Grand Challenge AI Applications

Hiroaki Kitano
NEC and Carnegie Mellon University
Lawrence Hunter
National Library of Medicine
Benjamin Wah
University of Illinois

Walther v.Hahn
University of Hamburg
Ryuuichi Oka
Real World Computing Program
Toshio Yokoi
Japan Electric Dictionary Research Inst.

Abstract

Grand challenge AI applications are applications with significant social, scientific and engineering impacts. This panel describes major research programs for grand challenge AI applications, and discusses why these applications are important, what technical challenges are, and what computing support levels are necessary. The panelists present grand challenge projects in the U.S., Europe and Japan.

1 Introduction

Grand challenge AI applications are *applications with significant social, engineering and scientific impacts, whose development will require substantial scientific and engineering break through, and use of high-performance computing and large-scale data resources.* In the high performance computers and communication act of 1991, the term **grand challenge** is defined as *a fundamental problem in science or engineering, with broad economic and scientific impacts, whose solution will require the application of high-performance computing resources.* Since this definition assumes grand challenges for scientific computing, or numeric computing field, it focuses on *a fundamental problem in science or engineering* rather than *applications.* The definition deliniated at the outset of this paper would clearify what are grand challenge AI applications.

In Feburary 1992, the National Science Foundation (NSF) hosted a workshop, chaired by Professor Benjamin Wah, titled *The Workshop on High Performance Computing and Communications for Grand Challenge Applications: Computer Vision, Natural Language and Speech Processing, and Artificial Intelligence.*

In October 1992, Kitano hosted a private workshop called *The Workshop on Grand Challenge AI Applications*, participated in by major Japanese AI researchers with several Japanese, American, and European observers [Kitano, 1992]. The next workshop will be held in October 2000, as *The International Workshop on Grand Challenge AI Applications: Accomplishments in the 20th Century and Prospects for the 21st Century.*

Apart from these workshops, there are several projects which are engaged in grand challenge AI applications. For example, Carnegie Mellon University, ATR Interpreting Telephony Research Laboratory, and other research institutions have been working on developing speech-to-speech translation systems — a representative grand challenge AI application. Already, several prototype systems, such as SpeechTrans, Dm-Dialog, SL-Trans, JANUS, ASURA, and Intertalker have been developed and publicly demonstrated. Recently, increasing numbers of research institutions and national projects have decided to tackle the grand challenge AI applications. Professor Hahn will describe the project verbmobil, a german national project to develop a speech-to-speech translation system.

Grand challenge AI applications often requires substantial involvement with non-AI, or even a non-computer science community. One typical example is the Human Genome project where molecular biologists play a major role. AI, however, is one of the critical aspects achieving success in the project. Dr. Hunter describes the international human genome project and how AI can contribute to the project, as well as getting benefits from the project.

Computational resources and technical breakthroughs would be necessary to implement the challenge. In the U.S., the high performance computing and communication project aims at development of high-performance massively parallel computers and ultra high-band width communication networks. In Japan, MITI sponsered the Real World Computing Program which was started in 1992. The projects aims at the development of high performance massively parallel computers and basic research for grand challenges.

High-performance computing and communication alone are not sufficient for the challenge. Availability of massive data resources is a critical factor in bringing the efforts to success. There are already several efforts to build large-scale knowledge-bases and data-bases. Dr. Yokoi describes the EDR's electric dictionary projects, and argues the need for the NOAH Knowledge Archive project.

Architectural Supports

Benjamin W. Wah
Coordinated Science Laboratory
University of Illinois
Urbana, IL 61801 USA
wah@crhc.uiuc.edu

Many grand challenge applications in AI are related to the representation, acquisition, storage, and processing of knowledge. As computer systems supporting these applications are generally built from knowledge and search, high performance computing and communication systems can make a dramatic difference in the amount of knowledge represented and the search performed. At the

NSF sponsored HPCC Workshop on Computer Vision, Speech and Natural Language Processing, and Artificial Intelligence last year (Wah, et al, 1993), participants identified many grand challenge applications that involve the use of AI methods. Examples of these applications include the following.

1) *Electronic librarian.* This involves the design of an electronic librarian that can navigate through a large amount of stored knowledge and information in an electronic library, interact with users in a natural language interface, and update the library automatically when new information is received. Relevant AI methods include knowledge representation, non-monotonic reasoning, and inexact searches. These methods must operate in large databases and adapt to different needs of users.

2) *Nation-wide job-bank.* This provides an online system that matches the skill profiles of job applicants against jobs listings. Relevant AI methods include constraint satisfaction, agent modeling, and possibly intelligent and interactive animation.

3) *Electronic market place.* This system provides a way for suppliers and customers to negotiate prices and conclude transactions.

4) *Intelligent planning and scheduling.* This involves the planning of resources in commercial applications, such as flight scheduling, and in manufacturing. It requires the modeling of a large amount of inexact information and the optimization of a multi-objective problem.

To understand how HPCC systems can be designed to support these applications, we must first understand the characteristics of these applications. In the following we first discuss the characteristics of AI application tasks. This is followed by a discussion of architectural requirements on HPCC systems for supporting these applications.

1) *Vast amount of knowledge and information.* Many AI applications deal with a large amount of knowledge and data. Examples include an electronic library and large corpora for supporting studies in this area. This information may be represented in many different schemes and may be interdisciplinary. It is anticipated that more than one internal knowledge representation schemes will be needed in an integrated system. Related architectural issues include the design of suitable representation schemes, hardware storage mechanisms, and storage methods that can adapt to different application requirements.

2) *Large number of users with different background.* A robust system supporting AI applications will need to have an intelligent interface that allows users with different background and goals interact with it. Advances in this area will require the design of natural language interfaces, interactive HDTV and computer graphics, and systems exploiting multimedia, visualization, and virtual reality technologies. These devices provide user friendly interfaces between users and systems.

3) *Inexact knowledge and information.* By definition, an intelligent system evolves with time. Consequently, knowledge stored in such systems will be updated incrementally, and existing knowledge may be invalidated with new knowledge acquired in the future. Non-monotonic reasoning is, therefore, an essential component of an intelligent system. Architectures suitable for such representation and reasoning may involve associative search. A familiar example is a hardware implementation of an artificial neural network.

4) *Knowledge acquisition and machine learning.* Machine learning is an important process for acquiring knowledge automatically in an intelligent system. Techniques used can be simple and involve the adaptation of parameters in an algorithm (such as a vision system), or can be very complex. Examples of the latter include the synthesis of vision algorithms (such as stereo) from examples of desired performance, the acquisition of models and problem solving heuristics, planning, and the acquisition of lexical, grammatical and discourse knowledge in natural language processing. Important issues to be considered in learning include the scaling and generalization of knowledge learned to more general application problems and a trade-off between the quality of the knowledge acquired and the amount of time taken to learn it. Computer systems supporting machine learning may include high speed massively parallel systems for inductive experimentation and a general purpose system for supporting high level reasoning.

5) *Multi-objective optimization.* Many AI applications involve the optimization of subgoals that are not precisely specified. Such problems are often posed as optimization problems with ill-defined objectives and imprecise constraints. Further, objectives and constraints may depend on measurable attributes that are too many to be included in the formulation. This ill-formedness leads to imprecise feedback (credit assignment) between actions and results. The problem here is analogous to a non-convex optimization problem over many hundreds or thousands of parameters. Techniques for solving these problems are often formulated as searches under syntactic and semantic constraints. Architectural studies in this area are directed towards load balancing, scheduling, and scalability.

6. *Integration of existing results.* In many AI applications known today, results have been demonstrated in a small scale in individual components. A lot of work remains to integrate these results into a single framework for real-world applications. For example, in spoken language translation, there is active research in integrating component modules, such as speech recognition, language analysis, message-content manipulation, message planning, language generation, and speech synthesis. Such integration may complicate the problem in an exponential fashion, demanding massive parallelism and high speed communication.

7. *Large computational requirements and distributed processing.* Many AI applications are characterized by algorithms that require large computational power, high communication bandwidth, and extensive memory. Examples include the search of large knowledge bases and the processing of a large number of distributed transactions and updates. Tasks may originate from applications with a large variety of features and objectives and exist in different granularity, requiring heterogeneous architectures that can support fine-grain as well as coarse-grain parallelism and that are connected by high speed communication networks. Moreover, these architecture may need to evolve with time when emerging technologies are available. Related and difficult issues include the design of software systems for supporting these diverse applications and the intelligent scheduling of computational resources among tasks.

The following are some architectural issues that are important in addressing the needs of grand challenge applications (Siegel, Abraham, et al., 1992).

- *Model formulation.* It is important to develop a model of computation that can abstract operational features and costs of parallel computation. Such a model can form the foundation for software developers and programmers to build their applications.

- *General purpose supercomputing.* Supercomputers are important tools for simulating prototypes before they are actually implemented. Researchers can test their algorithms on more realistic data before migrating to massively parallel systems and special purpose VLSI architectures.
- *Homogeneous architectures.* To effectively use massively parallel systems with homogeneous processors, it is critical to provide global (logical or physical) virtual memory that simplifies programming, language tools for expressing and for supporting parallelism and the interaction of concurrent activities, and compilers for optimizing system performance.
- *Special purpose VLSI architectures.* These are important when it is necessary to deploy systems that have stringent real-time requirements. In addition, emerging technologies may allow new problems to be solved in a reasonable amount of time.
- *Infrastructure support for prototyping architectures.* This consists of various simulation tools and prototyping facilities for computer architects to test their ideas. A particularly difficult problem involving AI applications is that it is often hard to instrument an operational system in order to test new ideas and algorithms.

Grand Challenge AI Applications Proposed in Japan

Hiroaki Kitano
NEC Corporation
Software Engieering Lab.
Shibaura, 2-11-5, Minato
Tokyo, 108 Japan hiroaki@cs.cmu.edu

At the workshop on grand challenge AI applications, which was held 6 October 1992, in Tokyo, various applications were proposed, which can be roughly categorized into six areas: games, science, evolutional systems, robots and replicant, communication, and infrastructure.

Games: Grand Master Computer Shogi and Go Systems.

Science: Anatomical Mapping of total brain wiring.

Evolutional Systems: Systems which evolve by themselves for use in building large-scale distributed and open systems, and to allow robots to be operated in a dynamically changing environment.

Robots and Replicants: Robots and programs which undertake daily matters, such as making coffee, answering the phone, cleaning floors, and delivering mails.

Infrastructure: Knowledgeable community, Global Architecture for Information Access (GAIA), and auditory scence analysis system.

Several interesting observations have been made. For example, the importance of Shogi and Go has been emphasized. As Chess has almost finished its role as a major technical challenge, Shogi and Go emerge as new challenges. While estimated search space for Chess is around 10^{120}, Shogi has about 10^{200} and Go has about 10^{230}. Rules for playing Shogi prevents search space to be narrowed down, even at the very end of the game. Go is a completely different game from Chess or Shogi. Developments of Grand Master Shogi and Go systems will be significant accomlishments, and the challenge is expected to produce various new technologies.

The goals for replicants and robots have been to carry out daily tasks. This is integrated with the challenge of achieving a highly flexible user interface research, since replicants and robots communicate using natural language and other media. While the goal of undertaking daily tasks does not have the flavor of *big science*, its accomplishment would bring major impacts to business and consumer markets. Even premature systems, if appropriate marketing can be carried out, could result in yet another workman.

Knowledgeable community and GAIA discuss the possiblity of building large scale and multi-purpose information infrastractures. While knowledgeable community stresses formal knowledge and multi-agent approach to support expert reasoning, GAIA aims at a more general information access infrastructure, reinforced by AI technologies.

Various breakthroughs in AI are necessary to meet the challenge. Basic technologies needed for the challenge are: massively parallel artificial intelligence, genetic algorithms, memory-based reasoning, parallel search algorithms, real-time planning and reasoning, integration of symbolic and connectionist processing, constraint-based processing, distributed AI, robust parsing, active vision, language learning, learning theory integrated with reasoning, situated theory, and flexible robot mechanism. This list contains a broad range of AI technologies, but it is not a random list. It selectively focuses on technologies to tackle real-world problems. In addition, technologies to handle incomplete and unstructured knowledge have been commonly recognized as an important factors. Building and sharing of large-scale data resources such as speech database, corpus, electric dictionaries, and image database, are considered to be of prime importance.

The list can be summarized as follows:

Central Nervous System: Massively parallel artificial intelligence technologies for reasoning, learning, search, real-time processing, integration of symbolic and numeric processing, and evolutional systems.

Knowledge Source: Very Large Knowledge-Bases, building and sharing of large-scale data resources. Technologies to represent and process incomplete and unstructured knowledge.

Input/Output: Presentation technologies, such as virtual reality. Recognition technologies, such as speech recognition, auditory scence recognition, and image understanding.

Motion Control: AI and robot mechanism integration.

Memory Storage: Tera-Peta byte size mass storage systems and efficient access methods.

Computers and communication systems should aim at achieving the following specifications, to be used as a basis for meeting the challenge.

Computers: Massively parallel machines with teraflops computing speed, tera-byte memory-space, peta or exa-byte scale mass storage, and tera-bits/second I/O. The CM-2/-5 equivalent systems should be desk-side size in 10 years.

Devices: Special purpose devices, such as WSI-MBR, WSI-NN, WSI-based audio stream separator, and vision devices, as well as Giga-byte memory chips and holographic memory.

Communication: Giga-bits/second end-to-end communication network. The first step should be B-ISDN with 622 Mbits/seconds. This should be up-graded to OC-192 of synchronous optical network (SONET). Then, it should be up-graded into a Giga-bits/seconds end-to-end network as the third phase.

AI and Grand Challenges in Biotechnology Computing

Lawrence Hunter
Director, Machine Learning Project
National Library of Medicine
Bethesda, MD 20894 USA
hunter@nlm.nih.gov

Several HPCC Grand Challenge Problems involve biomedical applications, and Artificial Intelligence methods are an important component of these approaches. The "Grand Challenges 1993" report identifies biotechnology computing and digital imaging as the two focus areas for development of advanced software technology and algorithms. It is in the area of biotechnology computing that AI has the potential for greatest impact.

The international Human Genome Project, underway since 1990, plans to produce complete genetic maps and sequences for human beings and a variety of model organisms by the year 2000. Currently, the sequence databases alone contain more than 500MB of sequence data, and are doubling roughly every 18 months. A DNA sequence, which can be represented as a string drawn from a four letter alphabet, is the blueprint of every living organism. The genome (complete DNA sequence) of a simple laboratory bacterium contains about 10,000,000 elements; human DNA contains about 3,000,000,000. Even the partial results obtained so far in this effort have had significant social and scientific importance. The causes of several inherited diseases have been identified, including Cystic Fibrosis, the most common inborn error of metabolism. These results have also led to improvements in therapy for such diseases, as well as for cancer, AIDS, heart disease and other major human health problems. In addition to health care, knowledge of genetic mechanisms and gene sequences has led to revolutionary changes in agriculture, materials science, pharmaceuticals and even law enforcement. It is probably not an exaggeration to say that the coming biotechnology revolution will be as rapid and significant as the computer revolution.

In fact, an important part of the underlying infrastructure for the explosive growth in our biological understanding and ability to do bioengineering involves the use of high power computation. Biotechnology computing involves a wide variety of specific applications, ranging from laboratory robotics in support of genome sequencing projects to the model-based simulation of cellular function and evolution. Many of these applications are well suited to AI techniques; they require finding patterns in large amounts of data, or qualitative modeling, or representation and reasoning about complex, structured knowledge. The connection is quite deep; renown biophysicist Harold Morowitz claims that "computers will be to biology as math is to physics." It is beyond the scope of this abstract to summarize all of this work, but see (Hunter, 1993), or the special issues of the Communications of the ACM (1991) or IEEE Expert (forthcoming) for surveys.

The amounts of data involved are often quite large, so that AI approaches must be scaled up in order to be effective. That is why the biotech area is such fertile ground for AI/HPCC research. Here I mention as examples three brief cases from my own work. First is the application of Bayesian classification to classifying protein structural elements (Hunter & States, 1992). This research involved the largest application of the AUTOCLASS III program, developed at NASA. The inputs to be classified were more than 10,000 protein fragments, each described by as many as 72 real numbered features. Running on a powerful SGI 4D280 server, the classification took more than three weeks of CPU time. Several other promising approaches to understanding protein structure with AI techniques are as yet computationally intractable. One approach has been the development of a massively parallel version of AUTOCLASS which should significantly expand the potential application area.

Another large scale AI application involved the development of a novel unsupervised learning algorithm that could be applied to massive, unsegmented datastreams (Harris, Hunter & States, 1992). This project involved finding clusters within the protein sequence databases. These sequence databases contain more than 70,000 sequences containing more than 150,000,000 elements. The cluster works by looking first at all the pairwise similarities (more than 2,000,000) of them, and then does an approximate transitive closure of the binary similarity measures. This application takes more than 10 hours of supercomputer time to complete, and needs to be updated periodically as the databases grow. The resulting classification is supplanting manually derived definitions of protein families and domains, and provides a valuable tool for automatically annotating the results of large scale genomic sequencing projects.

A final example identifies the potential synergistic role of AI techniques and other advanced algorithms. The bone disease Osteogenesis Imperfecta (OI) is caused by mutation in the gene for the protein collagen. The range of manifestations of the disease is quite broad: some mutations are lethal, and others are nearly asymptomatic. A goal of OI researchers is to understand the molecular mechanisms that determine the significance of a particular mutation. Molecular dynamics simulations of collagen in its entirety are computationally intractable, and none of the small models proposed have led to significant insights into the mechanisms of the disease. In current work (Klein, Hunter, Byers, unpublished) we have applied machine learning techniques to sort through the hundreds of potentially relevant features of the environment of the mutation to identify which are in fact relevant to lethality. The results of the machine learning application are then used to guide molecular dynamics experiments to simplified models that are both computationally tractible and capture the important molecular interactions. In this way, heuristic machine learning methods can be used to focus attention on important subproblems where more formal but computationally intensive methods can be applied.

A significant impediment to progress in this area is the lack of support for AI research on very high speed computers. For example, there are very few available supercomputer LISP environments available to researchers. This is due in large part to lack of demand from AI researchers themselves. Although *LISP was one of the first languages available for the massively parallel Connection Machine, demand for other languages (especially parallel fortran and C) have dramatically outstripped LISP users, and support at Thinking Machines for *LISP development has dropped off. Also, many of these problems require massive memory as well as high speed computers. The sequence clustering work is constrained by space as well as time tradeoffs. The method requires approximately 700M of memory for the current databases, and the databases are doubling in size every 18 months. The program is much faster if it doesn't have to page, so in the near future we will require machines with a gigabyte of RAM. We have found that large scale AI problems tend to require more memory than CPU speed, and that large memory machines are an important need.

In summary, both scaling up AI applications and applying intelligent techniques to guide the use of non-AI methods are likely to be important in addressing grand challenge problems. AI researchers must work hard at scaling their approaches to grand challenge problems, and demand the supercomputing resources (including memory and software, not just megaflops) necessary to accomplish their visions.

The Real World Computing Program

Ryu-ichi Oka
Tsukuba Research Center
Real World Computing Partnership
Tsukuba Mitui Bldg., Floors.17
1-6-1 Takezono,Tsukuba City, Ibaraki 305, Japan

The Real World Computing Program (RWC) is a new national research program funded by Japan's Ministry of International Trade and Industry (MITI) at a level $500 million for the next 10 years. RWC has as its overall technical objective the development of flexible and advanced information technologies that are capable of processing a variety of diversified information (such as images, speeches, texts, and so forth). RWC emphasizes technologies that match the flexibility of human information processing capabilities such as pattern recognition, handling incomplete information, and learning, all of which are manifested in the way people solve problems in the real world.

The approach to the objective of RWC research program is structurally explained as follows. The flexible processing will be based on four types of function; understanding, learning, inference, and control. Patterns and symbols are basic and common materials laid in the real world. The functions will be required to carry out the integration of symbols and patterns. The integration will be realized by exploring both problems of representation and algorithm which work in the real world. The representation is the framework for modeling the real world and also strongly related to parallel systems including optical systems. The algorithm is the dynamics for solving problems formulated in terms of the representation. Neural models are candidates that are expected to create both the representation and the algorithm.

Five major research themes in the following give more concrete descriptions of the technical objective.

Theoretical foundations: The theoretical foundations for flexible information processing will be a general framework for solving problems by the integration of symbols and patterns in conjunction with learning and self-organizational ones. In order to reveal relationships between the integration and conventional methods, theories for pattern recognition, multivariate data analysis, probabilistic and statistic inference, neural computing, and machine learning will be deepened and unified.

Novel functions: The goal is to develop the technologies for the flexible information processing that conventional technology lacks in terms of robustness, openness, and real-time capability. Specific research goals include (i) establishing the schemes for flexible recognition and understanding of multi-modal information, including moving images, voices, texts, and gestures, and developing interactive information systems by which humans can communicate with such information media, and (ii) developing flexible systems that are able to autonomously sense and understand and also control the environment through interactions with the real world.

Massively parallel computing: RWC-1/2, massively parallel systems which will be developed to support the "real world computation". Research goals include (i) providing multiple computation paradigms such as shared memory, message passing, data parallel, multithreading, neural networks, (ii) establishing flexible adaptation to the execution status of problems and computation resources, i.e. they will have the mechanisms for being adapted to the system load condition, variety of properties of information, requirements of realtimeness, robustness both for hardware and software, etc., (iii) providing computation and space capacity to permit an unpredictable amount of computation which will permit programmers the naive description of their problems, and (iv) providing fancy fast interconnection network and synchronization mechanisms both of which are essential for massive and unstable activities like neurons and elements of other novel models.

The software architecture of the massively parallel system consists of the kernel, implementation language, base language, operating system, and programming environment layers. The base language supports the notion of object-orientation, reflection, and time dependent programming. Using reflection, the programming language semantics as well as the resource management strategy is modified by the programming language itself. The new language functionalities might be easily constructed by the base language. Since we have to consider "time" in the real world, the base language provides the facility to describe the real-time system.

Neural network systems: Here we wish to establish a quite new type of neural model or computation which is different from conventional models such as Backpropagation, Boltzmann machine, and so forth. Connectionist models will be pursued as candidates that will provide a new scheme for representations of the real world. The domain of "neural" research should be extended to deal with materials in the real world.

Optical computing systems: This part of RWC aims at establishing basic technology for (i) optical interconnection devices and networks, (ii) optical neural models, devices, and systems, (iii) optical logic devices, circuits, and digital systems, and (iv) advanced opto-electronic integrated circuits development environments.

The objective of RWC meets the grand challenge problems in AI in both conceptual and computational aspects. First, the approach from the pattern side toward the integration of symbols and pattern should reveal new models or algorithms suitable to obtain new function with robustness and openness beyond ones of conventional pattern recognition and AI. Scalable algorithms with simplicity and transparency will be also expected to process a large mount of raw data such as images, voices, texts, and so forth. Second, real-time realization of those models or algorithms will require a new architecture of massively parallel computation which has a wide spectrum of I/O channels that contact tightly with the real world.

The typical application of the RWC will include the realization of humanoid assistant which is capable to support human activities interactively between users and the real world. The support means that i) to suggest a new viewpoint to a human user in his creative activities in terms of automatic search and inference of items in a large amount of data base including images and voices, ii) to take care of handicapped peoples in daily life by means of intelligent monitoring and so forth.

In conclusion, RWC is not a simple successor of the Fifth Generation Computer Project, but will promote challenging researches to solve intrinsic problems related to AI and pattern recognition. The key concept of RWC is the integration of symbols and patterns in the best sense of the words.

Very Large-Scale Knowledge Bases - From Lexical Knowledge to World Knowledge -

Toshio Yokoi
Japan Electronic Dictionary Research Institute (EDR)
Mita-Kokusai Building Annex, 4-28, Mita 1-chome, Minato-ku, Tokyo 108, Japan
yokoi@edr.co.jp

The major themes of very large-scale (massive) knowledge, very large-scale knowledge acquisition and very large-scale knowledge bases, are the basic "Grand Challenge" common to all attempts in the AI field. Two projects are presented here: the EDR Electronic Dictionary Project and the NOAH Project. The EDR Electronic Project, which is presently under way, deals with lexical knowledge. while the NOAH Project that is being planned will deal with world knowledge. The NOAH Project will pursue a new development in the technology by uniting knowledge processing and natural language processing. This project will need to be carried out under appropriate international cooperation. An international conference of KB&KS '93 will be held in December of this year in order to create such worldwide consensus.

1. EDR Electronic Dictionary Project

The EDR Electronic Dictionary Project is a nine-year project that started in 1986. It aims to research and develop large-scale computer dictionaries required for enabling computers to process and understand natural languages. The target languages are Japanese and English. Total funding for research and development is 14 billion yen (approximately $100 million). Seventy percent of this amount is financed by the Japanese Government while the remaining 30% by eight representative computer manufacturers. The Japan Electronic Dictionary Research Institute, Ltd. serves as the promoting nucleus.

All the results of the EDR project will be sold at reasonable prices. The same conditions regarding the use of the EDR Electronic Dictionaries will be applied to both domestic and overseas users. We will set the prices much lower than those of the machine-readable dictionaries that are currently on sale. Naturally, for academic users, such as universities and public research institutions, special measures will be arranged, including prices, etc. Discussions for realizing sound mutual cooperation with various organizations overseas have finally started. To ensure that this cooperation continues permanently and produces fruitful results, the results of the Electronic Dictionaries must be inherited by an organization established to develop them. Discussion of this issue has just started and has not yet reached the stage where any official comments have been expressed. Yet, a suggestion would be to establish an appropriate organization to expand and promote better use of the Electronic Dictionaries. Any corporation or person who wishes to use the Electronic Dictionaries should be a member of this consortium. The members would pay an entrance fee and an annual membership fee. The consortium would be run using the funds gathered from these fees and with the assets inherited from EDR. The membership would be roughly divided into two classes: one for commercial users and the other purely for academic users. The members, no matter whether they are domestic or overseas, should equally enjoy the benefits of the consortium. Furthermore, the consortium should actively promote the creation of an international network of electronic dictionaries by, for example, establishing a working relationship with similar overseas organizations.

2. NOAH (Knowledge Archives for Building and Sharing Very Large-Scale Knowledge) Project

In March of last year, the first edition of its plan was drawn up and official examinations of the project started. The NOAH Project centers around knowledge processing and natural language processing, and aims to research and develop new technologies related to knowledge itself. Natural language is considered knowledge representation language used widely and steadily. Documents expressing organized information in natural language are considered the basic source of knowledge. Therefore, natural language processing and text processing are considered the basic technology of knowledge processing. Based on these views, this project aims to develop various new functions that acquire and collect massive amounts of knowledge, share the knowledge, reuse the knowledge and support the creation of new knowledge. It is a project related to technologies for software, knowledge, and information widely associated with such fields as natural language processing, text processing, knowledge engineering, hypermedia, advanced database, software engineering, etc., and with the knowledge and information technologies. The sharing of knowledge is naturally to be carried out on a global scale. The NOAH Project aims to research and develop new technologies for realizing global-scale knowledge sharing environments. The foundation for this would be the sharing of knowledge related to the various languages, that is, the sharing of linguistic data such as the electronic dictionaries. By developing mutually cooperative international networks achieved by the Electronic Dictionaries project, the NOAH Project strives to seek the ideal form of new international cooperation. For this reason, an international conference is scheduled to open in Tokyo this year. This conference will be promoted from its planning stage with wide-range participation of governmental organizations and researchers from various countries. We strongly hope that the Knowledge Archives Project will stimulate similar projects within each country and play a significant role in the research and development of knowledge processing and natural language processing.

THE PROJECT VERBMOBIL: Mobile speech-to-speech translation

Walther v. Hahn
University of Hamburg, Germany
vhahn@nats4.informatik.uni-hamburg.de

The project verbmobil is a German national project for the research and development of speech-to-speech translation system. The project overview is as follows:

Background

- Increasing Language problems in the European Community
- Foreign language technology as an economical factor to raise the international market accessability of the industry.

Ambitious tasks are a challenge and will increase the overall awareness for the field as well as concrete research activities in natural language and speech processing. The project takes up previous and ongoing work being performed at CMU and ATR, and will work in direct cooperation with these institutions.

Specific Goals:

Mobile speech-to-speech translation device from and to German. Two steps are planned:

1. the contributions of the dialogue partners are both translated to written English on a screen,
2. additionally there is a target language generation and speech output.

Scenario: Negotiation dialogues starting with time schedule tasks.

Size and schedule: 20 partners from universities and industry will participate with a total funding budget of 10 - 15 million DEM per year. Additional financial inhouse support of the industrial partners is expected. The run-time of the project will be 8 - 10 years, beginning with 1993. Mile-stones are every 3 years.

Research Issues

- Software architectures for speech and language integration.
- Cognitive models of human translation and interpreting.
- Development of functional specifications for different types of MT systems.
- Description and processing of spontaneous language (hesitations, ellipses, prosody etc.).
- Evaluation and integration of symbolic and neural network models for different tasks (hybrid modules).
- Integration of domain knowledge, cultural knowledge, and linguistic knowledge.
- Large corpora research.
- Discourse representation

Central Design Decisions

- English as intermediate language in the sense that both partners are expected to have elementary passive skills and all contributions are translated into English (additionally) to provide control facilities.
- Transfer techniques, no interlingua.
- Interactive architecture for top-down and bottom-up processes

Expected Technical Effects

- Concentration of national efforts by one research issue in speech and language,
- Reinforcement of international cooperation,
- Opportunity for natural language processing to solve real-world tasks,
- Evaluation of what is feasible at the moment under real-application conditions,
- Professionalizing of the natural language processing work,
- Linguistic robustness for handling spontaneous language.

References

[CACM, 1991] Communications of the ACM (1991) Special issue on "The Human Genome Project and Informatics" 34(11), November 1991.

[Harris et. al., 1992] Harris, N., Hunter, L, and States, D. J., (1992) "Megaclassification: Discovering Motifs in Massive Datastreams," Proceedings of the 10th National Conference on Artificial Intelligence (AAAI-92), AAAI Press, Menlo Park, CA. pp. 837-842.

[Hunter and States, 1992] Hunter, L, and States, D. J., (1992) "Bayesian Classification of Protein Structure," IEEE Expert 7(4):67-75.

[Hunter, 1993] Hunter, L. (1993) "Artificial Intelligence and Molecular Biology," AAAI Press, Menlo Park, CA.

[IEEE Expert, 1993] IEEE Expert (forthcoming) Special issue on "The Applications of Artificial Intelligence to Molecular Biology." to appear in 1993.

[Kitano, 1992] Kitano, H. (Ed.), *Proceedings of the Workshop on Grand Challenge AI Applications,* Tokyo, 1992. — A collection of the revised version of papers is available as: Kitano, H. (Ed.), *The Grand Challenge,* Kyoritsu Publishing Co., Tokyo, 1993.

[Siegel et. al., 1992] Siegel, H. J., Abraham, S., et al., (1992) "Report of the Purdue Workshop on Grand Challenge Applications in Computer Architecture for the Support of High Performance Computing," *Journal of Parallel and Distributed Computing*, 16:199-211.

[Wah et. al., 1993] Wah, B. W., Huang, T. S., Joshi, A. K., Moldovan, D., et al., (1993) "Report on Workshop on High Performance Computing and Communications for Grand Challenge Applications: Computer Vision, Speech and Natural Language Processing, and Artificial Intelligence," *IEEE Transactions on Knowledge and Data Engineering,* 5(1):150-168.