本田財団レポート No.112
第２６回本田賞授与式　記念講演（２００５年１１月２５日）

「社会的要請にこたえるロボット工学と知能システム」

カーネギーメロン大学
コンピュータサイエンスアンドロボティクス教授
（モザービントナッサーユニバーシティプロフェッサー）

ラジ・レディ博士

Robotics and Intelligent Systems in Support of Societal Needs

Commemorative Lecture
at the Twenty-Sixth Honda Prize Awarding Ceremony
on the 25th November 2005 in Tokyo

Dr. Raj Reddy
The Mozah Bint Nasser University Professor
of Computer Science and Robotics at Carnegie Mellon University

財団法人　本田財団

ラジ・レディ博士

カーネギーメロン大学
コンピュータサイエンスアンドロボティクス教授
（モザービントナッサーユニバーシティプロフェッサー）

The Mozah Bint Nasser University Professor of Computer Science and Robotics at Carnegie Mellon University

Dr. Raj Reddy

■略歴

1937 年	インドに生まれる
1958 年	マドラス大学（インド）卒業
1960 年	オーストラリア・ニューサウスウェールズ大学修士課程終了
1966 年	米国・スタンフォード大学 コンピュータサイエンス専攻 Ph.D
1969 年	カーネギーメロン大学助教授
1973 年	同大学　コンピュータサイエンス部教授
1979 年	同大学　ロボット工学研究所初代所長
1984 年	同大学　教授
1991 年	同大学　コンピュータサイエンス部部長
2005 年	モザービントナッサーユニバーシティプロフェッサー

■受賞歴（1984 年以降）

1984 年　仏ミッテラン大統領よりレジオンドヌール勲位授与
1991 年　ＩＢＭよりサーチ・ラルフ・ゴモリーフェロー賞
1994 年　アメリカ計算機学会（ＡＣＭ）チューリング賞
2001 年　インド大統領よりパドマブシャン（最高栄誉賞）
2004 年　大川財団より大川賞

印スリ・ヴェンカテスワラ大学、
仏アンリ・ポワンカレ大学、
豪ニューサウスウェールズ大学、
印ジャワハーラル・ネルー工科大学、
米マサチューセッツ大学、英ワーウィック大学、印アンナ大学、
インド情報技術大学(アラハバッド)の各大学より名誉理学博士号

■会員歴

- 電気エレクトロニクス・エンジニア協会
- アメリカ音響学会フェロー
- アメリカ人工知能学会（会長就任：1987-1989）
- アメリカ国立工学アカデミー会員
- アメリカ芸術・科学アカデミー会員
- 米国大統領情報技術諮問委員会共同委員長（1999－2001）
- イスラエル・ペレス平和センター執行委員会役員

■Personal History

1937	Born in India
1958	BE degree from the Guindy Engineering College of the University of Madras, India
1960	MTech degree from the University of New South Wales Australia
1966	Assistant Professor of Stanford University
1969	Carnegie Mellon faculty as an Associate Professor of Computer Science
1973	Carnegie Mellon faculty as a Full Professor of Computer Science
1979	Carnegie Mellon University Director of the Robotics Institute
1984	Carnegie Mellon University Professor
1991	Carnegie Mellon University Dean of School of Computer Science
2005	The Mozah Bint Nasser University Prosessor

■Awards (1984~)

1984	The Legion of Honor by President Mitterand of France
1991	The IBM Research Ralph Gomory Fellow Award
1994	The ACM Turing Award
2001	Padma Bhushan by President of India
2004	Okawa Prize

He has been awarded honorary doctorates from SV University in India, University Henri-Poincare in France, University of New South Wales in Australia, Jawaharlal Nehru Technological University in India, University of Massachusetts in USA, University of Warwick in England, Anna University in India and the Indian Institute for information Technology (Allahabad). He serves on the Board of Governors of Peres Institute for peace in Israel.

■Professional honors

- Fellow of the Institute of Electrical and Electronics Engineers
- Fellow of the Acoustical Society of America
- President of the American Association for Artificial Intelligence from 1987 to 89.
- Member of the National Academy of Engineering
- Member of the American Academy of Arts and Sciences
- Co-chair of the President's Information Technology Advisory Committee (PITAC) from 1999 to 2001
- Board of Governors of Peres Institute for Peace in Israel

このレポートは、平成１７年（2005年）11月２５日　ホテルオークラにおいて行われた第２6回本田賞受賞記念講演の要旨をまとめたものです。

25 November 2005 Hotel Okura,Tokyo

Robotics and Intelligent Systems in Support of Society

Raj Reddy

Carnegie Mellon University

Abstract:

Over the past 50 years there has been extensive research into robotics and intelligent systems. While much of the research has been targeted to solving specific technical problems, advances in these areas have led to systems and solutions that will have profound impact on society. This talk provides several examples of the use of such eco-technologies in the service of humanity, in the areas of robotics, speech, vision, human computer interaction, natural language processing, and artificial intelligence.

Technology Trends:

Underlying most of the advances in Robotics and intelligent systems are the unprecedented exponential improvement of information Technology. 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 billion bits per second bandwidth, all available for less than two thousand dollars. Barring the creation of a cartel or some unforeseen technological barrier, we should see a tera-PC by the year 2015 and a peta-PC by the year 2030. (Fig. 1)

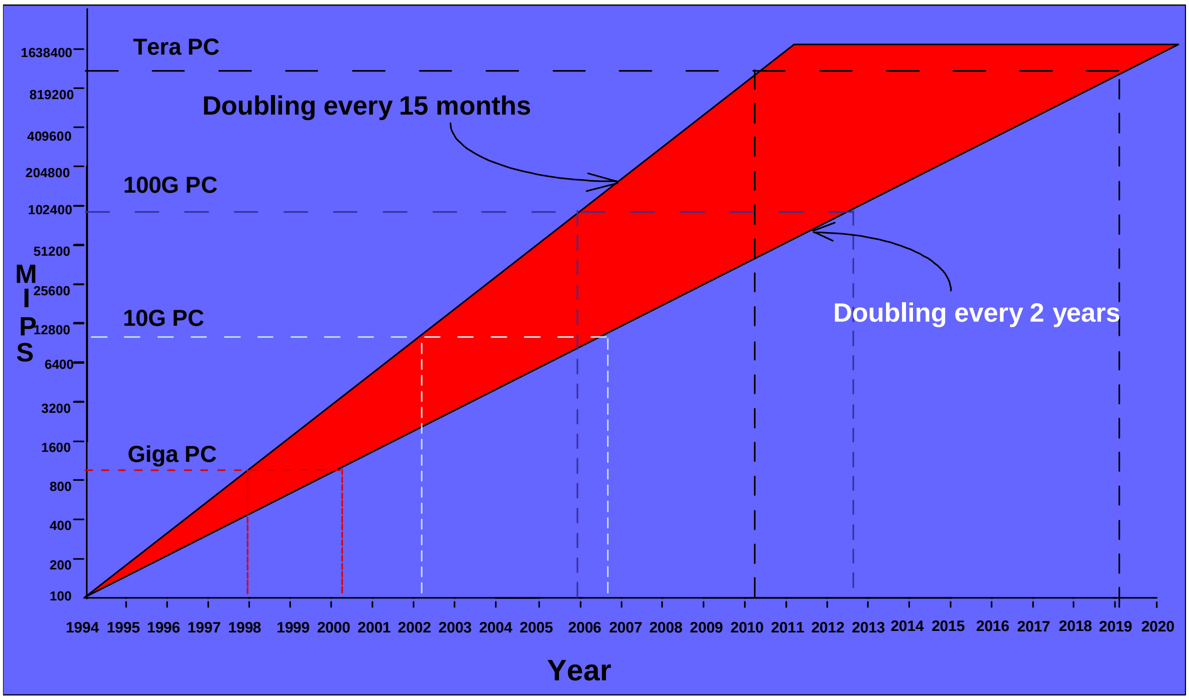

Fig 1. Exponential Growth Trends in Computer Performance

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, and the mating rituals will hardly be affected. Others, such as the way we learn, the way we work, the way we interact with each other and the quality and delivery of health care will undergo profound changes. Some of the computing power will be used to create self healing computers and networks that never fail or need rebooting

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

Most dramatic of all recent technological advances is the doubling of bandwidth every 9 months, propelled by the advances in fiber-optic technology. Today you can buy commercial systems that permit transmission of 1.6 terabits per second on a single fiber using dense wavelength division multiplexing (DWDM) technology. This technology uses 160 different wavelengths each capable of transmitting 10 giga-bits-per-second. Experimental systems are able to transmit as much as 25 tera-bits-per-second on a single fiber today!

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 can be carried on single fiber with room to spare. The main bottleneck today is not 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 protocols 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 have been transmitted before an acknowledgement can be received!

It is expected that the exponential doubling of memory and bandwidth will continue for the next 10 to 20 years, leading to the availability of terabyte disks and terabytes per second bandwidth at a cost of pennies per day! This qualitatively changes the way we think of how computers can be used in service to humanity. In this paper, we discuss several applications of robotics and intelligent systems that can be expected to transform the way we live, learn and work.

Robotics and the Aging Population:

As the life expectancy of the population of the world increases, it is expected that soon over 10% of the population will be over the age of 70. This age group can be expected to have minor disabilities impacting their quality of life. These disabilities can be grouped into three broad categories: sensory disabilities, cognitive disabilities, and motor disabilities and can be remedied through the use of robotic and intelligent systems technologies. In this Section, we will discuss the Nursebot project at Carnegie Mellon University capable of intelligent reminding, telepresence, data collection and surveillance, mobile manipulation and social interaction. Commercial availability of such solutions would pave the way for many senior citizens to lead productive lives much longer.

The cost of eldercare has grown dramatically over the last decade. Current nursing home costs range between $30,000 and $60,000 annually. Dramatic increase of the elderly population along with the explosion of costs poses extreme challenges to society. The current practices of providing care for the elderly population are already insufficient and in the future, there will be fewer young people to help older adults cope with the challenges of aging.

Robots for elder care must satisfying several requirements.

- Robots should must keep pace with the subject and not be either too fast or too slow. Current robots with fixed velocity lead to frustration on user's side.
- Safety while navigating in the presence of elderly people is a must. Given the limitations of the field of view of current vision systems, an eldercare robot may not always detect obstacles beyond its field of vision and risk accidentally hitting a person.
- A robot must be able accept and respond to voice commands. Current speech recognition and synthesis technologies are good enough to make this possible although there are several unsolved problems, such as the robot getting confused with what people speak, following the dialog, and noise in the environment.

The Pearl Robot developed at Carnegie Mellon (and currently continued at University of Michigan by Professor Martha Pollack [1]) has been demonstrated in several assistive care situations. (Fig. 2)

Fig 2. Encouraging Social Activity

It provides a research platform to test out a range of ideas for assisting elderly people. Two Intel® Pentium® 4 processor-based PCs run software to endow her with wit and ability to navigate; a differential drive system propels her; Wi-Fi wireless network connection helps her communicate as she rolls along; laser range finders, stereo camera systems, and sonar sensors guide her around obstructions; microphones help her recognize words; speakers enable others to hear her synthesized speech; an actuated head unit swivels in lifelike animation. Researchers hope that such autonomous mobile robots will one day live in the homes chronically ill elderly persons to perform a variety of tasks [2], such as (Fig. 3)

- Reminding elderly patients to visit the bathroom, take medicine, drink, or see the doctor.
- Connecting patients with caregivers through the Internet. The robot is a platform for tele-presence technology whereby professional caregivers can interact directly with remote patients, reducing the frequency of doctor visits.
- collecting data and monitoring the well-being of patients. Emergency conditions, such as heart failure or high blood sugar levels, can be avoided with systematic data collection.
- manipulating objects around the home such as the refrigerator, washing machine, or microwave. Researchers say arthritis is the main reason elderly give up independent living.
- taking over certain social functions. Many elderly people are forced to live alone, deprived of social contacts. The nursebot may help shut-ins feel less isolated.

Fig 3. Pearl Robot

At present these robots are very expensive. Pearl is a one-of-a-kind, costing close to $100,000. To ready a mass-market version of the automaton, researchers must overcome a rash of other issues as well, such as battery power and the robot's inability to navigate steps. In the Telepresence system, the caregiver should have access to more detailed status about elder's daily activities such as whether the patient took the correct medicine, ate meals etc. It should have the ability to monitor people and react to unusual behavior patterns such as alerting the caregiver if a person is not moving from a chair for a long time. The current system uses simple text strings for reminding a patient. As an elder's cognitive capacities diminish, they may require more detailed reminders. Even over shorter time spans, user's needs may change, for instance, during a period of illness. Intelligent reminders need to be based on the parameters like time of the day, timing of the previous interactions, user's mood, and actions that the user is supposed to do. Finally, future robots should be capable of pickup or move things as desired by the user.

As a society we need to find alternative ways of providing care to the elderly and chronically ill. Ways must be found not only lower the costs; they must also increase the comfort of living, and approach people with the level of dignity that our elderly deserve. Technological advances will no doubt make such Elder-care 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 the disaster locations are too dangerous or unreachable for human exploration. Additional constraints like the availability of rescuers, high or low temperatures, and hurricane force winds result in significant delays before human rescuers can start searching for victims. In most cases human rescuers need to retrieve victims within 48 hours to enhance survival rates. The disasters of the last decade have shown that advances like earthquake prediction are not enough. For example in the Kobe earthquake of 1995 several large structures including buildings and highways which were believed to be earthquake proof were destroyed. In both Kobe and Oklahoma bombing cases, the human rescue efforts alone were not enough resulting unnecessary loss of life. The huge robots available at the time of these disasters could not maneuver in the rubble and could not be used effectively.

Lessons learnt from these and other rescue situations motivated further research around the world leading to the creation of rescue robots that are small, light, and cheap. Today several organizations are actively participating in designing rescue robots that are small but can carry human sized payload. Some systems are able to maneuver over land and air and can detect sounds. Many of the robots have cameras with 360 degree rotation capability that can send back high resolution images. Some rescue robots are equipped with specialized sensors that can detect body heat or colored clothing. The 9/11 tragedy and the recent hurricanes have showcased the nascent robotics technology and the industry attempting to provide tools and systems for data gathering and rescue. (Fig. 4)

Fig 4. World Trade Center Site After 9/11

Current commercially available solutions are typified by The VGTV Xtreme [3], (Fig. 5) made by American Standard Robotics, is a 14 pound rescue robot that comes with a camera so searchers could see where it was too dangerous for them to venture.

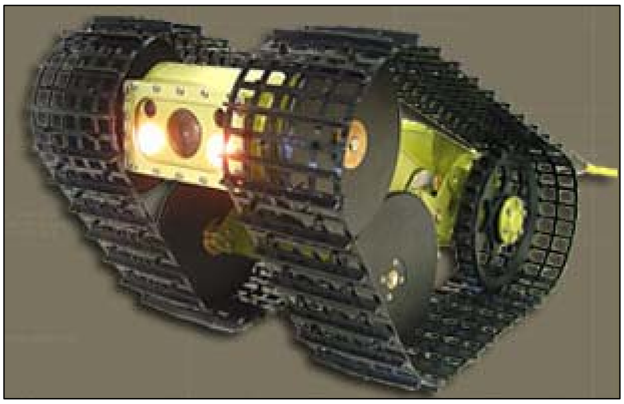

Fig 5. VGTV Xtreme Rescue Robot

Unlike the laboratory systems of the previous generation, this system is durable and flexible and stands up to the beating they take in the field. Current models have more ground clearance, so they can more easily go over rubble piles. Cameras are better, with longer zoom lenses, so even the robots don't have to journey as far into a structure. Operator controls have been improved and are compatible with equipment and gloves worn by rescuers. Also, the robots are waterproof, which makes decontamination easier.

Figure 6 showcases search and rescue systems that are being developed by The International Rescue System Institute (IRSI) of Japan [4].

Fig 6. Robots of International Rescue System Institute

These robots are capable of traveling hundreds of feet autonomously. Other examples are rescue robots like the "snake robot" from Carnegie Mellon University which helps assess contamination at waste storage sites or search for land mines. Similarly small and durable Tactical Mobile Robots (TMR) called PackBots were used by the military to explore Afghan Caves. The American Association for Artificial Intelligence conducts Rescue Robot competition annually (since 2000) in order to identify the capabilities and limitations of existing rescue robots.

Continued research and development is needed to enhance the capabilities and eliminate limitations rescue robotic systems. It is difficult to build robust rescue robots that can deal with the unexpected complexity of a disaster site. For example, the robots need to be adaptable to water clogged dark mines, collapsed building rubble, water clogged pipes etc. Also the robots should be able to adjust lighting, be resistant to extreme temperatures like a volcano or icy mountain environment.

A limitation observed in robot-rescue efforts at world trade center is that the robots can not go very far into the rubble because of short tether. At present, tele-operation from remote sites is often limited to about 100 feet from disaster site. Limitations of sensors, long setup times, human touch, ability to deliver food and medicines too an injured person etc needs to be addressed by next generation robots. (Fig. 7)

Fig 7. more Examples

Rescue robots can play a vital role in reducing the damage caused by disasters. Much work is underway to exploit ways in which technology can lead to robust systems that can aid in rescue missions, saving lives and keeping rescue worker from unsafe situations. Building rescue robots on a large scale that are affordable and efficient can be a challenge and demands increased participation and awareness from relevant government agencies. Progress is needed not only in improving the technology and research infrastructure but also in creating trained work force of First Responders who can effectively use such systems.

Speech and Reading Tutor:

Over a billion people in the world cannot read or write. Over 2 Billion people in the world are functionally illiterate in that they cannot understand the meaning of the sentences they could read. Advances in speech recognition and synthesis technologies provide an opportunity to create a computer based solution that can help to improve literacy. The solution involves creating 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, the speech recognition systems were not fast enough to recognize the connected spoken words in real time. For the reading tutor application, discovering and correcting mispronunciations requires not merely speech recognition, but also detecting deviations in stress, duration and spectral properties of a spoken word from that of native speakers.

The Reading Tutor developed by Dr Jack Mostow [5] at Carnegie Mellon University provides one such solution. To provide help in assisted reading, 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 Carnegie Mellon'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 modelled in part after expert reading teachers, but adapted to the capabilities and limitations of the technology. The current version runs under Windows XP on an ordinary personal computer with at least 128MB of memory. Though not (yet) a commercial product, the Reading Tutor [6] has been used daily (see photos below) by hundreds of children, as part of studies to test its effectiveness. (Fig. 8)

Fig 8. Reading tutor being used in a class room

Additional improvements are needed to fully realize the potential of this technology:

1. To make the Reading tutor a better listener, we must reduce speech recognition errors which are common in the recognition engine used. We also need systems that work with non-native speakers and local dialects and careless speech of children.

2. To improve assistance in learning, the user should get the help when they are stuck, better localization of reading mistakes, and an intuitive user interface for seeking help from the system.

3. To detect whether a student is engaged in answering questions, we need engagement tracing using response times. To improve on student performance the system must propose an approach by modelling the student behaviour.

4. To provide personalized help for each student, the system needs different levels of help based on the capabilities and learning abilities of the student. (Fig. 9)

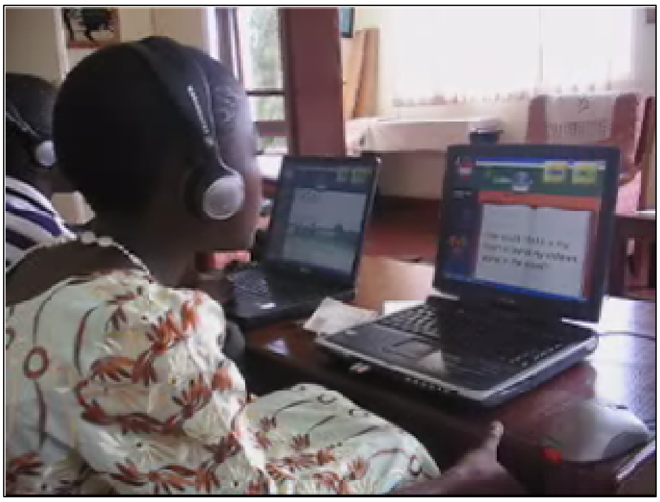

Fig 9. Reading tutor pilot in Ghana

Computer Vision and Intelligent Cruise Control:

Globally over a million people die annually in road traffic fatalities. In the U.S. this number is over 40,000 per year. The annual repair bill for the cars involved in these accidents in the U.S. exceeds over $55 billion. Sensing planning and control technologies made possible by the advances in computer vision and intelligent systems have the potential for reducing the deaths and the costs by over 80%.

It has been noticed that 40% of all vehicle crashes can be attributed in some form to reduced visibility due to lighting and weather conditions. With physical sensors monitoring the glare, fog and artificial light in the environment, alerts could be given to the driver so that he is more careful while driving through that particular area.

A majority of the road accidents [7] (over 70%) were found to be caused by human errors while driving. More often than not, these errors are due to speeding, driver fatigue or the driver being drunk. Over 240,000 accidents occur in the U.S. solely due to the driver falling asleep while driving. The casualties in such accidents can be reduced dramatically if there is an autopilot that can temporarily take control of the vehicle and navigate it to safe location.

Another major concern for drivers is the increasing number of traffic jams on highways and in cities. The Texas Transportation Institute estimates that in 2000 alone, major U.S. cities experienced 3.6 billion vehicle-hours of delay due to traffic congestion. This resulted in 21.6 billion litres of wasted fuel and $67.5 billion in lost productivity. Many of these jams are caused by excessive-braking due to driver panic or minor miscalculations while driving. Such irregular driving not only disrupts the traffic flow but also causes discomfort to the passengers travelling with the driver. In addition, it contributes to the wear of the vehicle resulting in higher fuel consumption. The other side-affect of irregular driving is that huge gaps are created on highways because the driver maintains a large distance from the vehicle in front. This leads

to a significant underutilization of the current roads.
Many of these problems can be avoided with the help of intelligent cruise control systems that regulate the speed of vehicle based on the time-to-collision.
In fact, studies show that majority of traffic jams can be completely avoided even if a fraction of vehicles are fitted with some form of driver assistance systems.
To create effective solutions to the above problems we need innovation in sensing technologies, control systems, actuators, route-planning algorithms and vision based navigation systems.
Dependable perception is a key component in building next generation cruise control systems. There is need for sensors that make use of light, sound or radio waves to detect and sense the physical environment. These sensors can be used to gather data like the speed, distance, shape and color of the objects surrounding the vehicle. Efficient object classification techniques need to be developed that can extract the shape of objects from this sensor data, and accurately classify them also based on color. Real-time object tracking systems are required to continuously monitor (and predict) the trajectories of these vehicles/people over time. Based on all this information, scene recognition algorithms can be developed to recognize the interactions between different objects in the environment and extrapolate them to a possible collision scenario in the immediate future.
In addition to the perception modules, there is also need for control systems and actuators to steer the vehicle. Feedback control systems with proportional control are required to help the vehicle maintain a constant speed by automatically adjusting the throttle based on the current speed. Efficient path planning and localization techniques are needed to enable the vehicle to autonomously navigate a set path.
Once these technologies are in place, it is possible to build collision warning systems with intuitive interfaces that warn drivers of the imminent danger ahead well in advance, so that he/she can stay clear. Collision avoidance and autonomous navigation systems can be designed that can plan and navigate the vehicle around the obstacles and danger scenarios without help from the driver.
Recent advances in sensor technology and scene recognition make it possible to develop collision warning systems that can give the driver sufficiently advance warning about the impending danger. Wherever such timely warning is not possible, collision avoidance systems take over and guide the vehicle to safety. Some current systems use laser-striping technology to detect the distance of surrounding obstacles for road-following and obstacle avoidance. Laser light is flashed on the scene and a camera is used to see the intersection of the light with objects in the scene. The distance between the vehicle and the object is then calculated by triangulation. Road-edges, side curbs and other obstacles can be detected using this technology. Omnicams and optical flow tracking are then used to build a 360 degree view of the world and track the movement of the edges along the motion of the vehicle. So whenever the vehicle

comes too close to the curbs or drifts off the road, the system recognizes it as a dangerous situation and warns the driver. The Vision and Autonomous Systems Center at the Carnegie Mellon University is building collision-warning systems for transit buses that also make use short-range sensors to detect and warn the driver of frontal and side collisions.

In the last decade, we also saw the emergence of fully autonomous navigation systems that can plan and navigate through real world traffic conditions. The key potential benefits of such systems are reduction of accident rate because of the collision-avoidance capabilities, reduction in driver fatigue, increase in fuel efficiency and better utilization of roads due to close-following (platooning). It has also been observed by a study published in Physics journal that a majority of traffic jams can be prevented if a mere 20% of vehicles on the road used adaptive cruise-control.

The Navlab vehicles [8] developed by Carnegie Mellon University's Field Robotics Center have demonstrated capability of autonomously operating at highway speeds for extended periods of time. These vehicles are equipped with perception modules to sense and recognize the surrounding environment. (Fig. 10)

Fig 10. Autonomous Navlab; System for Intelligent Cruise Control

Curbs are detected and tracked using a combination of lasers, radars and cameras. The location of curb is projected onto the visual image of the road and the image is processed to generate a template of the curb's appearance. Tracking the curb in the image along the motion of the vehicle facilitates the system to generate a preview of where the curb and road are headed. Accordingly the vehicles are able to differentiate between objects on the road and those that are on the side, and thus able to safely negotiate turns in the road. Sick single-axis ladders are used to build a map of the environment and track moving objects. The Sick scans a single line across the scene periodically with the help of radio waves and laser beams. As long as the vehicle carrying the sick does not move too irregularly, the moving objects can be tracked from frame-to-frame as the vehicle moves. Given the location of the moving and fixed objects, and the heading of the vehicle, it is possible to estimate the time of collision for each object. Stereo-vision systems are also used to get 3D information from a scene specially to detect and track human beings.

Drive-by-wire is another recent technology that has made it easy to implement autonomous control in next generation vehicles. The idea is to de-link 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 (eg. move the air-intake valve to allow more air for throttle). Drive-by-wire eliminates human errors like pressing the accelerator too hard or the clutch too lightly. At the same time it also optimizes the fuel-delivery strategy. As a result vehicles are safer to drive and more economical to maintain. It also makes driving more convenient because cruise control and gear shifting are optimized.

While these technologies have been demonstrated and available at Carnegie Mellon University and other research laboratories for over twenty years, legal and regulatory and reliability considerations have held back the rapid adoption of these technologies. (Fig. 11)

It is to be hoped that the legendary Honda innovation will make accidental avoiding cruise control systems affordable and accessible.

Fig 11. Autonomous Human Rescue System

Human computer interaction: Providing computer access to the people at the bottom of the pyramid.

There are 4 billion people at the bottom of the pyramid who subsist on less than $2000 income per year [9]. Most of them do not know English and over a billion are uneducated. Since most personal computer designers assume that the user knows English (or a few other languages of industrial nations), it might appear doubtful that people at the bottom of the pyramid [10] could ever benefit from the advances in Information Technology! However there is no reason for pessimism. Note that any one can learn to use a TV or a telephone or learn to drive a car, arguably some of the more complex technologies invented by society. The secret is to hide the complexity under a radically simple interface.

Thus, it is not the technology that acts as the barrier but rather the companies that target their designs at the affluent consumers. To be useful at the bottom of the pyramid, we need compelling content and an easy to learn user interface.

At Carnegie Mellon University, we have been developing a Multi Function Information Appliance that can be used by anyone. It primarily functions as an entertainment and communication device. It can be used not only as a TV or a DVR or a Video Phone or an IP phone, but also as a PC (hence the name PCtvt).

To improve accessibility to everyone on the planet, the system uses an iconic and voice based interface. It is easy to learn to use - with less than one minute learning time (like turning on a TV set and choosing a channel), requiring only 2 or 3 steps (two click paradigm), and operated using a TV remote. Figure 12 shows the system being used as TV along with the iconic user interface. The following video shows this system in operation.

Fig 12. 「PCtvt」 Multi-Function Appliance Two Click Interface

The basic technologies to add multimedia functionality to a PC have been around for over 15 years. By adding TV tuner chip a PC can become a TV and/or a DVR. Adding a camera and a microphone permits video and audio inputs enabling not only a telephone functionality (using VOIP protocols) but also capabilities such as Video Phone, Video conferencing, and Video and Voice email, capabilities that are currently unavailable to most users even in industrial countries.

The key improvement (or dumbing down!) is to insist on a radically simple user interface that can be used by even illiterate people– what we call an Appliance Model. Like a TV, the two click model and one minute learning time prescribe that any function should be as simple as power on and selecting a channel, and that anyone to learn to use it in less than a minute. Please see figure 13 for two click iconic interface of PCtvt.

Fig 13. Two Click Video Phone

It is interesting to note that an Illiterate person needs a more powerful PC than a PhD! She cannot use email to read or write, but can easily learn to use voice mail or video mail. She cannot use text based help system but can benefit from a video chat. Such solutions demand more bandwidth, computation and memory than conventional systems for the power users of PCs. It presents an interesting conundrum. To be useful and affordable at the bottom of the pyramid, we need a PC which has hundred times more computing power for 1/10 of the cost! There are no technological reasons (but some regulatory and political reasons) why increased memory and bandwidth required by a PCtvt cannot be provided at nominal costs! Projected exponential growth in computing and communication technologies should make this possible. It is a research challenge that is worthy of undertaking not only as an eco-technology but also as a good business. It would open up the large untapped customer base in the emerging markets.

Natural Language Processing and Digital Libraries:

Creating a universal, free to read, digital library containing all the books ever published is technically feasible today. Google, Yahoo and Microsoft have all announced their intention to scan and make available books of interest to public. Unfortunately most of these will be in English and will not be readable by over 80% of the world's population. Even when books in other languages become available online, their content will remain incomprehensible to most people on the planet. Natural Language Processing Technology for translation among languages is not yet perfect but promises to provide a way out of this conundrum. Such a resource furthers the democratization of knowledge by making this large digital library available to scholars around the world.

Natural Language Processing technologies provide tools for indexing and retrieval, translation, summarization, and document clustering and topic tracking [11]. Taken together these solutions are essential for successful utilization of vast amount of information contained in the future digital libraries. Even if we read a book every day, we cannot read more than 40,000 books in a life time. Thus having millions of books online and accessible creates an information overload. Our late colleague Nobel Laureate Herbert Simon used to say "we have a wealth of information and scarcity of (human) attention!" The online multilingual search technology provides a way out. It allows users to search very large data bases quickly and reliably independent of language and location and enhancing accessibility to relevant information.

The past few decades has seen major improvements in the area of Natural Language processing. Systems that perform well based on Statistical and Linguistic techniques have come up. Also systems that improve upon their performance using Machine Learning techniques are widely seen today. Natural Language processing which had been an area for linguistics and theoreticians is now being helped by the involvement of Computer Scientists and Mathematicians. A wide variety of linguistic resources have been built at various Universities furthering our understanding.

The Million Book Digital Library Project [12] at Carnegie Mellon University is a collaborative venture among many countries including USA, 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. The website at CMU, www.ul.cs.cmu.edu <http://www.ulib.org>, provides a sampling of this collection which is full-content searchable and readable. The following figures provide cover pages of two books, illustrating the multilingual nature of the collection. (Fig. 14) (Fig. 15)

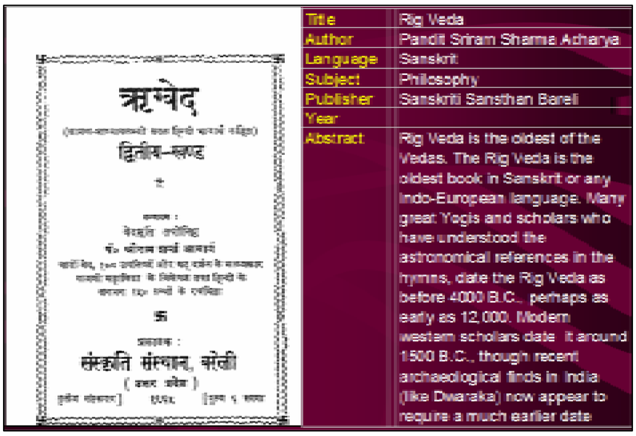

Fig 14. Sanskrit

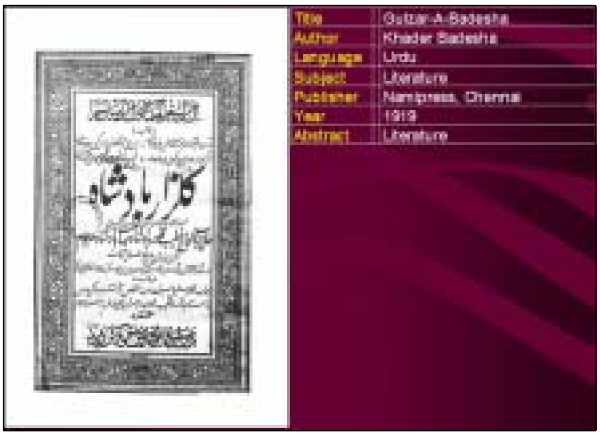

Fig 15. Urdu

Initiatives like Million Book Digital Library Project and other digitization projects will act as a rich test bed for validating concepts of Natural Language Research. A few decades ago even if such massive data were to be available, it would have been difficult to effectively utilize it, because NLP tools were not advanced enough. At present, machine translation (MT) systems like the AVENUE project at CMU aim to find quick, low-cost methods for developing MT systems for minority or endangered languages. We have question answering systems like "REQUEST" at IBM that provides natural language interfaces to Information Systems. Speech interfaces that use speech recognition and synthesis systems, permit the possibility of allowing physically challenged and elderly people (and even uneducated people) to benefit from the online data.

While basic technologies have been demonstrated using tools for English language, there are many languages where the necessary linguistics tools are not yet available. As a result, automatic translation and summarization among many of the languages is not yet satisfactory. We need to develop low cost, quick and reliable methods of producing Language processing systems for many languages of low commercial interest. Natural language processing systems have to better understand, interpret and resolve ambiguity in language. Ambiguities at the lexical, syntactic, semantic, contextual levels have to be resolved to generate to better quality output. Most language processing systems are very computationally intensive. Effective ways of redesigning these systems to be able to scale up in real time are needed.

Rapid access to relevant information and knowledge has become so important to society that that it has spawned a 100 billion dollar search industry, in the form of Google, Yahoo and MSN. As more and more information becomes searchable, findable and processable online, as is the intention of the Google Print and other such projects, we can expect digital libraries to become the most used eco-technology of all!

AI and Helping the Rural Digital Divide:

In many developing countries, one out of every 5 children dies before reaching the age of 5. Neonatal care is not readily available to many new born children because hospitals are inaccessible and costly. The root causes of many of these deaths are known and preventable. The major direct causes of neonatal death are infection, birth asphyxia/trauma, prematurity, and hypothermia. The underlying causes include poor pre-pregnancy health, inadequate care during pregnancy and delivery, low birth weight, breast-feeding problems, and inadequate newborn and post-partum care. In the systems currently in use, village health workers trained in neonatal care make home visits and manage such problems. (Fig. 16)

Fig 16. Making home visits by village health workers

These include a community-based integrated management of childhood illnesses. This includes programs such as control of diarrhoeal diseases, control of acute respiratory infection, immunization and nutrition including micro-nutrients and raising community awareness. Scalability and sustainability of the current solutions has been a problem, including ready accessibility to a health worker, identifying and training health workers, and providing the support and medicines in a timely manner.

It is believed that the poor, the sick and the uneducated masses on the other side of the rural digital divide would have more to gain in relative terms from information and communication technologies (ICT) than the billion people who already enjoy the benefits of this technology [13]. 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 may also prove to be effective in the developing world and provide timely intervention saving lives and cost.

One such example of direct relevance is the Vienna Expert System for Parenteral Nutrition of Neonates (VIE-PNN) [14]. The planning of an adequate nutrition support is a tedious time consuming calculation, needs practical expert knowledge and involves the risk of introducing possibly fatal errors. The aims of VIE-PNN were to avoid errors within certain limits, to save time, and to keep data for further statistical analysis. The daily fluids, electrolytes, vitamins, and nutrition requirements are calculated according to the estimated needs, the patients body weight, the age, and its clinical conditions (e.g. specific diseases, past and present day blood analysis).

Creating an expert system for rural environments for neonatal care is perhaps a hundred times more difficult. The cost of use must be pennies rather than dollars. There are also several technical challenges to be overcome.

- First is the availability and accessibility of an Information Appliance capable hosting an expert system and engaging in a dialog with a village house wife in a spoken local language. It is anticipated that within a few years, cell phones would be powerful enough to provide this functionality.
- Second is the development of recognition, synthesis and dialog systems for local language(s).
- Third is the creation of an expert system with a data base of answers to common problems, in the form of voice and video responses as well as text. When an answer is not available locally, a search is made on a global data base with more comprehensive frequently asked questions and answers from around the world. This in turn requires multi-lingual search and translation. If all else fails, the query is referred to human expert task force for generating an answer for immediate and future use.
- Finally we have the problem solving needed for providing the necessary therapy. For example, it is not enough to say to an AIDS patient that they need a three-drug cocktail, but we must also provide answers to “where to get it?”, “how to pay for it?”, and “what to do if you cannot afford it?”

ICT can be a powerful tool to facilitate and enable affordable solutions to those populations who are socio-economically deprived. It can provide a link to doctors and treatment using telemedicine, provide immediate access to health information through localized expert systems, help to reduce infant mortality and facilitate access to information about hygiene.

Conclusion:

In conclusion, the advances of the next fifty years will undoubtedly be dramatic. Capabilities such as videophones, accident-avoiding cars, and universal access to information and knowledge, entertainment on demand, learning on demand, and telemedicine will clearly come to pass. These advances will transform the way we live, learn, work and govern ourselves. Such capabilities can be used to further increase the gap between the haves and have-nots, or to help the poor, the sick and the illiterate. It is my hope that the Honda Foundation will lead the way to create eco-technologies for a "compassionate world in 2050"!

References:

1. M.E. Pollack, "Intelligent Technology for an Aging Population: The Use of AI to Assist Elders with Cognitive Impairment", AI Magazine, 26(2), September 2005.
2. Intel, "Robot nurse escorts and schmoozes the elderly", http://www.intel.com/employee/retiree/circuit/robot.htm , July 2004.
3. American Standard Robotics, "VGTV Xtreme**" ,** http://www.asrobotics.com/products/vgtv_xtreme.html
4. International Rescue System Institute, Japan, "", http://www.rescuesystem.org/tmp/NEW/en/framepage01.htm , November 2005.
5. J. Mostow, S. Roth, A. G. Hauptmann, and M. Kane, "A Prototype Reading Coach that Listens". Proceedings of the Twelfth National Conference on Artificial Intelligence (AAAI-94), American Association for Artificial Intelligence, August 1994, pp. 785-792.
6. J. Mostow and G. Aist, "Evaluating Tutors that Listen: An Overview of Project LISTEN. In K. Forbus and P. Feltovich (Eds.) Smart Machines in Education", MIT/AAAI Press, 2001
7. Chuck Thorpe, et. al. "Safe Robot Driving", Proceedings of the International Conference on Machine Automation (ICMA 2002), September 2002
8. Chuck Thorpe, M Hebert, T.Kanade, and S.Shafer "Vision and Navigation for the Carnegie Mellon NavLab", IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 10, No. 3, May 1988, pp. 362 - 373.
9. Computer Science and Telecommunications board, "National Research Council, "More than Screen Deep, Toward Every Citizen Interfaces to the Nation's Information Infrastructure." July 1997.
10. Raj Reddy, "The Role of Information Technology at the Bottom of the Pyramid" http://www.rr.cs.cmu.edu/it4pbp.ppt , October 2005.
11. Jaime G. Carbonell, Alon Lavie, Lori Levin and Alan W.Black "Language Technologies for Humanitarian Aid." Technology for Humanitarian Action pages 111-137.
12. Carnegie Mellon University, "Universal Library", http://www.ulib.org
13. Rahul Tongia, Eswaran Subrahmanian, V.S. Arunachalam, "Information and Communication Technology for Sustainable Development, Defining a Global Research Agenda", http://www.cs.cmu.edu/~rtongia/ICT4SD_Front_Pages.pdf , 2003.
14. Vienna Expert System for Parenteral Nutrition of neonates (VIE-PNN), " An Expert System for Parenteral Nutrition Composition of Newborn Infants" , http://www.ai.univie.ac.at/imkai/kbs/vie-pnn.html