J. Indian Inst. Sci., Mar.–Apr. 1990, 70, 91–99

A survey of work in AI at the Indian Institute of Science

V. V. S. SARMA AND S. J. NORONHA
Department of Computer Science and Automation, Indian Institute of Science, Bangalore 560012.

Received on March 29, 1990.

The fascinating field of artificial intelligence (AI), for long confined to the research laboratory, has in recent years generated remarkable interest among the industry and the general public. Much of this may be attributed to the maturing of several problem-solving and knowledge-representation techniques that has resulted in a profusion of expert systems and has lent engineering respectability to a field full of promises but with few instances of fulfillment. Work in AI at the Indian Institute of Science, however, has progressed over more than a decade, with developments in theoretical issues in AI as well as in applications. In this paper we attempt to survey past accomplishments as well as ongoing research at the Institute.

Artificial intelligence is the study of how to make computers do things at which, at the moment, people are better[1]. While computers outperform humans at numerical computation, tasks such as vision, planning, and understanding natural language which are so easy for people, present remarkable difficulties to a computer. A clue to the origin of these difficulties is the observation that the brain is the product of a million years of evolution and has been designed under the pressure of survival, whereas the computer has a fundamentally different 'hardware' architecture.

Nevertheless, the central assumption of AI is that at an appropriate level of abstraction what the brain does may be regarded as a form of computation. That is, every information processing system, human or mechanical, can be viewed as a physical symbol system, *i.e.*, a system containing processes that manipulate structures of symbols (physical patterns); and it is hypothesized that a physical symbol system has the necessary and sufficient means for general intelligent action.

An important consequence of this hypothesis is that while AI is concerned with effecting intelligent behaviour in computing systems, it is not committed to imitate the methods humans use. Notwithstanding this, the human brain and its mental processes are very often the inspiration for a particular AI technique or paradigm; a prominent example is the use of a massive network of simple computing units (neurons) to process information stored in the form of the weights of the connections between neurons (synapse strengths), in rough analogy with the functioning of the cerebral cortex. These artificial neural networks find effective use in highly parallel tasks such as vision.

The converse is also true: exploring information processing with the computer can contribute to the understanding of the human mind, and a second school of thought in AI attempts to study mental faculties through computational models. There is scope for strong interaction and cross-fertilization between AI and other disciplines such as psychology, neurobiology and linguistics, often leading to the development of new philosophies of the mind, *e.g.*, functionalism. Unfortunately, these disciplines are studied more or less independently and with limited interaction within the Institute.

While early work in AI held general problem-solving mechanisms to be the basis of intelligence, during the second generation of AI research it became increasingly clear that a large amount of specific domain knowledge was necessary in a competent AI system. From this insight and intensive work in knowledge representation rose the expert system, a computer program that uses extensive knowledge of a relatively narrow problem domain to display expert problem-solving performance in that domain. Intelligent decision support systems employ AI and expert systems technology to aid decision makers facing ill-structured problems.

1. Knowledge representation and problem solving

The application of AI techniques of knowledge representation and problem solving in decision support systems is being explored with project scheduling as a test domain[2]. These techniques include informed search methods such as the A* and the generalized branch and bound algorithms, and frame-based knowledge representation. Petri nets have also been proposed as a knowledge representation scheme and modeling tool for tackling the project scheduling problem, and are used in conjunction with a Petri net simulator to explore the generation of robust schedules and what-if analyses.

Over the last few years, there has been a sharp increase in interest in computational models of intelligence and information processing derived by analogy with the brain, such as neural networks, and in general in the cognitive sciences. One ongoing study explores mental representation of knowledge, *i.e.*, the organization of propositional and non-propositional knowledge in humans. Typical questions asked are: Are there any (computationally) useful insights to be derived from human knowledge representation structures? Do these structures place constraints on building intelligent systems with current symbolic and neural technology? What role do neural networks play as a possible cognitive architecture? Ongoing experiments study a two-level representation model, with a propositional scheme of representation on top of a non-propositional scheme[3].

2. Logic and commonsense reasoning

Traditional approaches to nonmonotonic reasoning with beliefs require a hierarchy of assertions, and reasoning is done on the basis of the most specific property that is known of an object. This poses difficulties to truth maintenance within the database of assertions, because less-specific earlier assertions which are contradicted by a new assertion cannot be deleted from the reasoning process. A scheme that identifies such assertions through a meta level of reasoning and maintains consistency has been proposed to tackle this problem[4].

Another problem with traditional logics derives from the principle 'ex contradictione quodlibet', *i.e.*, that everything follows from a contradiction, and therefore the database is required to start off consistent. This requirement may not be practicable, and we would like to reason with a database of facts even if it contains some inconsistencies. Paraconsistent logics, in which this principle does not hold, may be the answer, and the application of truth maintenance techniques to reasoning within these logics is under study.

3. Natural language understanding

Machine translation, one of AIs most unsuccessful ventures in the past, is once again back under the microscope. Efforts are on in the AI Laboratory to translate between English and the Indian languages[5]. Translation of simple sentences from English to Tamil, Kannada and Hindi, and from Tamil to English has been made possible, using syntactic and morphological analysis and synthesis. Facilities for multilingual input-output on computers, developing a production system-based parser for Kannada, reverse translation from Indian languages to English, and translation of more complex sentences are some of the topics currently being pursued.

In addition to syntactic analysis, semantic analysis is necessary in a natural language system that can carry on a conversation with humans. A system for conversing and carrying out actions in a restricted domain, the blocks world, has been designed. The system includes a morphological analyzer, grammar, procedural semantic analyzer, planner and a 3-D graphics interface.

Natural language processing techniques are being applied to develop database interfaces that can translate natural language queries to a relational database. Semantic grammars and logic grammars have been used to implement these interfaces, and present work is on applying context-free grammars. Future work will tackle the construction of database interfaces for Indian languages.

Several knowledge representation formalisms are being explored in the context of natural language understanding and generation. These include logical forms which resemble predicate logic, conceptual dependency notation, conceptual graphs, and frames. Of interest is the application of these formalisms to Indian languages.

In a separate project, software is being developed for computer-aided requirement analysis and design of engineered software (COMRADES)[6]. A prospective user presents the requirements for a program in natural language, and it is the task of COMRADES to detect errors in the specification, such as incompleteness, over- or under-specification, ambiguities and contradictions. It then refines the specifications with the help of the user.

4. Computer vision

A new computational model which has evolved from recent developments in empirical studies on human perception has been developed for the extraction of depth information from a pair of stereo images[7]. The computational theory is described in terms of the physical world constraints, *viz.*, uniqueness, compatibility, and figural continuity.

Machine perception requires the recognition of objects in the scene, and hence, understanding of shapes of objects. Hence shape recovery of objects from their single images plays a pivotal role in monocular computer vision. Geometric cues are used to fit shapes to features extracted from the gray level (shading) picture.

There have been several applications of neural networks in vision. A 4-layer neural network for textural segmentation has been simulated and applied to find textural boundaries. A new formulation of the stereo-matching algorithm based on feature vector consensus has been proposed and implemented using neural networks. Neural networks have also been used for illusory contour detection.

A model-based robotic vision system has been developed to recognize non-occluded objects which may be randomly positioned and oriented[8]. The 2-dimensional shape recognition technique involves creating a library of models of the various objects one may come across in a particular environment. The model of any object is a collection of its invariant features. When a scene is presented to the system, it extracts features of the various objects present in the scene, and tries to find the best match to the model of the desired object in the library. Once the presence of the desired object is established, the location of the object is computed and the coordinate information is passed on to the robotic arm.

5. Learning

The ability to learn is a central feature of intelligence, and hence there is a great deal of activity at the Indian Institute of Science in the area of machine learning. The tasks explored include learning to classify patterns and constraint satisfaction problems such as relaxation labeling. Pattern classification is the process of assigning to a particular input pattern the name of a class to which it belongs. This process may be implemented using a one-shot classifier, or by multistage methods[9]. In a multistage classifier several partial decisions are taken during the classification process, each transforming from a subfeature space to an intermediate decision set. The resulting phased decision structure is often referred to as a decision tree. Given the *a priori* class probabilities, class conditional feature distributions, and the misclassification set function, the optimal pattern classifier has been shown theoretically to be a decision tree only if the feature measurements have costs associated with them[10].

In general, a good definition of the classes may not be available, and in such a circumstance a learning program can be used to evolve the description of the class, *i.e.*, to learn to correctly classify any presented input. Learning class descriptions may occur in two ways, *viz.*, supervised and unsupervised learning. The former, learning from example inputs whose classes have been specified, is also known as concept learning, and the presence of noise due to classification or measurement errors makes it a hard problem. Unsupervised learning, or clustering, is the process of partitioning the set of patterns into meaningful groups, and does not have the benefit of training examples. Both these forms of learning have been addressed by research groups at the Institute.

Concept learning has been tackled *via* two main approaches: through stochastic models of learning such as learning automata, and through neural networks. Several M.Sc. and Ph.D. theses in the Department of Electrical Engineering have developed the game of learning automata concept and applied it to classification[11], relaxation labeling[12] and concept learning[13]. Current interest is in the study of systems consisting of many learning automata for problems in pattern recognition and concept learning, and their relation to the neural networks model which has also been used to solve the same problems. For example, the combination of a constraint satisfaction problem such as relaxation labeling with standard neural network models, such as the Hopfield network, has led to a new neural network model in which weights are dependent on the activation levels of the neurons.

Learning by observation, or clustering, has been explored by another group in the Department of Computer Science and Automation, led by Krishna and Narasimha Murty. Clustering has been applied to image segmentation, ECG classification, and classification of remote-sensed data. A new paradigm, knowledge-based clustering, has been proposed and used to group objects on the basis of their functionality[14]. This approach uses prior descriptions of the functions or purposes of a set of objects to deduce the classes into which they should be grouped.

A knowledge-based environment for clustering has been designed and several clustering algorithms which exploit contextual and domain-dependent knowledge have been proposed[15]. The environment generates and selects these algorithms based on the specification of the problem, and provides facilities for answering queries about the clusters. Current interest pursues a unified approach to knowledge-based clustering, and the relevance of nonmonotonic reasoning to this problem.

Introspective learning, in which the learners actively participate to acquire new concepts, is another type of learning under exploration. This has applications in knowledge acquisition and computer-aided learning.

6. Expert systems

In the area of expert systems, a complete EMYCIN-like expert system shell is being developed. At present the system comprises a rule and parameter specification language, a rule compiler which produces output in C, an inference engine, and a window interface. Future extensions include a knowledge-based editor and debugger, and options for generating compiled output in Lisp. The shell is currently being used to develop an expert system for cardiac diagnosis.

An expert recognition system has been developed to identify names of persons of Indian origin from among a large set of names[16]. This system extracts information from a database of publications by scientists from all countries and has been used to create a database of Indian nationals working in the area of biotechnology. The name recognition process has its basis in the observation that most Indian names have their origin in Indian languages, particularly Sanskrit, and therefore a major first step is the conversion of the textual name into a representation of its spoken form. A large amount of expertise in the form of rules is employed to break up the name into units called 'S-P parts', which are then matched with a pre-constructed database of S-P parts that are found in Indian names with high frequency. Additional knowledge goes into eliminating or accepting special cases and final decision making. The system displays a recall of about ninety per cent on actual databases. Current research explores an alternative implementation of this system using a neural network.

Design tools for generating a knowledge-based classifier for a given complex pattern-recognition problem are being developed[17]. The design tools themselves are knowledge based and use declarative definitional knowledge. The final classifier will be described in terms of 'plans' for suitable actions in a given situation. These plans are generated by analyzing the merit of each available piece of knowledge with respect to different situations. The time-space tradeoff in finding planning actions in an uncertain domain is met by using 'active plans' *i.e.*, plans which are extended upon demand. Techniques such as the Dempster-Shafer theory of uncertainty handling are being used to evaluate the effectiveness of alternative classification actions. The tools are being applied to the recognition of ethnicity of names, simultaneous digit and speaker recognition, and interpretation of Landsat images.

A rule-based expert system for analysis and diagnosis of high-voltage dc power system faults has been developed[18]. Information in the form of conduction patterns during normal and fault conditions is extracted and represented in the form of rules. When a fault occurs, symptoms such as the dc voltage and current are used by the rule-based system to diagnose the cause of the fault and suggest a control strategy for real-time recovery of the power system. The system is being extended to carry out deep reasoning about the converter. The application of pattern-recognition techniques to identify the faults from the conduction patterns is also being explored.

Knowledge acquisition for an expert system usually takes place with the help of a knowledge engineer or a knowledge acquisition program, and consumes a lot of effort. A novel way of looking at knowledge transfer, treating it as an educational activity has been suggested[19]. It is based on teaching, which is the most popular method of knowledge transfer in human education. Current findings indicate that knowledge teaching is advantageous and work is on to produce a knowledge teaching system.

A systematic methodology for knowledge acquisition and reasoning in the failure analysis of systems has been proposed[20]. This methodology combines enhanced versions of classical failure models with inference techniques from AI. The failure analysis processes that operate on these models utilize backtracking with constraint enforcement, implicit problem reduction, and bidirectional chaining of production rules.

An expert system for scheduling job orders in a flexible manufacturing system has been implemented. Heuristics specific to the particular system are employed in conjunction with a parallel scheduling algorithm to schedule the operations of several part types in such a way as to increase machine utilization, meet deadlines on jobs, and decrease the amount of time jobs remain within the FMS.

An expert system that learns to categorize the type of jaundice, given the patient's symptoms such as color of the eyes, drowsiness, blood pressure, etc., has been developed[21].

It uses a scheme of hierarchical counterfactual expressions to learn the four categories of jaundice from labeled examples.

The expert systems approach is being applied to the planning and usage of geostationary satellite orbit and radio resources[22]. Orbit-frequency planning and coordination is viewed as a multiperson coordination and multicriteria optimization problem. Two interacting models, the technical model and the human abstract system have been proposed to tackle the problem. Game-structured models are being used for the representation.and analysis of the latter.

Interpretation of remote-sensed images is another area in which the application of expert systems technology is being explored[23]. This problem requires the identification of land-use/land-cover mapping, mineral exploration, underground water resources, etc., from IRS satellite data. Several primitive keys or features, such as texture, pattern, context of interpretation, association between objects in the scene, etc., are used for this task. While classical pattern-recognition techniques can be used to extract texture information and to classify objects, they are inadequate for representing association and contextual knowledge. Heuristic knowledge is necessary for interpretation, and AI has a significant role to play here.

7. AI tools and languages

The Japanese Fifth Generation Computer programme has triggered major activity in AI all over the world. India's response has been the setting up of seven centres, including one at the Indian Institute of Science, for research and development in knowledge-based computing systems. At the Institute, the goal of this project is the development of computer architectures for symbolic computing. This includes research in hardware, systems software, and programming environments.

Hardware research aims at developing multicomputer systems with special accelerators for symbolic computing. Two prototype systems have been constructed and are being used with appropriate software environments for testing various parallel algorithms. A string-matching chip has also been developed. String matching is an important facility in AI systems, since it forms the basis for the execution of production systems. An interconnection network which supports on-the-fly pattern matching of messages has also been developed, and work is on to develop a unification chip to expedite execution of logic programs.

Other work under this project includes the design and implementation of a multi-paradigm programming language with facilities for unification, backtracking, and higher-order functions, and the design of systems software which supports intelligent manipulation of persistent objects. The major application being explored is an intelligent information retrieval system for accessing abstracts from databases of scientific and technical literature.

Another research team, in the Department of Electrical Engineering, has implemented a multiprocessor architecture for functional languages[24]. Functional languages are characterized by the fact that statements need not be executed sequentially. Further, they have no conventional assignment statements and hence are free from side effects. The applicative (*i.e.*, based on function applications) nature of these languages obviates many preprocessing transformations of the type used to extract parallelism, and hence they are attractive candidates for execution on parallel architectures. The architecture employed has six processing elements based on the Z80 microprocessor interconnected by a complete mesh network. Sublisp, a functional language which is subset of pure Lisp, has been implemented on this machine. The evaluation of programs in this language is based on dataflow principles and exploits the intrinsic parallelism expressed in a functional language program.

In yet another study, logic programming, one of the most popular programming paradigms in AI, has been the subject of investigation. An intelligent backtracking algorithm which uses variable- and data-based backtracking to reduce execution time has been proposed and implemented[25]. Ongoing research is aimed at applying model theory to study the parallel aspects of logic programs.

A five year project for research, training and development in AI and robotics has been under way at the Department of Computer Science and Automation since 1987. Sponsored by the DRDO, it pursues activities in three major areas, *viz.*, expert systems, computer vision, and natural language processing, concentrating on problems of relevance to defence.

References

1. Rich, E. — *Artificial intelligence*, 1983, McGraw-Hill.
2. Noronha, S. J. and Sarma, V. V. S. — Intelligent decision support systems for project management, *AI in industry and government*, (ed.) E. Balagurusamy, 1989, pp. 177–190, Macmillan India.
3. Srivastava, A. and Krishna, G. — Two-level representation for mental structures, submitted to *Cognition*.
4. Sridhar, V., Murty, M. N. and Krishna, G. — An extended nonmonotonic logic, *Proc. Int. Conf. Expert Systems for Development*, Kathmandu, Nepal, pp. 86–94, March 1989.
5. Santhi Seela, R. and Krishna, G. — Analysis of case-endings in Indian languages and bidirectional machine translation between Indian languages and English, *J. Indian Inst. Sci.*, 1990, This issue.
6. Sarma, V. V. S., Viswanadham, N., Negi, A., Roohie Mir and Mahesh, K. — *COMRADES: Computer aided requirements, analysis and design of engineered software*, Technical Reports 1 (Aug 1989) and 2 (Feb 1990), Deptt of Computer Science and Automation, Indian Institute of Science, Bangalore 560 012, India.
7. Banerjee, S. — *On stochastic relaxation paradigms for computational vision*, Ph.D. Thesis, Deptt of Electrical Engineering, Indian Institute of Science, Bangalore 560 012, India, 1990.
8. Rajeshwari, U. — *A vision system for a pick and place robot*, M.E. Thesis, Deptt of Electrical Engineering, Indian Institute of Science, Bangalore 560 012, India, 1989.
9. Kanal, L. N. and Dattatreya, G. R. — Pattern recognition, *Encyclopaedia of AI*, (eds) S. C. Shapiro and D. Eckroth, pp. 720–729, 1987, Wiley.
10. Dattatreya, G. R. and Sarma, V. V. S. — Bayesian and decision tree approaches for pattern recognition including feature measurement, *IEEE Trans.*, 1981, **PAMI-3**, 293–298.

11. THATHACHAR, M. A. L. AND SASTRY, P. S. — Learning optimal discriminant functions through a cooperative game of automata, *IEEE Trans.*, 1987, **SMC-17**, 73–85.

12. THATHACHAR, M. A. L. AND SASTRY, P. S. — Relaxation labeling with learning automata, *IEEE Trans.*, 1986, **PAMI-8**, 256–268.

13. RANJAN, S. R. — *Stochastic automata models for concept learning*, M.Sc. (Engng) Thesis, Deptt Electrical Engng, Indian Institute of Science, Bangalore 560 012, India, 1989.

14. SHEKAR, B., MURTY, M. N. AND KRISHNA, G. — Knowledge based learning using conceptual transformers, to appear in the *J. Intell. Robotic Systems Theory Appli.*, 1990.

15. NEGI, A AND MURTY, M. N. — A knowledge based approach to cluster analysis, *Proc. Indo-US Workshop on Signals and Systems*, Jan. 1988.

16. SRIVASTAVA, A. AND RAJARAMAN, V. — An expert system to create a database of potential consultants for a developing country, *Proc. Inter. Conf. on Expert Systems for Development*, Kathmandu, pp. 257–266, March 1–3, 1989.

17. CHITNIS, S. AND SARMA, V. V. S. — Knowledge based design of speech recognition systems, *Proc. Seventh Biennial Conv. of South East Asia Regional Computer Confederation*, pp. 111–118, 1988.

18. SWARUP, K. S. AND CHANDRASEKHARAIAH, H. S. — Expert system approach to fault diagnosis of HVDC systems, *J. Indian Inst. Sci.*, 1990, **70**(3) (to appear).

19. MOSHKENANI, M S. AND SARMA, V. V. S. — *Knowledge teaching: An alternative knowledge acquisition strategy for expert systems*, Technical Report CSA-89-21, Deptt of Computer Science and Automation, Indian Institute of Science, Bangalore 560 012, India, 1989.

20. HARI NARAYAN, N. AND VISWANADHAM, N. — A methodology for knowledge acquisition and reasoning in failure analysis of systems, *IEEE Trans.*, 1987, **SMC-17**, 274–288.

21. BHANDARU, MALINI KRISHNAN — Learning from examples using hierarchical counterfactual expressions, M.Sc. (Engng) Thesis, Deptt of Computer Science and Automation, Indian Institute of Science, Bangalore 560 012, India, 1988.

22. MOHANAVELU, K. S. — *Development and use of expert system approach for geostationary orbit-frequency planning and coordination*, Ph.D. Thesis, Deptt of Computer Science and Automation, Indian Institute of Science, Bangalore 560 012, India (forthcoming).

23. SARMA, L. C. AND SARMA, V. V. S. — Knowledge bases for remote sensing data interpretation: Issues and new directions, *Proc. of Prof. A. K. Choudhury Commemoration Symposium*, University of Calcutta, Feb. 21–23, 1990.

24. SANDHYA, K. — *A multiprocessor machine implementation for functional languages*, M.Sc. (Engng) Thesis, Deptt of Electrical Engineering, Indian Institute of Science, Bangalore 560 012, India, 1988.

25. RAJASEKAR, V. — Intelligent backtracking in logic programming, M.Sc. Thesis, Deptt of Computer Science and Automation, Indian Institute of Science, Bangalore 560 012, India, 1989.