manufacture significantly more complex software. Moore's law does not solve the problem of engineering 10 billion lines of code!

The simple reason we haven't witnessed Moore's law operate for software is that 32 million copies of the same line of code is just one more line of code — the DO loop. Thus today's supercomputers run the same sized programs as the supercomputers of the 1970s, which are the desktops of today. The applications can use lots of floating point multiplication, but the complexity of the tasks hasn't grown beyond word processing, spreadsheets, and animations. Faster and faster computers seem to encourage software companies to write less and less efficient code for the same essential functionality — Windows is just DOS with wallpaper.

We've learned this hard lesson from the field of software — which isn't even constrained by material cost or by physical reality: *there are limits on the complexity of achievable design*. This is true even when throwing larger and larger teams of humans at a problem, even with the best groupware CAD software, even with bigger computers. Therefore, assumptions that new fabrication methodologies will lead to a breakthrough in design complexity ought to be taken with a grain of salt.

Yet many nano pundits expect that smaller-scale manufacturing, rather than leading to homogenous materials competitive with wood and plastic, will automatically lead to artificial objects of extraordinary complexity and near life-like capacity. They seem to ignore the technical challenges of understanding and modeling cognition, plugging portals into our brains, and programming Utility Fogs of nanobots that are intelligent enough to swarm and perform coordinated missions. The reality is that making life-sized artifacts out of molecules may require the arranging of 10^{30} parts.

AI is stalled because it is starved of the much more complex blueprints than anyone has any clue how to build. Software engineering seems to have reached a complexity limit well below what computers can actually execute. Despite new programming languages and various movements to revolutionize the field, the size of programs today is about the same as it has been for 40 years: 10-100 million lines of code. Old code finally collapses under the cost of its own maintenance.

The high-level languages, object-oriented programming systems, and computer-assisted software engineering (CASE) breakthroughs have all seemed promising, yet each new breakthrough devolves back into the same old thing in new clothes: the Fortran compiler plus vast scientific libraries. The power of each new programming tool, be it PL/1, Turbo Pascal, Visual Basic, Perl, or Java, is located in the bundled collections of subroutine libraries, which eventually grow to surpass our merely human cognitive ability to remember or even look them up in burgeoning encyclopedias.

The problem illustrated here is still Brooks' Mythical Man Month: We can't get bigger and better software systems by putting more humans on the job. The best original software, whether DOS, Lotus 123, or Wordstar, have been written by one or two good programmers; large teams extend, integrate, copy, and maintain, but they do not create. The more programmers on a task, the more bugs they create for each other.

The opportunity available today is that the way out of this tarpit, the path to achieving both software and nano devices of biological complexity with tens of billions of moving parts, is very clear: it is through increasing our scientific understanding of the processes by which biologically complex objects arose. As we understand these processes, we will be able to replicate them in software and electronics. The principles of automatic design and of self-organizing systems are a grand challenge to unravel. Fortunately, remarkable progress has been shown since the computer has been available to refine the theory of evolution. Software is being used to model life itself, which has been best defined as that "chemical reaction, far from equilibrium, which dissipates energy and locally reverses entropy."

Much as logic was unconstrained philosophy before computer automation, and as psychological and linguistic theories that could not be computerized were outgunned by formalizable models, theories on the origin of life, its intrinsic metabolic and gene regulation processes, and the mechanisms underlying major transitions in evolution, are being sharpened and refuted through formalization and detailed computer simulation.

Beyond the basic idea of a genetic algorithm, the variety of studies on artificial life, the mathematical and computational bases for understanding learning, growth, and evolution, are rapidly expanding our knowledge and our know-how.

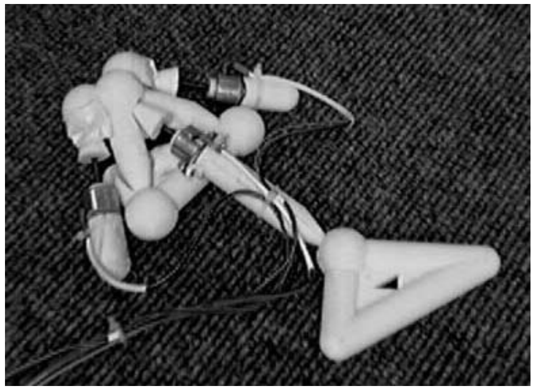

Figure B.13. Semi- and fully automatic design can help design complex systems like robot hardware and software.

My laboratory, which studies machine learning and evolutionary computation, has focused on how semi- and fully-automatic design can help design complex systems like robot hardware and software. We have used a collection of methods called "co-evolution," in which the idea is to create a sustained "arms-race" amongst or between populations of simple learning systems in order to achieve automatic design of various structures such as sorting nets, cellular automata rules, game players, and robot bodies and brains (Fig. B.13).

The field of evolutionary design, which aims at the creation of artifacts with less human engineering involvement, is in full force, documented by the books edited by Peter Bentley, as well as a NASA-sponsored annual conference on evolutionary hardware. Evolutionary robotics is a related field that started with Karl Sims' virtual robots and has grown significantly in the last five years.

So far, few artificial evolutionary processes have produced software or systems beyond those that can be designed by teams of humans. But they are competitive, and they are much cheaper than human designs. More importantly, thus far, they have not hit a barrier to complexity as seen in software engineering. Automatic design converts surplus computer time into complex design, and this will be aided by Moore's law. As inexpensive one-of-a-kind fabrication becomes possible, mass manufacture will no longer be necessary to amortize the fixed costs of engineering design, and automatic design will become necessary to generate complex designs with low cost. Success in this field holds keys to surpassing today's limits on complexity.

References

Bentley, P.J., ed.. 1999. *Evolutionary design by computers*. San Francisco: Morgan Kaufman.

Funes, P. and J.B. Pollack. 1998. Evolutionary body building: Adaptive physical designs for robots. *Artificial Life* 4: 337-357.

Holland, J.H. 1975. *Adaptation in natural and artificial systems*. Ann Arbor: University of Michigan Press.

Hornby, G.S., and J.B. Pollack. 2000. Evolving L-systems to generate virtual creatures. *Computers and Graphics* 25(6): 1041-1048.

Lipson, H., and J.B. Pollack. 2000. Automatic design and manufacture of robotic lifeforms. *Nature* 406:974-978.

Nolfi, S. and D. Floreano. 2000. Evolutionary robotics. The biology, intelligence, and technology of self-organizing machines. Cambridge, MA: MIT Press.

Pollack, J.B. and A.D. Blair. 1998. Co-evolution in the successful learning of backgammon strategy. *Machine Learning* 32(3): 225-240 (September).

Ray, T. 1992. An approach to the synthesis of life. In *Artificial Life II*, C. Langton, C. Taylor, J.F. and S. Rasmussen, eds. Reading, MA: Addison-Wesley.

Sims, K. 1994. Evolving virtual creatures. In SIGGRAPH 94 Conference Proceedings, Annual Conference Series, pages 15-22.

Thompson, A. 1997. "Artificial evolution in the physical world." In *Evolutionary Robotics: From intelligent robotics to artificial life*, T. Gomi, ed. Elsevier.

Enhancing Personal Area Sensory and Social Communication Through Converging Technologies

Rudy Burger, MIT Media Lab Europe

The next decade will see great strides in personal wearable technologies that enhance people's ability to sense their environment. This sensing will focus on at least two different areas:

a) *social sensing*, in which we may augment our ability to be aware of people in our immediate vicinity with whom we may wish to connect (or possibly avoid!)

b) *environmental sensing*, in which we may augment our ability to sense aspects of our environment (for example, the quality of the air we are breathing) that may be hazardous to us but that our normal senses cannot detect

Social Sensing

Few would question the remarkable extent to which the two pillars of modern day business communication — cell phones and email — enable us to effortlessly stay in touch with people on the other side of the planet. The paradox lurking behind this revolution is that these same technologies are steadily eroding the time and attention we devote to communicating with people in our immediate vicinity. The cost to the sender of sending an email or placing a cellular call is rapidly approaching zero. Unchecked, the cost to the recipient may rapidly become unmanageable, not in terms of financial cost, but rather in terms of demands on our time and attention. Witness the now common scene in airports and other public spaces — hundreds of people milling around in what appears to be animated conversation; on closer inspection, it turns out that they are not with each other, but rather with people connected to them via the near-invisible ear bud microphones they are wearing. Similarly, it is common to observe business colleagues in offices sitting just a few feet away from each other engaged in passionate debate. But the debate is often not verbal; rather, the only sound is the click-clack of keyboards as email flies back and forth. In desperation, some companies have resorted to the draconian measure if banning emails on certain days of the week ("email-free Fridays") or certain core hours of the day as the only way to pry their employees away from the email inboxes to engage in face-to-face dialog.

Why is it that so many people seem to find communication through email or cell phone more compelling than face-to-face dialog? Many value the fact that email permits asynchronous communication, enabling the recipient to respond only when it is convenient for them. Email also enables people to reinvent their personalities in ways that would be difficult or impossible for them socially. Cell phone technology has conquered geographical separation — anytime, anywhere communication. Rather than pursuing a chance encounter with the stranger standing next to me, it seems easier to talk to someone I know over a cell phone.

In contrast, technology has done little or nothing to enhance face-to-face communication. As we move from the current era of computing (so-called "personal" computers) to the next era (described variously as ambient intelligence or ubiquitous computing), help for social dialog will arrive in the form of next-generation personal information managers (PIMs) connected via wireless Personal Area Networks (PANs). PANs operate over a distance of just a few feet, connecting an individual to just those people within their immediate vicinity — their dinner companions, for example. First-generation PAN devices will be based on Bluetooth wireless technology.

Next generation PIMs arriving on the market over the next 24-36 months will store their owner's personal profile that will contain whatever information the owner may wish to share with others in their immediate vicinity. The information a user may wish to exchange in this way will obviously depend on the social context that the user is in at any given moment. In contrast to today's PIMs (where a lot of fumbling around will eventually result in a digital business card being exchanged between two devices), rich personal information will flow automatically and transparently between devices. It is quite likely that these PIMs will evolve to look nothing like today's devices. They may be incorporated into a pair of eyeglasses, or even in the clothes that we wear.

Widespread use of such devices will, of course, require that issues of personal privacy be resolved. However, peer-to-peer ad hoc networks of this type are inherently more respectful of individual privacy than client server systems. Users of PAN devices can specify either the exact names or the profiles of the people with whom they want their devices to communicate. They may also choose to have any information about themselves that is sent to another device time-expire after a few hours. This seems relatively benign compared to the information that can be collected about us (usually without our knowledge or consent) every time we browse the Web.

Many of us attend conferences every year for the purpose of professional networking. At any given conference of a hundred people or more, it is likely that there are a handful of potentially life-transforming encounters that could happen within the group. But such encounters are reliant on a chain of chance meetings that likely will not happen, due to the inefficiencies of the social network. Personal Area Network devices could dramatically improve our ability to identify the people in a crowd with whom we may wish to talk. Of course, we will want sophisticated software agents acting on our behalf to match our interests with the profiles of the people standing around us. We could even imagine a peer-to-peer Ebay in which my profile indicates that I am in the market to buy a certain type of car and I am alerted if anyone around me is trying to sell such a car. In Japan, it is already possible to buy a clear plastic key chain device that can be programmed to glow brightly when I encounter someone at a party whose interests are similar to mine. A high tech icebreaker!

The most profound technologies are the ones that “disappear” with use. Personal Area Network devices may enable nothing fundamentally new — they may just simplify what we already do

Environmental Sensing

We rely heavily on our natural senses (touch, sight, sound, smell) to keep us out of danger. Recent events are likely to have a lasting impact on the public’s awareness that there are an increasing number of hazards that our biological senses do not help us avoid. This desire for enhanced personal area environmental awareness is not simply a function of the anthrax scare. We will increasingly want to know more about the safety of air we breathe, the water that we drink, and the things we touch. This must be accomplished without bulky instrumentation and provide realtime feedback. I expect considerable commercial effort to be devoted towards transparent technology for personal environmental sensing. This may take the form of clothing that contains chemicals that change color in the presence of certain biohazards. Equally, we can expect a new generation of nano-sensors, custom-built to detect the presence of specific molecules, to be built into our clothing. Wearable technology presents great design challenges given the need to fold and wash the fabrics, maintain wearability, fashion, and light weight. For this reason, we should expect development in this arena to focus on chemical and nano-scale sensing. We have long expected our clothing to protect us from our surroundings — whether it be from the cold, UV radiation, or industrial hazards. Designing clothes that provide protection (through awareness) from other environmental hazards is a logical extension of the function of clothing to date.

THE CONSEQUENCES OF FULLY UNDERSTANDING THE BRAIN

Warren Robinett

We start with questions:

- How does memory work?
- How does learning work?
- How does recognition work?
- What is knowledge?
- What is language?
- How does emotion work?
- What is thought?

In short, how does the brain work?

We have nothing better than vague, approximate answers to any of these questions at the present time, but we have good reason to believe that they all have detailed, specific, scientific answers, and that we are capable of discovering and understanding them.

We want the questions answered in full detail — at the molecular level, at the protein level, at the cellular level, and at the whole-organism level. A complete answer must necessarily include an understanding of the developmental processes that build the brain and body. A complete answer amounts to a wiring diagram of the brain, with a detailed functional understanding of how the components work at every level, from whole brain down to ion channels in cell walls. *These are questions of cognitive science, but to get detailed, satisfying, hard answers, we need the tools of nanotechnology, biochemistry, and information technology.*

How important would it be if we did achieve full understanding of the brain? What could we do that we can't do now? How would it make our lives better? Unfortunately, scientific advances don't always improve the quality of life. Nevertheless, let's look at some possibilities opened up by a full understanding of how the brain works.

New Capabilities Enabled by Full Understanding of the Brain

We understand the input systems to the brain — the sensory systems — better than the rest of the brain at this time. Therefore, we start with ways of fooling the senses by means of electronic media, which can be done now, using our present understanding of the senses.

Virtual Presence

The telephone, a familiar tool for all of us, enables auditory-only virtual presence. In effect, your ears and mouth are projected to a distant location (where someone else's ears and mouth are), and you have a conversation *as if you were both in the same place*. Visual and haptic (touch) telepresence are harder to do, but nevertheless it will soon be possible to electronically project oneself to other physical locations and have the perceptions you would have if you were actually there — visually, haptically, and aurally, with near-perfect fidelity.

Tasks that could be accomplished with virtual presence include the following:

- meeting with one or more other people; this will be an alternative to business travel but will take the time of a telephone call rather than the time of a cross-country airplane flight
- interacting with physical objects in a distant location, perhaps a hazardous environment such as a nuclear power plant interior or battlefield, where actual human presence is impossible or undesirable
- interacting with objects in microscopic environments, such as in the interior of a human body (I have worked on a prototype system for doing this, the NanoManipulator; see http://www.WarrenRobinett.com/nano/)

Better Senses

Non-invasive, removable sensory enhancements (eyeglasses and contact lenses) are used now and are a useful first step. But why not go the second step and surgically correct the eyeball? Even better, replace the eyeball. As with artificial hips and artificial hearts, people are happy to get a new, better component; artificial sensory organs will follow. We can look at binoculars, night-vision goggles, and Geiger counters (all currently external to the body) to get an idea of what is possible: better resolution, better sensitivity, and the ability to see phenomena (such as radioactivity) that are normally imperceptible to humans. Electronic technology can be expected to provide artificial sensory organs that are small, lightweight, and self-powered. An understanding of the sensory systems and neural channels will enable, for example, hooking up the new high-resolution electronic eyeball to the optic nerve. By the time we have a full understanding of all human sensory systems, it is likely we will have a means of performing the necessary microsurgery to link electronic signals to nerves.

Better Memory

What is the storage mechanism for human memory? What is its architecture? What is the data structure for human memory? Where are the bits? What is the capacity of the human memory system in gigabytes (or petabytes)? Once we have answers to questions such as these, we can design additional memory units that are compatible with the architecture of human memory. A detailed understanding of how human memory works, where the bits are stored, and how it is wired will enable capacity to be increased, just as you now plug additional memory cards into your PC. For installation, a means of doing microsurgery is required, as discussed above. If your brain comes with 20 petabytes factory-installed, wouldn't 200 petabytes be better?

Another way of thinking about technologically-enhanced memory is to imagine that for your entire life you have worn a pair of eyeglasses with built-in, lightweight, high-resolution video cameras which have continuously transmitted to a tape library somewhere, so that every hour of everything you have ever seen (or heard) is recorded on one of the tapes. The one-hour tapes (10,000 or so for every year of your life) are arranged chronologically on shelves. So your fuzzy, vague memory of past events is enhanced with the ability to replay the tape for any hour and date you choose. Your native memory is augmented by the ability to reexperience a recorded past. Assuming nanotechnology-based memory densities in a few decades (1 bit per 300 nm^3), a lifetime (3×10^9 seconds) of video (10^9 bits/second) fits into 1 cubic centimeter. Thus, someday you may carry with you a lifetime of perfect, unfading memories.

Better Imagination

One purpose of imagination is to be able to predict what will happen or what might happen in certain situations in order to make decisions about what to do. But human imagination is very limited in the complexity it can handle. This inside-the-head ability to simulate the future has served us very well up to now, but we now have computer-based simulation tools that far outstrip the brain's ability to predict what can happen (at least in certain well-defined situations). Consider learning how to handle engine flameouts in a flight simulator: you can't do this with unaugmented human imagination. Consider being able to predict tomorrow's weather based on data from a continent-wide network of sensors and a weather simulation program. This is far beyond the amount of data and detail that human imagination can handle. Yet it is still the same kind of use of imagination with which we are familiar: predicting what might happen in certain circumstances. Thus, our native imagination may be augmented by the ability to experience a simulated future. At present, you can dissociate yourself from the flight simulator — you can get out. In future decades, with enormous computing power available in cubic micron-sized packages, we may find personal simulation capability built-in, along with memory enhancement, and improved sensory organs.

Now the Really Crazy Ones

Download Yourself into New Hardware

Imagine that the brain is fully understood, and therefore the mechanisms and data structures for knowledge, personality, character traits, habits, and so on are known. Imagine further that, for an individual, the data describing that person's knowledge, personality, and so forth, could be extracted from his brain. In that case, his mind could be "run" on different hardware, just as old video games are today run in emulation on faster processors. This, of course, raises lots of questions. What is it that makes you *you*? (Is it more than your knowledge and personality?) Is having the traditional body necessary to being human? Nevertheless, if you accept the above premises, *it could be done*. Having made the leap to new hardware for yourself, many staggering options open up:

- No death. You back yourself up. You get new hardware as needed.
- Turn up the clock speed. Goodbye, millisecond-speed neurons; hello, nanosecond-speed electronics.
- Choose space-friendly hardware. Goodbye, Earth; hello, galaxy.

Instant Learning

If the structure of knowledge were fully understood, and if we controlled the "hardware and software environment" of the mind, then presumably we would understand how new knowledge gets integrated with old knowledge. The quaint old-fashioned techniques of "books" and "school" would be reenacted sometimes for fun, but the efficient way would be to just get the knowledge file and run the integrate procedure. Get a Ph.D. in Mathematics with "one click."

Hive Mind

If we can easily exchange large chunks of knowledge and are connected by high-bandwidth communication paths, the function and purpose served by individuals becomes unclear. Individuals have served to keep the gene pool stirred up and healthy via sexual reproduction, but this data-handling process would no longer necessarily be linked to individuals. With knowledge no longer encapsulated in individuals, the distinction between individuals and the entirety of humanity would blur. Think Vulcan mind-meld. We would perhaps become more of a hive mind — an enormous, single, intelligent entity.

Speed-of-Light Travel

If a mind is data that runs on a processor (and its sensors and actuators), then that data — that mind — can travel at the speed of light as bits in a communication path. Thus, Mars is less than an hour away at light speed. (We needed a rocket to get the first receiver there.) You could go there, have experiences (in a body you reserved), and then bring the experience-data back with you on return.

Self-Directed Evolution

If mind is program and data, and we control the hardware and the software, then we can make changes as we see fit. What will human-like intelligence evolve into if it is freed from the limits of the human meat-machine, and humans can change and improve their own hardware? It's hard to say. The changes would perhaps be goal-directed, but what goals would be chosen for self-directed evolution? What does a human become when freed from pain, hunger, lust, and pride? (If we knew the answer to this, we might be able to guess why we haven't detected any sign of other intelligences in the 100 billion stars of our galaxy!)

User-Interface Olympics: Using Competition to Drive Innovation

Warren Robinett

Has bicycle racing improved bicycles? Yes, it has. We humans like to win, and like Lance Armstrong pedaling through the Alps in the Tour de France, we demand the best tools that can be made. The competition, the prestige of being the world champion, the passion to win, publicity for the chosen tools of the winners — these forces squeeze the imaginations of bicycle engineers and the bank accounts of bicycle manufacturers to produce a stream of innovations: lighter and higher-strength materials, more efficient gearing, easier and more reliable gear-shifting, aerodynamic improvements such as farings and encased wheels... the list goes on and on.

Competition spawns rapid improvements. Sounds a bit like evolution, doesn't it? *Lack of competition* can lead to long periods of quiescence, where nothing much changes. (Did you know the QWERTY keyboard was designed 100 years ago?)

This principle that *competition spawns improvement* could be applied to drive innovations in user-interface design. We call the proposed competition the *User-Interface Olympics*. Here is a sketch of how it might work:

- It would be an annual competition sponsored by a prestigious organization — let's say, the U.S. National Science Foundation.
- The winners would get prestige and possibly prize money (like the Nobel Prize, Pulitzer Prize, Emmies, Academy Awards, Oscars, and so on).
- The competition would be composed of a certain number of events, analogous to Olympic events. Individual contestants, or teams of contestants, compete for the championship in each event. User-interface events would be such things as
  - a timed competition to enter English text into a computer as fast as possible. (Surely someone can do better than the QWERTY keyboard!)
  - a timed competition to select a specified series of items from lists. (Can we improve on the 40-year-old mouse?)
- Contestants would provide their own tools. This is analogous to the equipment used by athletes (special shoes, javelin, ice skates). However, for the User-Interface Olympics, the tools are the hardware and software used by each competitor.
- Since the goal is to stimulate innovation, contestants would have to fully disclose the working of their tools. A great new idea would get you one gold

medal, not ten in a row. This is similar to the patent system, in which rewards during a limited period are bartered for disclosure and dissemination of ideas.

- An administrative authority would be needed, analogous in the Olympic Committee and its subordinate committees, to precisely define the rules for each event, for qualifying for events, and many other related matters. This Rules Committee would monitor the various events and make adjustments in the rules as needed.
- We would expect the rules of each event to co-evolve with the competitors and their tools. For example, the rule against goal tending in basketball was instituted in response to evolving player capabilities; in the 100-meter dash, precise rules for false starts must be continually monitored for effectiveness. Winning within the existing rules is not cheating, but some strategies that players may discover might not be really fair or might circumvent the intent of the competition. Of course, some competitors do cheat, and the rules must set reasonable penalties for each type of infraction. The Rules Committee would therefore have to evolve the rules of each event to keep the competition healthy.
- New events would be added from time to time.

These contests would be similar to multiplayer video games. The contestants would manipulate user-input devices such as the mouse, keyboard, joystick, and other input devices that might be invented. The usual classes of display devices (visual, aural, and haptic) would be available to the contestants, with innovations encouraged in this area, too. Most malleable, and therefore probably most fertile for spawning innovations, would be the software that defined the interaction techniques through which the contestant performed actions during the contest.

If we set things up right, perhaps we could tap some of the enormous energy that the youth of the nation currently pours into playing video games.

The rules for each contest, which would be published in advance, would be enforced by a computer program. Ideally, this referee program could handle all situations that come up in a contest; whether this actually worked, or whether a human referee would be needed, would have to determined in real contests. Making the referee completely automated would offer several advantages. Contests could be staged without hiring anyone. Computer referees would be, and would be perceived to be, unbiased. Early qualifying rounds could be held using the Internet, thus encouraging many contestants to participate. Figure B.14 shows a system diagram.

If this idea is to be attempted, it is critical to start with a well-chosen set of events. (Imagine that the Olympics had tried to start with synchronized swimming and sheep shearing!) A small, well-justified set of events might be best initially, just to keep it simple and try out the idea. One way to identify potential events for the UI Olympics is to look at input devices that currently are widely used:

- computer keyboard — suggests a text-entry event
- computer mouse — suggests an event based on selecting among alternatives
- joystick, car steering wheel — suggest one or more events about navigating through a 2-D or 3-D space

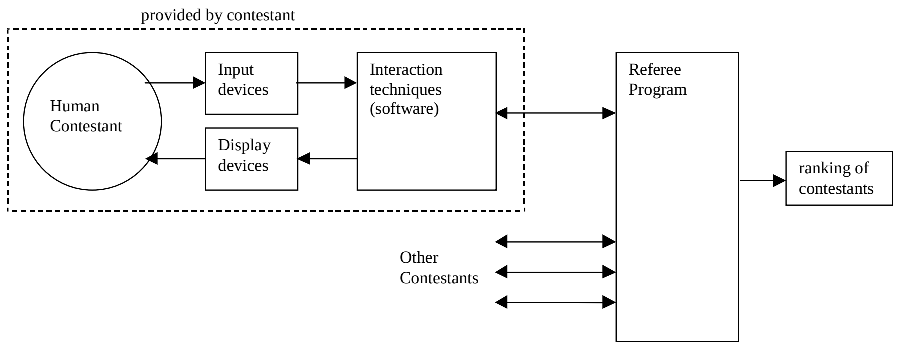

Figure B.14. System Diagram for a contest in the User-Interface Olympics, mediated by an automated referee program, with several contestants participating. The contestants provide their own hardware and software.

The real Olympics has events based both on raw power, speed, and stamina (weight lifting, races, and the marathon) and also events based on more complex skills (skiing, badminton, baseball). Similarly, the User-Interface Olympics could complement its events based on low-level skills (text entry, navigation) with some events requiring higher-level thinking. There are many kinds of "high-level thinking," of course. One class of well-developed intellectual contests is the mathematical competition. There are a number of well-known competitions or tests we can consider as examples: the MathCounts competitions run among middle schools and high schools; the Putnam Mathematical Competition run for undergraduates, and the math portion of the Scholastic Aptitude Test (or SAT, the college entrance test). Another similar competition is the annual student programming contest sponsored by the Association for Computing Machinery. One or more events based on solving well-defined categories of complex problems, using tools chosen by the contestant, would be desirable.

Strategy board games, such as chess and go, are another class of contests requiring complex skills. The rules for these games have already evolved to support interesting, healthy competitions and cultures. To focus on chess for a moment, by making chess an event in the User-Interface Olympics, we have an opportunity to reframe the false dichotomy between a human chess player and a chess-playing computer — we introduce a third possibility, a human contestant combined with her chess-analysis software. I personally believe that the combination of a good chess player, a good chess program, and a good user interface to integrate the two could probably beat both Deep Blue and Garry Kasparov. At any rate, this is a well-defined and testable hypothesis.

Therefore, the following events are proposed for the initial User-Interface Olympics:

- Text-entry speed competition
- Selection-among-alternatives race
- Navigation challenge: a race through a series of waypoints along a complex racecourse

- Timed math problems from the SAT (or equivalent problems)
- Timed chess matches

Each of these events would need precisely-formulated rules.

The strategy needed to achieve this vision of a thriving, well-known, self-perpetuating User-Interface Olympics that effectively drives innovation in user interface hardware and software is this:

- Fund the prizes for the first few years — let's say $100,000 for each of the four events
- Set up a governing committee and carefully choose its chairman and members. Give the committee itself an appropriate level of funding.
- Set an approximate date for the first User-Interface Olympics.

If the User-Interface Olympics were to become successful (meaning it had the participation of many contestants and user interface designers, it spawned good new ideas in user interface design, it had become prestigious, and it had become financially self-supporting), the benefits which could be expected might include the following:

- rapid innovation in user-interface hardware and software
- recognition for inventors and engineers — on a par with scientists (Nobel Prize), writers (Pulitzer Prize), and actors (Academy Award)
- improved performance on the tasks chosen as events

Sometimes prizes can have an inordinately large effect in relation to the amount of money put up. Witness the prize for the first computer to beat the (human) world chess champion (Hsu 1998; Loviglio 1997). Witness the prize for the first human-powered flying machine (Brown et al. 2001). A million dollars or so in prize money to jump-start the User-Interface Olympics might be one of the best investments ever made.

References

Brown, D.E., L.C. Thurow, and J. Burke. 2001. *Inventing modern America: From the microwave to the mouse.* MIT Press. Also at http://web.mit.edu/invent/www/ima/maccready_bio.html.

Hsu, F.H. 1998. Computer chess: The Deep Blue saga. http://www.stanford.edu/class/ee380/9798sum/lect03.html.

Loviglio, J. 1997. "Deep Blue Team Awarded $100,000 Fredkin Prize." New York Times CyberTimes, July 30. http://www.rci.rutgers.edu/~cfs/472_html/Intro/NYT_Intro/ChessMatch/DeepBlueTeamAwarded.html.

ACCELERATING CONVERGENCE OF NANOTECHNOLOGY, BIOTECHNOLOGY, AND INFORMATION TECHNOLOGY

Larry Todd Wilson, IEEE

My goal is to focus on a single NBIC-oriented idea that, if actualized, would unleash massive capabilities for improving all human performance. This single thing would have extreme interrelated, multiplicative effects. It's a bit like an explosion that starts consequential, far-reaching chain reactions. Furthermore, the one thing should accelerate and strengthen all other biotech ideas and fulfill a self-referential quality for advancing itself. It is difficult to negate the notion that some ideas, actions, or objects are more important than others. This perspective is characterized by statements like, “This is what should come first because if we had that ability or understanding, then we could (achieve these results)… and if we had those results, then we could actualize…”

The “One Thing”is, *Nullify the constraints associated with a human’s inherent ability to assimilate information.*

Why should this receive favorable positioning? Advances in thinking performance are more important than advances in artifacts. This is due to the fact that the advances in artifacts are always a function of the human thinking system. The dynamics of innovation must be managed by human consciousness before it is “externally” managed at all. There are many naturally occurring phenomena that are not apparent to the senses or the imagination. However, a technology does not become a technology until it enters the realm of human consciousness.

Examples below deliver “as-is” versus “could be” explanations of the importance of enhancing how we assimilate information. From the examples, it is not difficult to imagine the transformations that may result due to the ripple effects. Overall, the focus on ways to enhance how humans assimilate information will result in significant increases in a human’s ability to approach a complex need, achieve comprehension, and accomplish an intended result. Increased ability equates to gaining faster comprehension, better comprehension, comprehension in a situation that previously was unfathomable, faster solutions, and better solutions, and to finding solutions to problems that seemed unsolvable.

Assimilating information is a kind of human intellectual performance. There are three and only three types of human performance that could be the focus of improvement:

- intellectual performance (such as thinking, deciding, learning, and remembering)
- physical performance (such as moving, reaching, and lifting)
- emotional performance (feeling)

All human experiences are variations of one or more of these three.

Candidates of the “best thing” could be evaluated according to either criteria or questions like these:

- Is this idea/action/object fundamental to all dimensions and expressions of human performance (thinking, feeling, and moving)?
- Does this thing have a multiplicative nature in regards to all other biotech ideas, actions, and objects? Does this one thing produce fission-oriented and fusion-oriented results? Does its presence cause a reaction that in turn creates energy associated with pragmatic NBIC inventions and discoveries?
- *A priori*, does it have validity on its face? Does a listener agree that this one thing will indeed impact everything else?
- *A posteriori*, does it have perceptible, significant advances in several other areas? Did this one thing deliver a high return on investment? How do we

know? What is measured? Does its presence actually increase the rate of all biotech inventions and discoveries?

Table B.1

AS IS	COULD BE
The span of judgment and the span of immediate memory impose severe limitations on the amount of information that we are able to receive, assimilate, and remember. In the mid-1950s, this was labeled as “seven, plus or minus two.”	The innate limitations of human short-term memory are irrelevant due to the synergistic reliance upon “external” working memory, which is embedded in everything around us.
Short-term memory is working memory that works to retain sensory information presented by the mechanism of attention. No human being can hold many concepts in his head at one time. If he is dealing with more than a few, he must have some way to store and order these in an external medium, preferably a medium that can provide him with spatial patterns to associate the ordering, e.g., an ordered list of possible courses of action.	Increase the size and capability of working memory. Deliberate consideration of the items in external working memory can be called to mind upon demand. Manage how linguistic coding influences thought processes. Quantitatively measure stimulus (primarily in the form of linguistic-based prompts) and response (reactions in the form of decisions or feelings or movements).
Material is lost from short-term memory in two ways; it will not be committed to long-term memory if interference takes place or time decay occurs. One of the by-products related to the limitations of short-term memory is that there is great relief when information no longer needs to be retained. Short term memory is like a series of input and output buffers in which intermediate data can be stored during any thinking activity; this memory has very limited capacity and can be easily overloaded. In order to alleviate the anguish of overload, there is a powerful desire to complete a task, reduce the memory load, and gain relief. This event is referred to as “closure,” which is the completion of a task leading to relief.	Minimize the losses that naturally occur. Consciously add or delete items in working memory. Regulate the need for closure because the human is confident that it’s “still there” (although I don’t remember exactly what *it* is). Increase the number and rate of working memory instances. Engineer a seamless human mind/external memory interface, and thereby make human and machine intelligence coextensive. Basic analysis and evaluation of working memory contents are achieved in partnership or alone.

AS IS	COULD BE
Bounded rationality refers to the limitations inherent in an individual's thought processes when there are more than a few alternatives being considered at the same time. Bounded rationality occurs because an individual has limited, imperfect knowledge and will seek satisfaction rather than strive for optimal decisions.	Effectively unbound "bounded rationality." The number and interrelationships of evaluations are dramatically expanded.
Individual thinking repertoires are limited (in their usefulness) and limiting (in their applicability).	Codify the elemental and compound thinking processes. Use the external working memory to manage the objects of the attention with novel ways of orchestrating the human's awareness of them. Increase the frequency, quantity (novel combinations), and throughput of these compounds. Gather more and more intelligence about the signals — the contextual nuances associated with variations of the compounds. Examples of compounds are Abstract Accept Accommodate Adopt Advise Agree Align Apply Appraise Approve Arrange Assign Assimilate Assume Authenticate Authorize Calculate Catalogue Categorize Change Check Choose Classify Close Compare Compile Compute Conclude Conduct Confirm Consider Consolidate Construct Contrast Contribute Coordinate Create Decide Decrease Deduce Define Delete Deliberate Deliver Deploy Derive Describe Determine Develop Differentiate Direct Disagree Disapprove Discern Distinguish Elaborate Eliminate Emphasize Enable Enhance Enrich Establish Estimate Examine Exclude Execute Expand Explore Extrapolate Facilitate Find Focus Formulate Generalize Group Guess Guide Hypothesize Imagine Include Incorporate Increase Index Induce Infer

AS IS	COULD BE
	Inform Initiate Insert Inspect Interpret Interview Invent Judge Locate Match Measure Memorize Merge Modify Monitor Observe Optimize Organize Originate Outline Pace Predict Prepare Presume Prevent Prioritize Probe Promote Provide Question Rank Rate Reason Receive Recognize Recommend Refine Reflect Regulate Reject Remove Report Resolve Respond Scan Schedule Scrutinize Search Seek Serve Settle Show Solicit Solve Sort Speculate Submit Support Suppose Survey Synthesize Translate Validate Verify Visualize.
Specialists often miss the point. The point is to swap advances among different disciplines. It's all about permutations and combinations. Discoveries from biology and chemistry are hooked up with synthesis and fabrication tools from engineering and physics. Each discipline has its own sets of problems, methods, social networks, and research practices. There are no effective ways in which the intellectual results of subdisciplines can be managed and thereby accelerate consilience and cross-disciplined performance breakthroughs.	Progress towards a new sense of the complex system. The most obvious change will be the benefits of working with many kinds of associations/relations. More people will be able to perceive loops and knots. Sense the complex system with a set of universal constructs for systematically managing the interrelationships among disciplines. Accurate visualization of many kinds of relations (not just parent-child relations) will shift the reliance of the satisficing mode of hierarchical interpretations to the closer-to-reality heterarchical structure. Continue to splinter the subdisciplines and achieve convergence when needed for important insights.
Today, many physicists spend time translating math into English. They hunt for metaphors that can serve as a basis for enhancing comprehension of relatively imperceptible physical phenomena.	Integrate mathematics, verbal, and visual languages in order to allow individuals to traverse the explanation space. Aid the acceleration of new ways for more people to abandon their intuitive (perhaps innate) mode of sensory perception associated with the macro world. Achieve integration (and concise translation) between our symbol sets

AS IS	COULD BE
	(math, verbal, and visual) and open up the chance to address more, apparently paradoxical, phenomena. The assumption is that many of these paradoxes are just illusions created when you look at an n-dimensional problem through a three-dimensional window.
Linguistic-based messages, which plod along the user's tolerance for listening, govern the rate of assimilation.	Establish the path more directly because all forms of intelligence, whether of sound or sight, have been reduced to the form of varying currents in an electric circuit.
Imaging modalities don't offer a concise way of observing the dynamics of how we assimilate information. PETs are more accurate in space, and EEGs are more accurate in time. EEGs can capture events on the scale of milliseconds, but they're only accurate to within centimeters. Scans are like slow motion — a thousand times slower — but they're accurate to the millionth of an inch.	Extend the visual languages to the actual visualization of localized neuronal activity. Understand the spatial-temporal nature of assimilation with a realtime movie stage where we watch thoughts as they gather and flow through the brain. Understand how the human perception of mind arises from the brain. Formalize in neural network models operating on traditional hardware. Thus, intelligences akin to humans will reside in the Internet. These intelligences, not being physically limited, will merge and transform themselves in novel ways. The notion of discrete intelligence will disappear.