Robotics and Intelligent Systems in Support of Society

Raj Reddy, *Carnegie Mellon University*

Over the last 50 years, there has been extensive research into robotics and intelligent systems. Although much of the research has targeted specific technical problems, advances in these areas have led to systems and solutions that will profoundly impact society. Underlying most of the advances is the unprecedented exponential improvement of information technology. Computer scientists expect the exponential growth of memory and bandwidth to continue for the next 10 to 20 years, leading to terabyte disks and terabytes-per-second bandwidth at a cost of pennies per day (see the "Technology Trends" sidebar).

Several applications of robotics and intelligent systems could profoundly impact the well being of our society, transforming how we live, learn, and work.

The question is, what will we do with all this power? How will it affect the way we live and work? Many things will hardly change—our social systems, the food we eat, the clothes we wear, our mating rituals, and so forth. Others, such as how we learn, work, and interact with others, and the quality and delivery of healthcare, will change profoundly. Here I present several examples of using intelligent technologies in the service of humanity. In particular, I briefly discuss the areas of robotics, speech recognition, computer vision, human-computer interaction, natural language processing, and artificial intelligence. I also discuss current and potential applications of these technologies that will benefit humanity—particularly the elderly, poor, sick, and illiterate.

Robotics for elder care and search and rescue

As the life expectancy of the world's population increases, soon over 10 percent of the population will be over age 70.[1] This age group will likely have minor disabilities impacting its quality of life, which we can group into three broad categories: sensory, cognitive, and motor. Fortunately, robotic and intelligent systems can remedy these disabilities.

Robots are also paving the way for many rescue missions, traveling through smoldering ruins to places rescuers can't reach. The 9/11 tragedy and recent hurricanes have showcased the nascent robotics technology and the industry attempting to provide tools and systems for data gathering and rescue.

Helping the aging population

A dramatic increase in the elderly population along with the explosion of nursing-home costs poses extreme challenges to society. Current care for the elderly is insufficient, and in the future, there will be fewer young people to help older adults cope with the challenges of aging.[1]

Robots for elder care must satisfy several requirements. Current robots, with their fixed velocities, can frustrate users. We need robots that can keep pace with the subject, moving neither too fast nor too slow. Safety while navigating in the presence of an elderly person is also important. Given the limitations of current vision systems, an eldercare robot might not always detect obstacles beyond its field of view and could accidentally hit a person. Also, the robot must be able to understand and respond to voice commands. Current speech recognition and synthesis technologies are sufficient to make this possible, but several problems exist, such as not being able to comprehend a continuous open-domain speech consisting of confusing words, having trouble following who is speaking when multiple people are present for the conversation, and not being able to block out environmental noise.

The Pearl Robot, originally developed at Carnegie Mellon University (CMU) and currently being used for research by Martha Pollack at the University of

1541-1672/06/$20.00 © 2006 IEEE
Published by the IEEE Computer Society

Technology Trends

Back in 1994, I created a chart predicting exponential growth trends in computer performance (see figure A). In 2000, as expected, we saw the arrival of a giga-PC capable of a billion operations per second, a billion bits of memory, and a billion bits per second bandwidth, all available for less than US$2,000. Barring the creation of a cartel or some unforeseen technological barrier, we should see a tera-PC by 2015 and a peta-PC by 2030.

Advances in magnetic disk memory have been even more dramatic. Disk densities have been doubling every 12 months, leading to a thousand-fold improvement every 10 years. Today, you can buy a 100-gigabyte disk memory for about $100. By 2010, we should be able to buy 100 terabytes for about the same price. At that cost, each of us will be able to have a personal library of several million books and a lifetime collection of music and movies—all on our home PC.

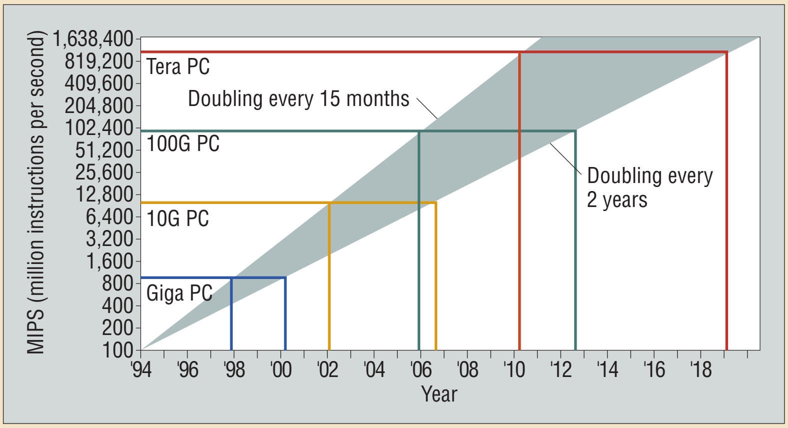

Figure A. Exponential growth trends in computer performance. (figure courtesy of Carnegie Mellon University and Raj Reddy)

Most dramatic of all recent technological advances is the doubling of bandwidth every nine months, propelled by the advances in fiber-optic technology. Today you can buy commercial systems that will transmit 1.6 terabits per second on a single fiber using dense wavelength division multiplexing technology. DWDM technology uses 160 different wavelengths, each capable of transmitting 10 gigabits per second. Experimental systems can currently transmit as much as 25 terabits per second on a single fiber.

What can you do with 1.6 terabits per second bandwidth? It would take about 50 seconds to transmit all the books in the Library of Congress. All the phone calls in the world could be carried on single fiber with room to spare. The main bottleneck today isn't the bandwidth but rather the computer speed capable of accepting and switching data packets arriving at terabit data rates. The speed of light is also proving to be a problem. The maximum sustainable bandwidth using TCP/IP is governed by the round-trip delay times. At terabit rates, with round-trip times of about 30 ms across the US, 30 billion bits would be transmitted before an acknowledgment was received.

Michigan,[2] has been demonstrated in several assistive-care situations (see figure 1a). Pearl provides a research platform to test a range of ideas for assisting the elderly. Two Intel Pentium 4 processor-based PCs run software to endow her with wit and the ability to navigate, and a differential drive system propels her. A Wi-Fi network connection helps her communicate as she rolls along, while laser range finders, stereo camera systems, and sonar sensors guide her around obstructions. Microphones help her recognize words, and speakers enable others to hear her synthesized speech. An actuated head unit swivels in lifelike animation.

Currently, Pearl can trek across carpets and kitchen floors, and she includes a handheld device that prompts people to do certain (preprogrammed) things. She also acts as a walker, guiding people through pathways. Researchers hope that such autonomous mobile robots will have endless possibilities and one day live in the homes of chronically ill elderly people to perform a variety of tasks,[3] such as

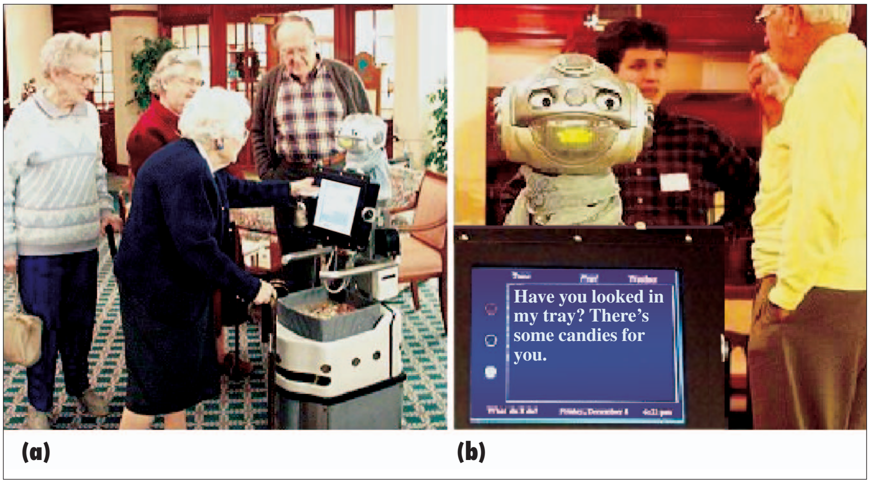

Figure 1. The Pearl Robot (a) being used in an eldercare environment and (b) providing text messages on a screen as a reminder. (figure courtesy of the Carnegie Mellon University Nursebot Project and Martha Pollack)

- monitoring people and reminding them to visit the bathroom, take medicine, drink, or see the doctor (see figure 1b).
- connecting patients with caregivers through the Internet. The robot is a platform for telepresence technology whereby professional caregivers can interact directly with remote patients, reducing the frequency of doctor visits.
- collecting data and monitoring patient well-

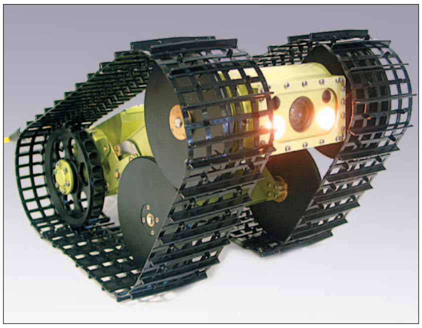

Figure 2. The VGTV Xtreme rescue robot. (figure courtesy of American Standard Robotics)

being. Emergency conditions, such as heart failure or high blood-sugar levels, can be avoided with systematic data collection.
- operating appliances around the home such as the refrigerator, washing machine, or microwave.
- taking over certain social functions. Many elderly people are forced to live alone, deprived of social contacts. The robot might help people feel less isolated.

Although Pearl can't yet perform such tasks, she's a one of a kind costing close to US$100,000. To ready a mass-market version of the automaton that would be useful for the elderly, researchers must address the list of desirables just presented as well as overcome a rash of software and hardware issues, such as battery power and Pearl's inability to navigate stairs.

Furthermore, to act as a telepresence system, Pearl should be able to offer the caregiver detailed status reports about the elder's daily activities, such as whether the patient took the correct medicine and ate his or her meals. The system should be able to monitor people and react to unusual behavior patterns—for example, alerting the caregiver if a person hasn't moved from a chair for a long time.

Also, Pearl currently uses simple text strings for patient reminders, but as an elder's cognitive capacities diminish or during periods of illness, the person might require more detailed reminders. The system should be able to monitor the elder's needs and distribute intelligent reminders based on parameters such as time of day, timing of the previous interactions, the user's mood, and required user actions. Finally, future robots should be capable of picking up or moving things for the user.

As a society, we must find alternative ways of providing care to the elderly and chronically ill, providing them with the dignity they deserve. Technological advances will no doubt make such eldercare robots available commercially in the near future.

Robotics for rescue

Natural and manmade disasters lead to unique requirements for effective collaboration of human and robotic systems. Often disaster locations are too dangerous for human exploration or are unreachable. Additional constraints such as the availability of rescuers, extreme temperatures, and hurricane-force winds result in significant delays before human rescuers can start searching for victims. In most cases, rescuers need to retrieve victims within 48 hours to enhance survival rates.

Disasters in the last decade showed that advances such as earthquake prediction are not enough. For example, the 1995 Kobe earthquake destroyed several large structures—including buildings and highways believed to be earthquake proof. In both Kobe and in the Oklahoma bombing case, human rescue efforts alone weren't enough, resulting in unnecessary loss of life. The huge robots available at the time couldn't maneuver in the rubble or be used effectively.

Lessons learned from these and other rescue situations motivated further research around the world, leading to the creation of small, light, and cheap rescue robots.[4] Today several organizations are actively participating in designing small rescue robots that can carry a human-sized payload. Some systems can maneuver over land and air and detect sounds. Many of the robots have cameras with 360-degree rotation that can send back high-resolution images. Some rescue robots are equipped with specialized sensors that can detect body heat or colored clothing.

Figure 2 shows the VGTV Xtreme, American Standard Robotics' commercially available 14-pound rescue robot (see www.asrobotics.com/products/vgtv_xtreme.html). Unlike the previous generation's laboratory systems, this rescue robot's system is durable and flexible and can stand up to a beating in the field. Current models of the VGTV Xtreme robot have more ground clearance, so they can more easily go over rubble piles. Cameras have longer zoom lenses, so the robots don't have to journey as far into a structure to provide searchers with high-resolution images. Operator controls have been improved and are compatible with the rescuers' equipment and gloves. Also, the robots are waterproof, so decontamination is easier.

Rescue robots are highly invaluable in disaster relief efforts, and this particular area of robotics has received much attention in the last few decades. Several rescue robots exist today with varying capabilities suited to specific domains. The VGTV Xtreme robots can travel hundreds of feet autonomously. Other rescue robots, such as CMU's snake robot, help assess contamination at waste storage sites or search for land mines (see the "Robotics for Land Mine Detection" sidebar). The military has used similarly small and durable tactical mobile robots called PackBots to explore Afghan caves. Also, since 2000, the American Association for Artificial Intelligence has conducted an annual rescue robot competition to identify the capabilities and limitations of existing rescue robots (see http://palantir.swarthmore.edu/aaai06).

Continued research and development is needed to enhance rescue robotic systems and eliminate their limitations. It's difficult to build robust rescue robots that can deal with a disaster site's unexpected complexity. For example, the robots must be adaptable to water-clogged dark mines, collapsed building rubble, water-clogged pipes, and so forth. Also, the robots should be able to adjust to lighting and be resistant to extreme temperatures such as in a volcano or on an icy mountain.

A limitation observed in robot rescue efforts at the World Trade Center was that the robots couldn't go very far into the rubble because of short tether. At present, teleoperation from remote sites is often limited to about 100 feet from the disaster site. Next-generation robots must address issues such as sensor, power, and commuting limitations and long set-up times.

Furthermore, large-scale, affordable, and efficient deployment of these robots demands increased participation and awareness from relevant government agencies. We not only must improve the technology and research infrastructure but also train a work force of first responders who can effectively use such systems.

Speech recognition: The reading tutor

Over a billion people in the world can't read or write (see www.literacyonline.org/explorer/overview.html), and it's likely that as many as two billion are functionally illiterate in that they can't understand the meaning of the sen-

Robotics for Land-Mine Detection

Around 60 to 100 million land mines infest over 60 countries following a century of conflicts. The presence of these mines has devastating consequences for the civilian populations, resulting in millions of acres of land that can't be reclaimed. Thousands of innocent people are maimed or killed every year in land-mine accidents.

Currently, most land-mine detection is done manually, using metallic and other types of sensors (see figure B1). Often, the sensor indicates the presence of metal when there is none, leading to false alarms. Even worse, with the proliferation of mines with little to no metallic content, the use of the classical metal detector is limited. For every 5,000 mines removed, one deminer is killed or maimed. More sophisticated techniques are needed, and robotics technologies offer a unique opportunity to create a safe solution.

Autonomous robots equipped with appropriate sensors can systematically search a given area and mark all potential land mine areas, leading to significantly higher detection rates and fewer casualties (see figure B2). Newer sensors such as quadruple resonance and explosive odor detectors called "sniffers" should be available in a few years. Quadruple resonance, a radio-frequency imaging technique similar to an MRI but without the large external magnet, is a promising technological alternatives. Aerial electronic sniffers can rapidly detect the absence or presence of mines.

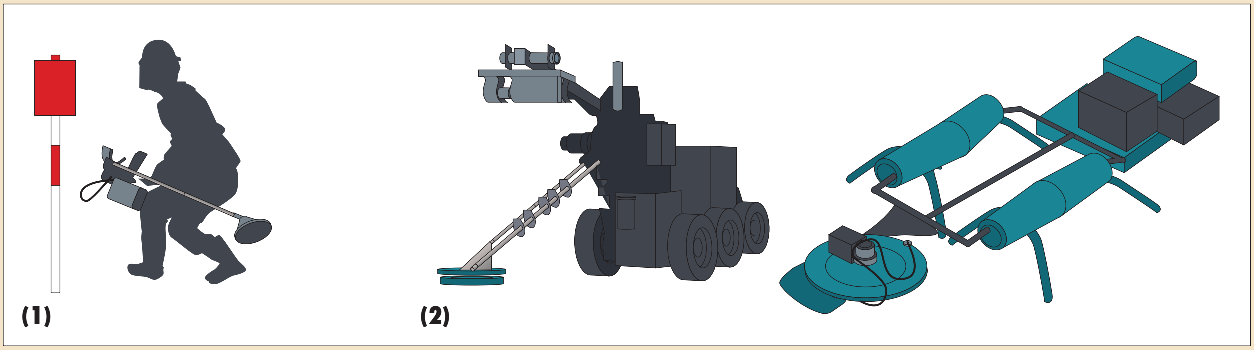

Figure B. (1) General land-mine detection and (2) autonomous robots for land-mine detection.

tences they read. Advances in speech-recognition and synthesis technologies provide an opportunity to create a computer-based solution to illiteracy. The solution involves an automated reading tutor that displays stories on a computer screen and listens to children read aloud using speech-recognition technology.

Until 10 years ago, speech recognition systems weren't fast enough to recognize the connected sentences spoken in real time. A reading tutor application that can discover and correct mispronunciations requires not merely speech-recognition capabilities; it must also be able to detect deviations in stress, duration, and spectral properties from that of native speakers.

Figure 3 shows the Reading Tutor, developed by CMU's Jack Mostow.[5] The Reading Tutor lets the child choose from a menu of high-interest stories from *Weekly Reader* and other sources, including user-authored stories. It adapts CMU's Sphinx speech-recognition system to analyze the student's oral reading. The Reading Tutor intervenes when the reader makes mistakes, gets stuck, clicks for help, or is likely to encounter difficulty. It responds with assistance modeled after expert reading teachers but adapted to the technology's capabilities and limitations. The current version runs on Windows XP on a PC with at least 128 Mbytes of memory. Although it's not yet a commercial product, hundreds of children have used the Reading Tutor daily as part of study to test its effectiveness.[6]

Figure 3. A student using the Reading Tutor in Ghana. (figure courtesy of CMU and Jack Mostow)

To fully realize this technology's poten-

Figure 4. A Navlab vehicle developed at Carnegie Mellon University can autonomously operate at high speeds for an extended period of time. (figure courtesy of CMU and Red Whittaker)

tial, we must reduce speech-recognition errors and develop systems that can understand nonnative speakers, local dialects, and the speech of children. We must also track user engagement using response times and enable the system to propose a mentoring approach for a specific student by modeling that student's behavior.

Computer vision and intelligent cruise control

Over a million people die annually in road traffic fatalities—40,000 in the US alone.[7] The annual repair bill for the cars involved in these accidents in the US exceeds over $55 billion. Sensing, planning, and control technologies made possible by advances in computer vision and intelligent systems could reduce deaths and repair costs by a significant percent.

Forty percent of vehicle crashes can be attributed in some form to reduced visibility due to lighting and weather conditions.[8] Physical sensors could alert a driver when glare, fog, or artificial light exists in the environment. Furthermore, over 70 percent of accidents are caused by human driving errors—usually speeding, driver fatigue, or drunk driving.[9] An autopilot that could temporarily control the vehicle in such situations and navigate it to safety could dramatically reduce casualties. However, this would require dependable sensitivity and perception—a key component in building next-generation cruise-control systems.

To accomplish this, we need sensors that use light, sound, or radio waves to detect and sense the physical environment. These sensors could gather data such as the speed, distance, shape, and color of objects surrounding the vehicle. We then need efficient object-classification techniques to extract the shape of objects from this sensor data and accurately classify them based on color. We'd also require real-time object-tracking systems that can continuously monitor (and predict) the trajectories of these vehicles and drivers over time. Based on this information, researchers should be able to develop scene-recognition algorithms that can recognize interactions between objects in the environment and extrapolate them to a possible collision scenario.

In addition to perception modules, we also need control systems and actuators to steer the vehicle. We need feedback control systems with proportional control to help the vehicle maintain a constant speed by automatically adjusting the throttle based on the current speed. We need efficient path-planning and localization techniques so the vehicle can autonomously navigate a set path.

Once these technologies are in place, we'll be able to build collision-warning systems with intuitive interfaces that warn drivers of imminent dangers well in advance. We can design collision-avoidance and autonomous navigation systems that can navigate the vehicle around obstacles and dangerous scenarios without help from the driver.

Such systems could also help reduce traffic. The Texas Transportation Institute estimates that in 2000 alone, major US cities experienced 3.6 billion hours of total delay and 5.7 billion liters of wasted fuel.[10] This resulted in an aggregate of $67.5 billion in lost productivity.[10] Excessive braking due to driver panic or minor miscalculations caused many of these jams. Such irregular driving not only disrupts traffic flow but also causes discomfort for passengers. Also, when drivers maintain a large distance between their vehicle and the one in front of them, they significantly underutilize the road. Intelligent cruise-control systems that regulate vehicle speed based on time-to-collision could help us avoid many of these problems. A majority of these traffic jams could be prevented if a mere 20 percent of vehicles on the road used adaptive cruise control.[11]

The Vision and Autonomous Systems Center at the CMU is building collision-warning systems for transit buses that use short-range sensors to detect and warn the driver of frontal and side collisions.

Furthermore, CMU's Field Robotics Center has developed Navlab vehicles that can autonomously operate at highway speeds for extended periods of time (see figure 4).[12] These vehicles are equipped with perception modules to sense and recognize the surrounding environment.

A Navlab vehicle can detect and track curbs using a combination of lasers, radar, and cameras. The vehicle can project a curb's location onto the visual image of the road and process the image to generate a template of the curb's appearance. Tracking the curb in the image as the vehicle moves lets the system generate a preview of the curb and road's path. Accordingly, the vehicles can differentiate between objects on the road and those on the side of the road, so they can safely negotiate turns.

Navlab builds a map of the environment and tracks moving objects. It also scans a single line across the scene periodically using radio waves and laser beams. As long as the vehicle doesn't move too irregularly, it can track moving objects from frame to frame as the vehicle moves. Given the locations of the moving and fixed objects, and the heading direction of the vehicle, it's possible to estimate the time-to-collision for each object. Stereo-vision systems obtain 3D information from a scene to detect and track pedestrians.

Drive-by-wire is another recent technology that has made it easy to implement autonomous control in next-generation vehi-

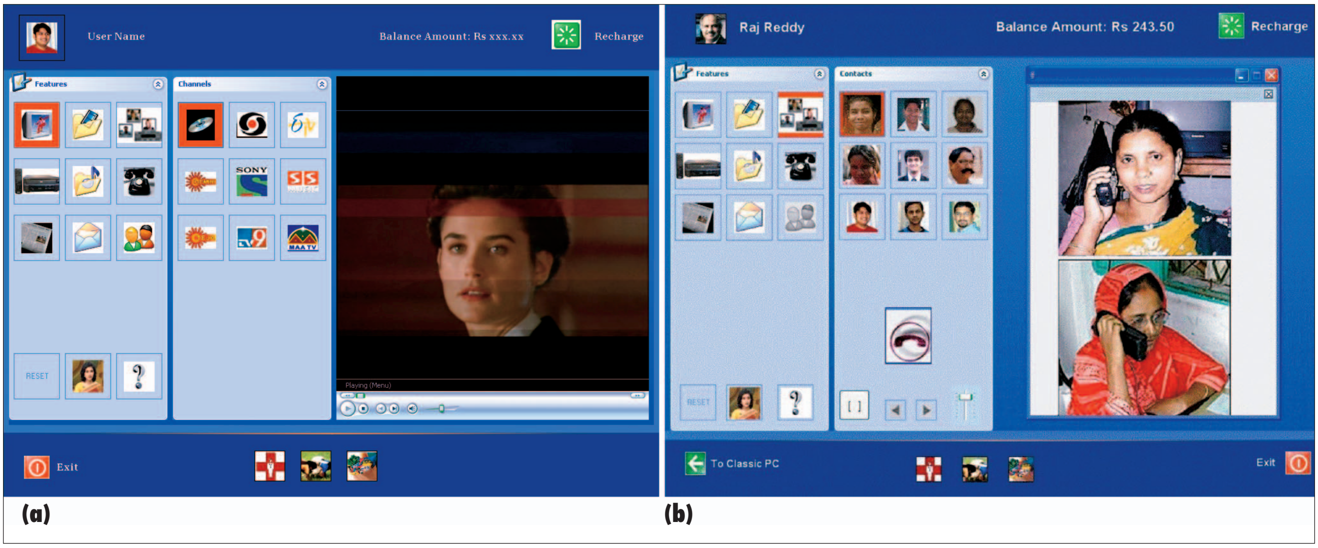

Figure 5. The PCtvt, a multifunction information appliance: (a) to watch TV, the user clicks on the picture of the TV and then selects a channel; (b) for the video phone, the user clicks on the video-phone picture and then selects one of the faces. (figure courtesy of CMU and Raj Reddy)

cles, such as Navlab. The idea is to delink the vehicle's mechanical controls from the actual operation of the associated devices. Mechanical movement of the accelerator, steering, brake, clutch, or gear controls generate different electric signals depending on the amount of pressure applied. The signals are sent to a central computer, which uses preset instructions to control the vehicle accordingly. Drive-by-wire eliminates human errors, such as pressing the accelerator too hard or the clutch too lightly. At the same time, it also optimizes fuel delivery. As a result, vehicles are safer to drive and more economical to maintain. It also makes driving more convenient because it optimizes cruise control and gear shifting.

The Grand Challenge race organized by DARPA illustrates current technology for fully autonomous driving capabilities. Five teams completed the Grand Challenge in 2005; four of them under the 10-hour limit. Stanford University's Stanley took the prize with a winning time of 6 hours, 53 minutes.

Although these technologies have been demonstrated at CMU and other research laboratories for more than 20 years, legal, regulatory, and reliability considerations have held back their rapid adoption. I hope that companies in collaboration with researchers will make accident-avoiding cruise-control systems affordable and accessible.

Human-computer interaction

Four billion people in the world subsist on less than $2000 per year.[13] Most of them don't know English and more than a billion are uneducated. Yet most PC designers assume that the user knows English (or a few other languages of industrial nations). Thus, it's not necessarily the technology that's the barrier to wide acceptability of and accessibility to computers, but rather companies that target their designs at affluent and literate consumers in industrial nations. Anyone can learn to use a TV or telephone or learn to drive a car, even though these are arguably some of the more complex technologies invented by society. The secret is to hide the complexity under a radically simple interface

At CMU, we've developed the PCtvt, a multifunction information appliance that primarily functions as an entertainment and communication device (see figure 5). It can be used as a TV, a DVR, a video phone, an IP phone, or a PC. To improve global accessibility, the system uses an iconic and a voice-based interface. It takes less than one minute to learn to use it (like turning on a TV set and choosing a channel), requires only two or three steps, and can be operated using a TV remote.

The basic technologies to add multimedia functionality to a PC have been around for over 15 years. By adding a TV tuner chip, a PC can become a TV and DVR. Adding a camera and microphone permits video and audio inputs, which enable telephone functionality (using voice-over-IP protocols) as well as video phone, video conferencing, and video and voice email—capabilities currently unavailable to most users even in industrial countries. The key improvement is to insist on a radically simple user interface that doesn't require literacy—what we call an *appliance model*.

So, even if users can't read or write, they can easily learn to use voice or video mail. If they can't use a text-based help system, they can benefit from the video chat. Of course, such solutions demand more bandwidth, computation, and memory, which present an interesting conundrum. To be useful and affordable for those who are poor and illiterate, we need a PC with 100 times more computing power at one-tenth the cost.

There are no technological reasons why we can't provide increased memory and bandwidth at a nominal cost. Projected exponential growth in computing and communication technologies should make this possible. It's a research challenge that is worth undertaking not only as an ecotechnology but also as good business. It would open up the large untapped customer base in emerging markets.

Natural language processing

Natural language processing provides tools for indexing and retrieval, translation, summarization, document clustering, and topic tracking, among others.[14] Taken together, such tools will be essential for successfully using the vast amount of information being added to digital libraries.

Although basic technologies and tools for natural language processing have been demon-

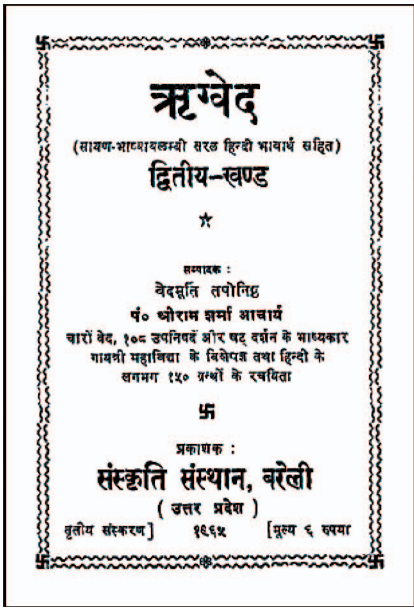
ऋग्वेद
(सायण-भाष्यावलम्बी सरल हिन्दी भावार्थ सहित)
द्वितीय-खण्ड
सम्पादक :
वेदमूर्ति तपोनिष्ठ
पं० श्रीराम शर्मा आचार्य
चारों वेद, १०८ उपनिषदें और षट् दर्शन के भाष्यकार
गायत्री महाविद्या के विशेषज्ञ तथा हिन्दी के
लगभग १५० ग्रन्थों के रचयिता
प्रकाशक :
संस्कृति संस्थान, बरेली
(उत्तर प्रदेश)
तृतीय संस्करण] १९६५ [मूल्य ६ रुपया

Title	Rig Veda
Author	Pandit Sriram Sharma Acharya
Language	Sanskrit
Subject	Philosophy
Publisher	Sanskriti Sansthan Bareli
Year	
Abstract	Rig Veda is the oldest of the Vedas. The Rig Veda is the oldest book in Sanskrit or any Indo-European language. Many great Yogis and scholars who have understood the astronomical references in the hymns, date the Rig Veda as before 4000 B.C., perhaps as early as 12,000. Modern western scholars date it around 1500 B.C., though recent archaeological finds in India (like Dwaraka) now appear to require a much earlier date

Figure 6. The cover page of Sanskrit book Rig-Veda that was digitized in the CMU's Million Book Project. (figure courtesy of CMU and Raj Reddy)

strated for the English language, the necessary linguistic support isn't yet available in many other languages. As a result, automatic translation and summarization among many languages isn't yet satisfactory. We must develop low-cost, quick, and reliable methods for producing language-processing systems for many languages of low commercial interest. Natural language processing systems must better understand, interpret, and resolve ambiguity in language. We must resolve ambiguities at the lexical, syntactic, semantic, and contextual levels to generate better-quality output. Also, most language-processing systems are computationally intensive, so we need effective ways of redesigning these systems to scale up in real time.

The first step to realizing such ambitious goals in natural language processing is to create a rich digital content repository. Many initiatives have begun to act as a testbed for validating concepts of natural language research. One such project is CMU's Million Book Digital Library Project, a collaborative venture among many countries including the US, China, and India. So far, over 400,000 books have been scanned in China and 200,000 in India. Content is made available freely from multiple sites around the globe (see www.ulib.org for a sampling of this collection). Figure 6 provides the cover page of a book in Sanskrit that has been digitized for free access to everyone.

Google, Yahoo, and Microsoft have all announced their intention to scan and make available books of interest to the public. Unfortunately, most of these will be in English and thus won't be readable by over 80 percent of the world's population. Even when books in other languages become available online, their content will remain incomprehensible to many people. Natural language processing technology for translation among languages isn't perfect, but it promises to provide a way out of this conundrum. A resource like a universal digital library, supported by natural language processing techniques, furthers the democratization of knowledge by making digital libraries usable and available to scholars around the world.

Furthermore, having millions of books accessible online creates an information overload. Our late colleague at CMU, Nobel Laureate Herbert Simon used to say, "We have a wealth of information and scarcity of (human) attention!" Online multilingual search technology provides a way out. It lets users search very large databases quickly and reliably independent of language and location, enhancing accessibility to relevant information. Additionally, machine-translation systems such as CMU's Avenue project aim to find quick, low-cost methods for developing machine-translation systems for minority or endangered languages. Also, question-answering systems such as IBM's Request provide natural language interfaces to information systems. Speech interfaces that use speech-recognition and synthesis systems will let the physically challenged, elderly, and uneducated benefit from online data.

Rapid access to relevant information and knowledge has become so important to society that that it has spawned a $100 billion search industry in the form of Google, Yahoo, and MSN. However, a lot of information exists in nondigital, multilingual formats that we must capture, preserve, and provide to users. Natural language processing techniques will help immensely as an indispensable front-end interface to such repositories of multilingual information. As more information becomes searchable, findable, and processable online, as Google Print and other such projects intend, we can expect digital libraries and natural language processing to become the most used technology of all.

Artificial intelligence

In many developing countries, neonatal care isn't readily available to many newborn children because hospitals are inaccessible and costly. The causes of many newborn deaths—infection, birth asphyxia or trauma, prematurity, and hypothermia—are preventable. The underlying issues include poor prepregnancy health, inadequate care during pregnancy and delivery, low birth weight, breast-feeding problems, and inadequate newborn and postpartum care.

Currently, village health workers trained in neonatal care make home visits to control diarrheal diseases and acute respiratory infection, offer immunizations, and provide better nutrition through micronutrients and community awareness. Unfortunately, scalability and sustainability have been problems, including ready access to a health worker, identifying and training health workers, and providing support and medicines in a timely manner.

Many believe that the poor, sick, and uneducated masses on the other side of the rural digital divide have more to gain from information and communication technologies than the billion people who already enjoy this technology.[15] Specifically, artificial intelligence-based approaches such as expert systems and knowledge-based systems that have been used in medical diagnosis and therapy applications since 1970s might also prove to be effective in the developing world and provide timely intervention, saving lives and money.

One such example is the *Vi*enna *E*xpert System for *P*arenteral *N*utrition of *N*eonates (VIE-PNN; www.ai.univie.ac.at/imkai/kbs/

vie-pnn.html). Planning adequate nutritional support is a tedious and time-consuming calculation that requires practical expert knowledge and involves the risk of introducing possibly fatal errors. VIE-PNN aims to avoid errors within certain limits, save time, and keep data for further statistical analysis. The system calculates daily fluid, electrolyte, vitamin, and nutrition requirements according to estimated needs as well as the patient's body weight, age, and clinical conditions (such as specific diseases and past and present-day blood analysis).

Creating an expert system for neonatal care in rural environments is perhaps a hundred times more difficult. The cost of use must be pennies rather than dollars, and several technical challenges exist. First, we need an information appliance that can host an expert system and engage in a dialog with a village housewife in a spoken local language (cell phones should be powerful enough to provide this functionality within a few years). We also must develop recognition, synthesis, and dialog systems for local languages.

Second, we must create an expert system with a database of answers to common problems, in the form of voice and video responses as well as text. When an answer isn't available locally, the user should be able to search a global database with more comprehensive frequently asked questions and answers. This in turn requires multilingual search and translation. If all else fails, the query should be referred to a human-expert task force for generating an answer for immediate and future use.

Finally, we must provide the necessary therapy. For example, it's not enough to tell AIDS patients that they need a three-drug cocktail. We must also tell them where to get it and how to pay for it.

Although I've presented several technologies and intelligent system applications that could significantly impact the well-being of our society—helping, in particular, those who are poor, sick, or illiterate—much work remains before these tools are routinely available and widely deployed. The productization of these technologies requires additional research and development with the aim of increasing reliability, reducing cost, and improving ease of use. This is the grand challenge for the next generation of researchers and entrepreneurs. It is my hope that governments, foundations, and multinational corporations will lead the way in creating technologies for a compassionate world in 2050.

The Author

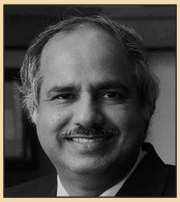

Raj Reddy is the Mozah Bint Nasser University Professor of Computer Science and Robotics at Carnegie Mellon University. His research interests include the study of human-computer interaction, robotics, and artificial intelligence. He received his PhD in computer science from Stanford University. He's a Centennial Fellow of the IEEE and an ACM Turing Award winner. Contact him at Wean Hall 5325, School of Computer Science, Carnegie Mellon Univ., Pittsburgh, PA 15213-3981; rr@cmu.edu.

Acknowledgments

This article is based in part on the Honda Prize lecture I gave in Nov. 2005. Praveen Garimella, Vamshi Ambati, and Hemant Rokesh helped prepare the article.

References

1. "World Population Aging: 1950-2050," tech. report, Population Division, 2002 World Assembly on Aging, 2002.
2. M.E. Pollack, "Intelligent Technology for an Aging Population: The Use of AI to Assist Elders with Cognitive Impairment," *AI Magazine*, vol. 26, no. 2, 2005, pp. 9–24.
3. D. Jajeh, "Robot Nurse Escorts and Schmoozes the Elderly," *Circuit*, 24 Aug. 2004, www.intel.com/employee/retiree/circuit/robot.htm.
4. R.R. Murphy and J.L. Burke, "Up from the Rubble: Lessons Learned about HRI from Search and Rescue," *Proc. 49th Annual Meetings of the Human Factors and Ergonomics Society*, vol. 49, 2005, p. 437; http://crasar.csee.usf.edu/research/Publications/CRASAR-TR2005-10.pdf.
5. J. Mostow et al., "A Prototype Reading Coach that Listens," *Proc. 12th Nat'l Conf. Artificial Intelligence* (AAAI 94), AAAI Press, 1994, pp. 785–792.
6. J. Mostow and G. Aist, "Evaluating Tutors that Listen: An Overview of Project Listen," *Smart Machines in Education*, K. Forbus and P. Feltovich, eds., MIT/AAAI Press, 2001, pp. 169–234.
7. "Remarks of Jacqueline Glassman," Presentation at the Intelligent Transportation Systems' World Congress, Nov. 2005, www.its.dot.gov/press/pdf/ITSspeech2005.pdf.
8. "Examination of Reduced Visibility Crashes and Potential IVHS Countermeasures," US Dept. of Transportation, Jan. 1995, p. 21; www.itsdocs.fhwa.dot.gov//JPODOCS/REPTS_TE/1BD01!.PDF.
9. C. Thorpe et al., "Safe Robot Driving," *Proc. Int'l Conf. Machine Automation* (ICMA 02), 2002; www.ri.cmu.edu/pub_files/pub3/thorpe_charles_2002_2/thorpe_charles_2002_2.pdf.
10. "The Short Report," tech. report, Urban Mobility Study, Texas Transportation Inst., 2002.
11. L. Bailey, "U-M Physicist: Smart Cruise Control Eliminates Traffic Jams," University of Michigan, 12 July 2004, www.umich.edu/news/index.html?Releases/2004/Jul04/r071204.
12. C. Thorpe et al., "Vision and Navigation for the Carnegie Mellon NavLab," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 10, no. 3, 1988, pp. 362–373.
13. Computer Science and Telecommunications board, *More than Screen Deep: Toward Every Citizen Interfaces to the Nation's Information Infrastructure*, Nat'l Academies Press, 1997; http://www7.nationalacademies.org/cstb/pub_screendeep.html.
14. J.G. Carbonell, A. Lavie, and A. Black, "Language Technologies for Humanitarian Aid," *Technology for Humanitarian Action*, Fordnam Univ. Press, 2005, pp. 111–137.
15. R. Tongia, E. Subrahmanian, and V.S. Arunachalam, "Information and Communication Technology for Sustainable Development, Defining a Global Research Agenda," 2003, www.cs.cmu.edu/~rtongia/ICT4SD_Front_Pages.pdf.

For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.