Introduction

Artificial intelligence (AI) — over the last two decades and more — has grown to be a wide-ranging and active sub-discipline of Computer Science and Engineering. The broad scope of this sub-discipline can be gauged by listing the specializations that are covered in international AI conferences. Perhaps the largest such conference is **IJCAI** (International Joint Conference on AI) which is held every second year. The sessional headings from IJCAI-89 (the 11th such conference) included the following specializations:

* Automated deduction
* Cognitive models
* Commonsense reasoning
* Foundations
* Intelligent tutoring systems
* Knowledge representation
* Machine learning and knowledge acquisition
* Parallel and distributed processing
* Planning, scheduling and reasoning about action
* Real time and high performance
* Search
* Speech and natural language
* Tools
* Vision and robotics

Judged from the perspective of such conferences, this special issue on AI is not by any means a status report of this sub-discipline. The scope of this issue is highly limited and the intention is not to provide a balanced introduction to AI as a sub-discipline of computer science. For that purpose, one should consult the increasing number of good text-books on AI that are becoming available; (see in this context the review of the book *The Age of Intelligent Machines* by Chandrasekar, **page 434**).

The focus of this special issue on AI is clearly circumscribed by the first paper by Narasimhan (**page 361**). He argues in this paper that AI as science is primarily concerned with the computational modelling of agentive behaviour. Just as the physical sciences use mathematical formalisms to theorize about physical objects and physical events, AI, as the science of agents and agentive behaviour, uses computational (i.e. information processing) models to theorize about the capabilities and the intentional behaviour of agents. All biological organisms are agents in this sense. However, historically AI has been preoccupied with modelling the agentive behaviour of human beings. This special issue on AI, then, is delimited to the computational characterizing of human agents and their behaviour. The current status of, and open problems in, the computational modelling of the several aspects of human agents are discussed by the various authors contributing to this issue. Unfortunately, it has not been possible to include a paper on the computational modelling of the motor-behaviour of agents, i.e. on manipulation, locomotion, and navigation by humans. However, in this area the actual accomplishments are more in AI as engineering (i.e. robotics) rather than as science.

In their excellent and comprehensive tutorial survey of the architecture of intelligence, Chandrasekaran and Josephson (**page 366**) take as their working hypothesis the view that 'it is possible that intelligence can be explained or simulated without necessarily explaining and simulating other aspects of mind', specifically 'emotional states and subjective consciousness'. They refer to this as '*The Separability Hypothesis*'. After exhaustively surveying various computational models that have been proposed to account for the 'deliberative' and 'sub-deliberative' levels of cognitive behaviour, they come to the conclusion that 'in a biologically evolved object like the human brain ... a clear separation between levels of architecture and between hardware and software is impossible'. Their conclusion is that although we can potentially model each individual function of cognition, there may be no abstract Platonic engine which accounts for all and only cognitive, or all and only mental, behaviour. There may well be just various cognitive functions and various machines that can be used to explain these functions.'

As Chandrasekaran and Josephson discuss in their survey, so far as 'sub-deliberative' behaviour is concerned. *connectionism* (or neural network modelling) has been the most actively pursued alternative to the explicit symbol-manipulation techniques of conventional AI. Smolensky, one of the leading research workers in the connectionist camp, wrote a seminal paper in 1988 titled, 'On the proper treatment of connectionism'. His view was that connectionist models are neither models at the symbolic, conceptual-level, nor are they models at the neural-level. In other words, connectionism attempts neither a model of the mind nor that of the brain. Connectionist models are in-between the two in the modelling hierarchy. Smolensky referred to such models as 'sub-symbolic' or 'sub-conceptual' models.

Following the successes of connectionist models to account for some of the perceptual-level behaviour, some research workers thought that what are really needed are hybrid-models; connectionist models to account for the tacit knowledge underpinning perceptual-motor behaviour and mainstream symbol-manipulating AI models to account for higher-level cognition (e.g. problem-solving, deliberative thinking, and so on). However, a serious problem confronting the designer of a hybrid model is the task of abstracting symbolic encodings out of connectionism, i.e. coming to grips with the knowledge interface between connectionism and mainstream AI. This remains very much an unexplored territory.

In their latest paper, Smolensky, Legendre and Miyata (**page 381**) confront this problem and argue that the solution really lies in two-level interpretations of one and the same connectionist model — subsymbolic at the lower-level, and symbolic at the higher-level. They have tried out their methodology to tackle some open (and so far recalcitrant) issues in the grammar of French (construed as a natural language). The outcomes of their research work are discussed by them in their paper and they provide a rather extensive bibliography to substantiate their claims.

Natural language behaviour confronts us with some really deep problems. It is unclear whether formal grammars are really the central issues in modelling natural language behaviour. As Narasimhan discusses in his paper, natural language is used not only to describe the situational aspects of the world available to us through our (non-language) sensory modalities but natural language is also used to instruct and control behaviour. Language expressions making-up utterances play a dual role. They are 'things' (objects) and 'acts' simultaneously. As 'things' they are capable of forming the bricks for symbolic description. As 'acts' they constitute behavioural chunks. We exteriorize and deploy technology (i.e. scripts, pictures, computer programs, etc.) to freeze and make permanent such language descriptions for later contemplation at leisure. AI — whether mainstream or connectionist — is yet to come to grips with *literacy*, in the above sense, and its cultural manifestations.

From an evolutionary point of view, again, natural language behaviour exhibits anomalous traits. 'Having language' seems to be an all-or-none feature of human beings. (Ignoring vocabulary), we do not have any evidence (historically or at present) of language communities with partially developed language competence. Darwin says in his *The Origin of Species*, 'if it could be demonstrated that any complex organ existed which *could* not possibly have been formed by numerous, successive, slight modifications, my theory would absolutely break down.' And yet, we have no evidence of incremental developments associated with 'having language'. Either we must assume that, like dinosaurs, species with partially developed language competence have all vanished, or, what is perhaps more likely, the origins of the language modality should be sought in non-language modalities. In other words, what we take to be characterizing features of the language modality (e.g. productivity, compositional semantics, etc.) should actually underpin our behaviour in general and are not specific to the language modality. Some scientists believe that the origins of language behaviour should really be sought in our competence to devise gestural languages. This may very well be true and, if so, this fact should have architectural and procedural impacts on AI in rather fundamental ways.

Joshi, in his paper on 'natural language processing', describes (**page 393**) on-going work in interfacing language inputs to realize and control computer simulations of human acts. In other words, he discusses issues that arise in using language to specify and control in a realistic manner the motor-behaviour of computer-simulated human figures. Such issues directly bear on 'instructability' of humans in behaviourally relevant ways. In the first part of his paper, Joshi discusses the technology of natural language processing (NLP) as applied to very large corpora (consisting of hundreds of millions of sentences). His paper also discusses active research problems in the areas of formal grammars and parsers.

Language and vision play central roles in human beings considered as 'situated intelligences', i.e. intelligences interacting with the environments in which they find themselves located. Approximately half the cortex of human beings is devoted to coping with visually mediated behaviour. That we are primarily spatially-oriented creatures is demonstrated by a variety of features of our behaviour. Developmentally, 'spatial' perception and related motor behaviour are the earliest to mature. Spatial metaphors play a significantly large role in the way we analyse and refer to a variety of temporal and social relationships. Yet, natural language behaviour and vision are two of the poorly studied areas in traditional AI.

Zucker, in his paper (**page 407**) on vision, argues that computational modelling of vision has made little progress in traditional AI studies for two reasons: (1) vision has been sought to be studied divorced from visually mediated behaviour of a 'situated' intelligence; (2) approaches adopted to the modelling of early vision (i.e. low-level visual processing) were defective. Zucker discusses in his paper alternate approaches that he and his associates have been developing. He illustrates the successes of their methodology through a variety of computer-processed images. He also emphasizes that different visually mediated behaviour (e.g. object recognition versus navigation) need not be based on a single uniform processing of visually given information.

By now it is commonplace for critics of AI to point out that computers can only exhibit behaviour that is pre-specified by a programmer. Computers cannot be 'original' and cannot 'surprise' the person who programmed the computer in the first place. (Discounting the use of random number generators in the program) this could only be true of computers that function as isolated islands delinked completely from their environments. But it could not be true of a computer that interacts with its environment. For instance, trivially a programmer could include the instruction: 'Repeat the last conversational item keyed-in during the just-completed interaction.' Clearly, unless the programmer is a constant companion of the computer during its interaction with its environment, there is no way in which he would be able to predict what the computer would 'say'. Weizenbaum's Eliza program using a somewhat sophisticated version of this trick could engage in a realistic 'conversation' with its interacting partner.

Boden, in her contribution 'Creativity and Computers', considers (**page 419**) what it takes for a computer to be creative. In other words, she addresses the question: 'Is it possible to synthesize viable *computational models* of creative behaviour?' That this question can be tackled meaningfully is well-illustrated by the reproductions of paintings that adorn this special issue. These were all generated by an original drawing program called AARON written by Harold Cohen. The black-and-white line drawings produced by AARON were later manually coloured by Cohen. Boden, in addition, considers the issues that arise in the computational modeling of creativity in other domains such as music, story-writing, and so on.

Knowledge-based computing is, perhaps, the area where AI-technology has been convincingly deployed to tackle real-life problems. Expert-systems form the core of knowledge-based computing systems. The Japanese Government-sponsored Fifth Generation Computer Project had as its goal the design of specialized hardware and software to realize efficient knowledge-based computing in a variety of application areas. Chandrasekar and Anjaneyulu briefly describe (**page 391**) the Japanese Fifth Generation Computer Project and its accomplishments. Finally, the review by Chandrasekar of the book and video-tape (**page 436**) titled *The Age of Intelligent Machines* by Raymond Kurzweil complements nicely the contributions of the other authors to this special issue.

ACKNOWLEDGEMENTS. I must first thank Prof. S. Ramaseshan for inviting me to be the Guest Editor of this Special Issue of *Current Science* on AI. My next indebtedness is to the several authors for agreeing readily to contribute and for being prompt in delivering their papers. I must record my special thanks to Harold and Becky Cohen for graciously agreeing to the reproduction of the creations by AARON. Finally, it would have been practically impossible to produce the Special Issue without the generous e-mail and other computer facilities made available by the National Centre for Software Technology in Bombay. I thank Dr S. Ramani, Director of NCST, for this. I am personally indebted to R. Chandrasekar for logistic support in efficiently benefiting from all these infrastructural aids. I also thank the editorial staff of *Current Science* who have been most cooperative in this venture.

R. NARASIMHAN