FUTURE DIRECTIONS OF ARTIFICIAL INTELLIGENCE IN A RESOURCE-LIMITED ENVIRONMENT

by

R. S. Michalski
D. C. Littman

Chapter in the book, Future Directions in Artificial Intelligence, P.A. Flach and R.A. Meersman (Eds.), pp. 63-69, North-Holland, 1991.

NORTH-HOLLAND

FUTURE DIRECTIONS IN ARTIFICIAL INTELLIGENCE

P.A. FLACH
R.A. MEERSMAN
Editors

IFIP

FUTURE DIRECTIONS IN ARTIFICIAL INTELLIGENCE

IFIP TC12 Founding Workshop Collected Papers

Edited by

PETER A. FLACH
INFOLAB B 304
Tilburg University
Tilburg, The Netherlands

ROBERT A. MEERSMAN
INFOLAB B 304
Tilburg University
Tilburg, The Netherlands

1991

NORTH-HOLLAND
AMSTERDAM · NEW YORK · OXFORD · TOKYO

Table of Contents

List of Contributors xi

Editors' preface xiii

METHODOLOGY 1

The need for a formal education in Artificial Intelligence 3

Ranan B. Banerji

1. Experimental Computer Science? 4
2. The early days 4
3. Some symptoms 5
4. The way it hurts 6
5. A plea to the schools 8

Informatics as the scientific environment for Artificial Intelligence 11

Jozef Gruska

1. Introduction 12
2. AI and Computer Science — their ups and downs 12
3. Informatics as a new fundamental science and new methodology for science and technology 16
4. Implications for AI 20
5. Conclusions 23

Towards a reintegration of Artificial Intelligence research 25

John F. Sowa

1. Effects of specialization 26
2. Towards a reintegration 30

PARADIGMS 35

Paradigm shifts in Artificial Intelligence 37

Armand de Callataÿ

1. The changing environment for AI 38
   1.1 What was AI? 38
   1.2 AI and Computer Science 38
2. Paradigm shifts 39
   2.1 New definition of a "simple" computing instruction 39
   2.2 Processing directly in memory 39
   2.3 Computation on processors which are not universal 40
   2.4 Limitations of rule-based systems 41
   2.5 Reviving "system theory" 42
   2.6 Primitives and emergent properties 42
   2.7 Knowledge base of procedures 42
   2.8 Study of defective machines 43
   2.9 Learning by experience as in animals 44
3. Research related to brains 44
4. New computer architectures for AI 45
5. AI research on hybrid systems 46
6. A future direction for AI research 49
7. Conclusions. 51

Why today's computers don't learn the way people do 53

William J. Clancey

1. Introduction 54
2. Learning in AI programs 55
3. An alternative view 58
4. Two examples 58
5. Summary 61

Future directions of Artificial Intelligence in a resource-limited environment 63

Ryszard S. Michalski & David C. Littman

1. Introduction 64
2. An outline of future intelligent systems 64
   2.1 Functionality 64
   2.2 Inference capabilities 65
   2.3 Engineering considerations 65
3. Paradigms for AI research 65

4. Symbolic preeminence 67
5. Conclusions 68

Beyond the symbolic paradigm 71

Leon S. Sterling, Randall D. Beer & Hillel J. Chiel

1. Introduction 72
2. Coping with a changing, unpredictable environment 74
3. Intelligence as adaptive behavior 75
4. Lessons from biology 76
5. Computational ethology 78
6. Conclusions 79

TRENDS 81

Research trends in knowledge representation 83

Luigia Carlucci Aiello & Daniele Nardi

1. Premise 84
2. The problem 84
3. Hybrid reasoning 85
4. Reasoning with incomplete knowledge 87
5. Reasoning with contradictory knowledge 89
6. Reasoning about knowledge and reasoning 90

"Artificial" interfaces for knowledge acquisition: a futuristic scenario 93

R. Chandrasekar & S. Ramani

1. Introduction 94
2. Limitations in current AI systems 94
3. Prescription for the future 95
4. The structure of the ideal interface 96
   4.1 New modes of control and communication 96
   4.2 Intelligent interfaces 97
   4.3 The case for multi-modal communication 98
   4.4 New devices: a peek at what is possible 98
5. Lessons from child language research 99
   5.1 Language behaviour 99
   5.2 What Motherese offers: an aid to learning in children 100
   5.3 Motherese: where else can it be applied? 101
   5.4 The cognitive development of computer programs 101
6. An example: a natural language understanding system 102
7. Conclusions 103

Intelligent distributed and networking systems 105

Jacek Maitan

1. Current status 106
   1.1 Object-oriented model of distributed systems 106
   1.2 Implementation problems 107
2. Emerging needs 108
3. Scope of the problem 108
4. Research directions 109
   4.1 Nature of distributed problems 109
   4.2 Implementation and maintenance of distributed systems 110
   4.3 Theoretical models of computation and communication 111
5. Summary 112

Distributed Artificial Intelligence 113

Zhongzhi Shi

1. Introduction 114
2. The key issues 114
   2.1 Parallel distributed processing 115
   2.2 Knowledge representations 116
   2.3 Task decomposition 117
   2.4 Cooperative strategies 117
3. The principle techniques 118
   3.1 Coordinating via organizational structuring 118
   3.2 Contract network 119
   3.3 Task centralization 119
   3.4 Partial global planning 119
   3.5 Distributed knowledge base management system 119
4. Applications 120
5. Conclusions 120

PROSPECTS 121

On the future (and present) state of Artificial Intelligence 123

Yves Kodratoff

1. Existing features that will become better acknowledged 124
   1.1 Knowledge intensiveness 124
   1.2 Transparent box and explanations 124
2. Existing topics that will develop 125
   2.1 Knowledge intensive deduction 125
   2.2 Knowledge intensive induction 126

3. Future topics 126
 3.1 Explanations 126
 3.2 Multi-agent interactions 127
 3.3 Symbolic vs. numeric 127
 3.4 Analogy 127
4. Applications 128
 4.1 Application of AI to vision 128
 4.2 Application of expert systems by naive users 129
 4.3 Certification of expert systems by machine learning techniques 129
5. Conclusions 130
Acknowledgments 130

Artificial Intelligence needs its Eisenstein and Chaplin **131**

András Márkus & Elöd Knuth

1. The personification of software 132
 1.1 Individual features 132
 1.2 The program's lost identity and the need for archivation 133
2. The growing importance of transformations between representations 133
3. New integration and separation of cognitive activities 134
 3.1 Generating natural phenomena 134
 3.2 AI for and against the deterioration of traditional skills 134

AI technology, non-existent or extinct? **137**

Peter van Lith

The future of research in Artificial Intelligence **141**

Laurent Siklóssy

1. Introduction 142
2. The future of research 142
3. The future of AI 143
4. Research in AI 145
5. Future research in AI 145
6. Conclusions: more of the same! 146

EPILOGUE 149

The dialectics of Artificial Intelligence 151

Peter A. Flach

1. Introduction 152
2. What is Artificial Intelligence? 152
3. Is the mind a computer program? 153
4. Are sub-symbolic representations necessary? 155
   - 4.1 Scientific paradigms 156
   - 4.2 Symbolism vs. connectionism 157
5. Conclusions 158

References 161

Index 177

Future directions of Artificial Intelligence in a resource-limited environment

Ryszard S. Michalski
David C. Littman

Artificial intelligence research operates in a resource-limited environment, and therefore choices regarding the types of research to support are inevitable. In this context, we present an outline of future intelligent systems, and we discuss various research paradigms for guiding AI research in the future.

Future Directions in Artificial Intelligence
P.A. Flach & R.A. Meersman (Editors)
Elsevier Science Publishers B.V. (North-Holland)

1. Introduction

Artificial intelligence research, as any other human endeavor, operates in a resource-limited environment. It becomes increasingly necessary to make choices regarding the types of research to support and the directions in which to encourage young professionals to develop their interests and expertise. It is probably self-evident that these directions of AI depend upon the response to three general issues:

Issue 1: *What functionality should future intelligent systems have?*

Issue 2: *What inference capabilities are required to achieve the desired functionality?*

Issue 3: *What engineering considerations should be respected to achieve this functionality?*

2. An outline of future intelligent systems

2.1 Functionality

In our view, the capabilities of future intelligent systemshould include:

- Instructability, i.e., it should be relatively easy to communicate all kinds of knowledge to them, as well as to teach them all kinds of skills.
- The ability to explain the reasoning behind any conclusion, plan or task execution in human understandable terms and constructs, except in such situations where such an explanation is clearly unimportant. For example, in situations pertaining to human medical treatment, economical advice, or military defense actions, such an explanation would be necessary; for situations regarding robot manipulation or a delivery task, an explanation may not be important if the task was performed satisfactorily.
- The ability to interact with humans through their own media, e.g., texts, speech and images.

2.2 Inference capabilities

As to inference capabilities, future AI systems should be capable of performing the type of inference that humans would label plausible or commonsense. These systems should be capable of drawing likely conclusions from imperfect premises, performing meta-knowledge reasoning, creating descriptions or hypotheses at different levels of abstraction, evaluating and selecting different representational schemes, and constructing and evaluating goals and plans. All these capabilities would engage diverse kinds of knowledge.

To perform such types of inference, intelligent systems would have to ultimately possess various parts of technical as well as intellectual knowledge that humanity has accumulated during its existence. To this end, we believe that these systems would have to be able to acquire knowledge from human texts (at least partially by themselves), as well as from human speech.

2.3 Engineering considerations

As for engineering considerations, AI systems would have to be cost-effective, maintainable, reusable and generalizable. These features are neccessary in order to be economically viable and find practical applications.

3. Paradigms for AI research

In order to make good use of limited resources and funds, one must evaluate the implications of these three issues in selecting paradigms for AI research and development, and weigh the proper balance of effort.

At present, two general approaches to AI are predominant — the symbolic paradigm and the connectionist paradigm, and their relative importance is increasingly the subject of discussion. Much of this discussion results from the lack of a clear understanding of their respective capabilities. The symbolic paradigm, which emphasizes explicit, localized concept representation, and symbolic manipulation of these representations to derive inferences, is well established. It has had many successes, in particuar in building expert and advisory systems. It faces, however, the growing difficulty when building large-scale knowledge bases and implementing powerful multi-type inferences.

The fledgling connectionist paradigm, which focuses on distributed non-perspicuous knowledge representations, and their modification through changing weights of inter-unit connections, is becoming increasingly popular. The research in this paradigm is highly experimental, and avoids, at least at the present time, the issues of representing/manipulating large amounts of knowledge and/or diverse types of knowledge. Because it focuses on systems that are highly data-oriented and relatively knowledge-poor, it is not difficult for one researcher or a few researchers to make a series of experiments, and publish a research contribution in a short period of time. The latter factor seems to contribute to the attractiveness of research in this paradigm.

Research on either the connectionist paradigm or the symbolic paradigm, may primarily address one of two different yet interrelated goals: (*i*) a cognitive one, concerned with increasing our understanding of the functions of a human or an animal brain that are responsible for intelligent behavior; or (*ii*) an engineering one, concerned with developing systems that are able to perform some useful task, or serve as a useful tool for human society. While the cognitive goal is very important from the scientific viewpoint, the engineering one has a more direct economic impact In this paper, we will restrict ourselves to discussing these paradigms primarily from the engineering perspective.

As a step towards dealing with resource allocation, we have identified five candidate views for guiding future AI development. We also suggest which candidate we feel is most likely to succeed in this goal.

1. *Connectionist hegemony* holds that most, if not all, aspects of intelligent information processing can and should be performed in a substrate, as similar as possible to the low-level structure of the nervous system.

2. *Symbolic hegemony* holds that most, if not all, aspects of intelligent information processing can and should be performed using symbolic representations..

3. *Connectionist / symbolic equality* holds that connectionist and symbolic methods are equally necessary and that, a priori, there is no reason for preferring the connectionist or the symbolic approach in implementing any particular intelligent functionality.

4. *Connectionist preeminence* holds that, while symbolic methods may be useful to solve certain types of problems, such as performing numerical computations or some forms of reasoning, connectionist methods are more likely than symbolic methods to permit the implementation of the desired functions in future intelligent systems.

5. *Symbolic preeminence* holds that, while certain aspects of intelligent information processing might be best implemented in a connectionist system, such as low-level control functions or knowledge-limited signal processing, most of the functions responsible for intelligent behavior (such as context-dependent plausible reasoning, high-level abstraction, complex goal formulation, conceptual learning, and planning and selecting a representation) are most appropriately and cost-effectively performed symbolically.

In the remainder of this paper, we present some reasons that convinced us that symbolic preeminence is the best bet for guiding AI research in the future.

4. Symbolic preeminence

Symbolic processing systems can perform with ease a larger variety of complex information processing operations than can connectionist systems. In contrast to the connectionist approach, the symbolic processing paradigm provides many high-level, powerful and easily implementable functions. These functions include copying, modifying or erasing various parts of knowledge, operating on names representing components of knowledge, and explicitly performing "conceptual" inference operations such as generalizing, specializing or similizing selected knowledge segments. We define a "conceptual" inference operation as a knowledge transformation that is easy to understand, easy to explain in terms of human concepts, or for which one can easily build a mental model. For example, generalizing a decision rule by removing one of the conditions is a "conceptual" operation of generalization. On the other hand, generalizing a set of equations by replacing one numerical coefficient by another numerical coefficient, so that the modified set of equations describes a larger set of entities, is not a conceptual generalization.

The basic operation underlying the function of connectionist systems is parameter modification, i.e., a modification of the strength of the connections between various units. Through this modification, structural properties of knowledge represented by a network are modified indirectly. Such an operation is relatively easy to implement in biological systems, namely, through biochemical processes affecting synapses of neurons. Computer systems, however, offer information processing capabilities very different from those of biological systems. Consequently, it is likely that a possible future intelligent machine will be very different from its human counterpart. For example, it is easy for a computer to copy knowledge to (or from) another computer, or to erase, *on command*, undesirable segments of knowledge from its memory. Both of these operations are not possible for a human brain.

Focusing on a cost-effectiveness argument, let us suppose that the connectionist approach would overcome various important limitations of the current systems (e.g., how to represent complex structural descriptions with quantified variables), so that we could assume that symbolic and connectionist systems are computationally equivalent. This argument, however, does not resolve any issue. For example, the Turing Machine is theoretically equivalent to the contemporary computer, but from the pragmatic viewpoint, these two computational devices are vastly different. Because the symbolic paradigm is able to perform effectively more operations, this approach would likely lead to more cost-effective engineering solutions than the connectionist paradigm, given the assumption that these solutions would implement the capabilities stated earlier as being desirable in future intelligent systems.

Finally, let us use an analogy from the history of science. The success in the development of modern aircraft was primarily due to engineering experimentation and the development of general principles of flight independent of biological considerations, rather than the study and imitation of the biological properties that enable birds to fly.

In sum, we believe that there is a clear need to fully determine the capabilities and limitations of the connectionist and symbolic approaches. However, one should not be deceived by believing that there exists a "shortcut" to building an intelligent machine. With any new technology, the old dream comes back that somehow machines can self-organize themselves and learn on their own to do everything for us.

5. Conclusions

The following issues remain central for the development of AI:

- how to implement commonsense and plausible reasoning
- how to develop powerful learning capabilities that are able to take advantage of all kinds of prior knowledge and to explain to humans what was learned
- how to introduce and organize large amounts of human knowledge in machines, and how to update or improve that knowledge in the process of its normal use
- how to recognize objects and concepts from incomplete, variable and context-dependent cues.

These issues, which are basic items on the agenda of the symbolic approach, cannot be avoided if one wants to ultimately develop intelligent machines, and thus should be given a high priority in research in this area.