


Artificial Intelligence

A Modern Approach



Emilia Stones

 **Larsen & Keller**

Artificial Intelligence: A Modern Approach

Edited by
Emilia Stones

Artificial Intelligence: A Modern Approach

Edited by Emilia Stones

ISBN: 978-1-63549-032-9 (Hardback)

© 2017 Larsen & Keller

 Larsen & Keller

Published by Larsen and Keller Education,

5 Penn Plaza,

19th Floor,

New York, NY 10001, USA

Cataloging-in-Publication Data

Artificial intelligence : a modern approach / edited by Emilia Stones.

p. cm.

Includes bibliographical references and index.

ISBN 978-1-63549-032-9

1. Artificial intelligence. I. Stones, Emilia.

Q335 .A78 2017

006.3--dc23

This book contains information obtained from authentic and highly regarded sources. All chapters are published with permission under the Creative Commons Attribution Share Alike License or equivalent. A wide variety of references are listed. Permissions and sources are indicated; for detailed attributions, please refer to the permissions page. Reasonable efforts have been made to publish reliable data and information, but the authors, editors and publisher cannot assume any responsibility for the vailidity of all materials or the consequences of their use.

Trademark Notice: All trademarks used herein are the property of their respective owners. The use of any trademark in this text does not vest in the author or publisher any trademark ownership rights in such trademarks, nor does the use of such trademarks imply any affiliation with or endorsement of this book by such owners.

The publisher's policy is to use permanent paper from mills that operate a sustainable forestry policy. Furthermore, the publisher ensures that the text paper and cover boards used have met acceptable environmental accreditation standards.

Printed and bound in China.

For more information regarding Larsen and Keller Education and its products, please visit the publisher's website www.larsen-keller.com

Artificial Intelligence: A Modern Approach

Preface

Artificial intelligence is a growing field of information technology. It has transformed the world we live in. It has made the world more accessible, more social, more advanced and is developing the globe at a rapid speed. It has enabled human beings to study the minute and intricate concepts of science and has facilitated us to create better and more advanced machinery for medical and business purposes. This book contains the topics of utmost importance with regard to artificial intelligence. It aims to provide thorough insights into this subject and give detailed information about the various uses and methods applied in this area. As this field is emerging at a rapid pace, the contents of this text will help the readers understand the modern concepts and applications of the subject.

Given below is the chapter wise description of the book:

Chapter 1- The intelligence demonstrated by machines through programming is known as artificial intelligence. Artificial general intelligence is the intelligence that gadgets or machines possess by which they can accomplish tasks that humans being can also perform. This chapter will provide an integrated understanding of artificial intelligence.

Chapter 2- The principles and methods of artificial intelligence mentioned in the following section are swarm intelligence, cybernetics, symbolic artificial intelligence, superintelligence, computational intelligence etc. Swarm intelligence is the behavior of artificial systems and usually consists of a population of simple agents. These agents/boids communicate locally with one another or with their environment. The following section serves as a source to understand the basic methods of artificial intelligence.

Chapter 3- The philosophy of artificial intelligence questions artificial intelligence in a very philosophical manner. The questions raised mostly are, can a machine act intelligently? Is it possible for artificial intelligence to be the same as human intelligence? These questions including the ethical questions help in questioning how these artificially created intelligent beings can be used to harm humans and how they can also be useful. The topics elaborated in this section will help in gaining a better perspective of artificial intelligence.

Chapter 4- Adaptive system is a term that refers to the adaption of organisms, both artificial and natural organisms. Along with adaptive systems, the other applications of artificial intelligence are complex systems, expert systems and rule-based systems. This chapter is a compilation of the various branches of artificial intelligence and its applications that form an integral part of the broader subject matter.

Chapter 5- Artificial intelligence is used in many spheres of life; some of the various applications of artificial intelligence are robotics, virtual reality, game theory, concept mining, data mining and inference engine. The topics discussed in the section are of great importance to broaden the existing knowledge on artificial intelligence.

Chapter 6- Natural language processing is an important field of artificial intelligence. It is majorly concerned with the interactions between humans and computers. Along with natural language processing, this section explains to the reader other allied fields of artificial intelligence as well. Some of the allied fields explained are knowledge representation and reasoning, computational creativity and distributed artificial intelligence. This section has been carefully written to provide an easy understanding of the allied fields of artificial intelligence.

Chapter 7- The history of artificial intelligence is very important in order to understand the progress artificial intelligence has made over a period of years. The aspects of artificial intelligence elucidated in the chapter are machine learning, deep learning, AI effect and AI winter. In order to completely understand artificial intelligence, it is necessary to understand the evolution of artificial intelligence.

Indeed, my job was extremely crucial and challenging as I had to ensure that every chapter is informative and structured in a student-friendly manner. I am thankful for the support provided by my family and colleagues during the completion of this book.

Editor

Table of Contents

Preface	VII
Chapter 1 Introduction to Artificial Intelligence	1
a. Artificial Intelligence	1
b. Artificial General Intelligence	23
c. Weak AI	32
Chapter 2 Principles and Methods of Artificial Intelligence	34
a. Swarm Intelligence	34
b. Cybernetics	38
c. Symbolic Artificial Intelligence	49
d. Superintelligence	50
e. Computational Intelligence	54
f. Schismogenesis	58
g. Nouvelle AI	61
h. Diagnosis (Artificial Intelligence)	63
i. Collective Intelligence	65
j. Action Selection	79
Chapter 3 Philosophy and Ethics of Artificial Intelligence	85
a. Philosophy of Artificial Intelligence	85
b. Ethics of Artificial Intelligence	96
c. Artificial Psychology	101
d. Artificial Intelligence and Law	102
e. Friendly Artificial Intelligence	105
Chapter 4 Artificial Intelligence Systems and its Applications	109
a. Adaptive System	109
b. Complex Systems	111
c. Expert System	121
d. Rule-based System	128
Chapter 5 Various Applications of Artificial Intelligence	131
a. Applications of Artificial Intelligence	131
b. Robotics	135
c. Virtual Reality	158
d. Optical Character Recognition	178
e. Game Theory	184
f. Concept Mining	199
g. Data Mining	200
h. Inference Engine	210

Chapter 6	Allied Fields of Artificial Intelligence	213
	a. Natural Language Processing	213
	b. Knowledge Representation and Reasoning	222
	c. Computational Creativity	230
	d. Artificial Intelligence (Video Games)	241
	e. Distributed Artificial Intelligence	247
Chapter 7	Evolution of Artificial Intelligence	252
	a. History of Artificial Intelligence	252
	b. Machine Learning	268
	c. Deep Learning	279
	d. AI Effect	305
	e. AI Winter	307

Permissions

Index

Introduction to Artificial Intelligence

The intelligence demonstrated by machines through programming is known as artificial intelligence. Artificial general intelligence is the intelligence that gadgets or machines possess by which they can accomplish tasks that humans being can also perform. This chapter will provide an integrated understanding of artificial intelligence.

Artificial Intelligence

Artificial intelligence (AI) is intelligence exhibited by machines. In computer science, an ideal “intelligent” machine is a flexible rational agent that perceives its environment and takes actions that maximize its chance of success at some goal. Colloquially, the term “artificial intelligence” is applied when a machine mimics “cognitive” functions that humans associate with other human minds, such as “learning” and “problem solving”. As machines become increasingly capable, facilities once thought to require intelligence are removed from the definition. For example, optical character recognition is no longer perceived as an exemplar of “artificial intelligence” having become a routine technology. Capabilities currently classified as AI include successfully understanding human speech, competing at a high level in strategic game systems (such as Chess and Go), self-driving cars, and interpreting complex data.

AI research is divided into subfields that focus on specific problems or on specific approaches or on the use of a particular tool or towards satisfying particular applications.

The central problems (or goals) of AI research include reasoning, knowledge, planning, learning, natural language processing (communication), perception and the ability to move and manipulate objects. General intelligence is among the field’s long-term goals. Approaches include statistical methods, computational intelligence, soft computing (e.g. machine learning), and traditional symbolic AI. Many tools are used in AI, including versions of search and mathematical optimization, logic, methods based on probability and economics. The AI field draws upon computer science, mathematics, psychology, linguistics, philosophy, neuroscience and artificial psychology.

The field was founded on the claim that human intelligence “can be so precisely described that a machine can be made to simulate it.” This raises philosophical arguments about the nature of the mind and the ethics of creating artificial beings endowed with human-like intelligence, issues which have been explored by myth, fiction and philosophy since antiquity. Attempts to create artificial intelligence have experienced many setbacks, including the ALPAC report of 1966, the abandonment of perceptrons in 1970, the Lighthill Report of 1973 and the collapse of the Lisp machine market in 1987. In the twenty-first century AI techniques became an essential part of the technology industry, helping to solve many challenging problems in computer science.

History

While the concept of thought capable artificial beings appeared as storytelling devices in antiquity, the idea of actually trying to build a machine to perform useful reasoning may have begun with Ramon Llull (c. 1300 CE). With his *Calculus ratiocinator*, Gottfried Leibniz extended the concept of the calculating machine (Wilhelm Schickard engineered the first one around 1623), intending to perform operations on concepts rather than numbers. Since the 19th century, artificial beings are common in fiction, as in Mary Shelley's *Frankenstein* or Karel Čapek's *R.U.R. (Rossum's Universal Robots)*.

The study of mechanical or “formal” reasoning began with philosophers and mathematicians in antiquity. In the 19th century, George Boole refined those ideas into propositional logic and Gottlob Frege developed a notational system for mechanical reasoning (a “*predicate calculus*”). Around the 1940s, Alan Turing's theory of computation suggested that a machine, by shuffling symbols as simple as “0” and “1”, could simulate any conceivable act of mathematical deduction. This insight, that digital computers can simulate any process of formal reasoning, is known as the Church–Turing thesis. Along with concurrent discoveries in neurology, information theory and cybernetics, this led researchers to consider the possibility of building an electronic brain. The first work that is now generally recognized as AI was McCulloch and Pitts' 1943 formal design for Turing-complete “artificial neurons”.

The field of AI research was founded at a conference at Dartmouth College in 1956. The attendees, including John McCarthy, Marvin Minsky, Allen Newell, Arthur Samuel and Herbert Simon, became the leaders of AI research. They and their students wrote programs that were, to most people, simply astonishing: computers were winning at checkers, solving word problems in algebra, proving logical theorems and speaking English. By the middle of the 1960s, research in the U.S. was heavily funded by the Department of Defense and laboratories had been established around the world. AI's founders were optimistic about the future: Herbert Simon predicted that “machines will be capable, within twenty years, of doing any work a man can do”. Marvin Minsky agreed, writing that “within a generation ... the problem of creating ‘artificial intelligence’ will substantially be solved”.

They failed to recognize the difficulty of some of the remaining tasks. Progress slowed and in 1974, in response to the criticism of Sir James Lighthill and ongoing pressure from the US Congress to fund more productive projects, both the U.S. and British governments cut off exploratory research in AI. The next few years would later be called an “AI winter”, a period when funding for AI projects was hard to find.

In the early 1980s, AI research was revived by the commercial success of expert systems, a form of AI program that simulated the knowledge and analytical skills of human experts. By 1985 the market for AI had reached over a billion dollars. At the same time, Japan's fifth generation computer project inspired the U.S and British governments to restore funding for academic research. However, beginning with the collapse of the Lisp Machine market in 1987, AI once again fell into disrepute, and a second, longer-lasting hiatus began.

In the late 1990s and early 21st century, AI began to be used for logistics, data mining, medical diagnosis and other areas. The success was due to increasing computational power (see Moore's

law), greater emphasis on solving specific problems, new ties between AI and other fields and a commitment by researchers to mathematical methods and scientific standards. Deep Blue became the first computer chess-playing system to beat a reigning world chess champion, Garry Kasparov on 11 May 1997.

Advanced statistical techniques (loosely known as deep learning), access to large amounts of data and faster computers enabled advances in machine learning and perception. By the mid 2010s, machine learning applications were used throughout the world. In a *Jeopardy!* quiz show exhibition match, IBM's question answering system, Watson, defeated the two greatest Jeopardy champions, Brad Rutter and Ken Jennings, by a significant margin. The Kinect, which provides a 3D body-motion interface for the Xbox 360 and the Xbox One use algorithms that emerged from lengthy AI research as do intelligent personal assistants in smartphones. In March 2016, AlphaGo won 4 out of 5 games of Go in a match with Go champion Lee Sedol, becoming the first computer Go-playing system to beat a professional Go player without handicaps.

Research

Goals

The general problem of simulating (or creating) intelligence has been broken down into sub-problems. These consist of particular traits or capabilities that researchers expect an intelligent system to display. The traits described below have received the most attention.

Deduction, Reasoning, Problem Solving

Early researchers developed algorithms that imitated step-by-step reasoning that humans use when they solve puzzles or make logical deductions (reason). By the late 1980s and 1990s, AI research had developed methods for dealing with uncertain or incomplete information, employing concepts from probability and economics.


For difficult problems, algorithms can require enormous computational resources—most experience a “combinatorial explosion”: the amount of memory or computer time required becomes astronomical for problems of a certain size. The search for more efficient problem-solving algorithms is a high priority.

Human beings ordinarily use fast, intuitive judgments rather than step-by-step deduction that early AI research was able to model. AI has progressed using “sub-symbolic” problem solving: embodied agent approaches emphasize the importance of sensorimotor skills to higher reasoning; neural net research attempts to simulate the structures inside the brain that give rise to this skill; statistical approaches to AI mimic the human ability to guess.

Knowledge Representation

Knowledge representation and knowledge engineering are central to AI research. Many of the problems machines are expected to solve will require extensive knowledge about the world. Among the things that AI needs to represent are: objects, properties, categories and relations between objects; situations, events, states and time; causes and effects; knowledge about knowledge (what we know about what other people know); and many other, less well researched domains. A representation

of “what exists” is an ontology: the set of objects, relations, concepts and so on that the machine knows about. The most general are called upper ontologies, which attempt to provide a foundation for all other knowledge.



An ontology represents knowledge as a set of concepts within a domain and the relationships between those concepts.

Knowledge representation and knowledge engineering are central to AI research. Many of the problems machines are expected to solve will require extensive knowledge about the world. Among the things that AI needs to represent are: objects, properties, categories and relations between objects; situations, events, states and time; causes and effects; knowledge about knowledge (what we know about what other people know); and many other, less well researched domains. A representation of “what exists” is an ontology: the set of objects, relations, concepts and so on that the machine knows about. The most general are called upper ontologies, which attempt to provide a foundation for all other knowledge.

Among the most difficult problems in knowledge representation are:

Default reasoning and the qualification problem

Many of the things people know take the form of “working assumptions.” For example, if a bird comes up in conversation, people typically picture an animal that is fist sized, sings, and flies. None of these things are true about all birds. John McCarthy identified this problem in 1969 as the qualification problem: for any commonsense rule that AI researchers care to represent, there tend to be a huge number of exceptions. Almost nothing is simply true or false in the way that abstract logic requires. AI research has explored a number of solutions to this problem.

The breadth of commonsense knowledge

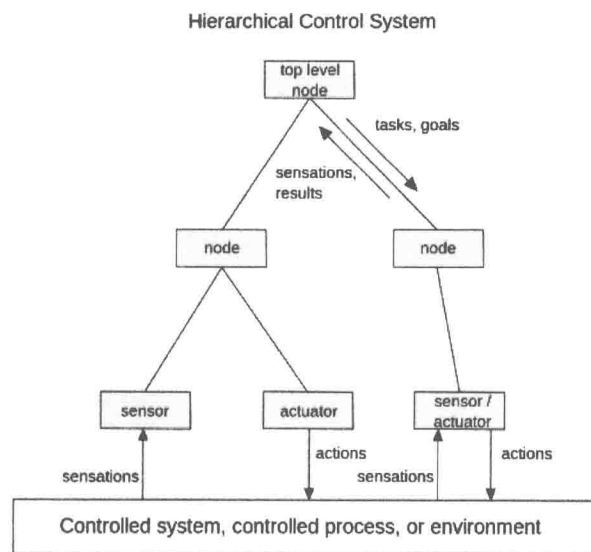
The number of atomic facts that the average person knows is astronomical. Research projects that attempt to build a complete knowledge base of commonsense knowledge (e.g., Cyc) require enormous amounts of laborious ontological engineering—they must be built, by hand, one complicated concept at a time. A major goal is to have the computer understand enough concepts to be able to learn by reading from sources like the Internet, and thus be able to add to its own ontology.

The subsymbolic form of some commonsense knowledge

Much of what people know is not represented as “facts” or “statements” that they could express verbally. For example, a chess master will avoid a particular chess position because it “feels too exposed” or an art critic can take one look at a statue and instantly realize that it is a fake. These are intuitions or tendencies that are represented in the brain non-consciously and sub-symbolically. Knowledge like this informs, supports and provides a context for symbolic, conscious knowledge. As with the related problem of sub-symbolic reasoning, it is hoped that situated AI, computational intelligence, or statistical AI will provide ways to represent this kind of knowledge.

Planning

Intelligent agents must be able to set goals and achieve them. They need a way to visualize the future (they must have a representation of the state of the world and be able to make predictions about how their actions will change it) and be able to make choices that maximize the utility (or “value”) of the available choices.



A hierarchical control system is a form of control system in which a set of devices and governing software is arranged in a hierarchy.

In classical planning problems, the agent can assume that it is the only thing acting on the world and it can be certain what the consequences of its actions may be. However, if the agent is not the only actor, it must periodically ascertain whether the world matches its predictions and it must change its plan as this becomes necessary, requiring the agent to reason under uncertainty.

Multi-agent planning uses the cooperation and competition of many agents to achieve a given goal. Emergent behavior such as this is used by evolutionary algorithms and swarm intelligence.

Learning


Machine learning is the study of computer algorithms that improve automatically through experience and has been central to AI research since the field's inception.

Unsupervised learning is the ability to find patterns in a stream of input. Supervised learning includes both classification and numerical regression. Classification is used to determine what category something belongs in, after seeing a number of examples of things from several categories. Regression is the attempt to produce a function that describes the relationship between inputs and outputs and predicts how the outputs should change as the inputs change. In reinforcement learning the agent is rewarded for good responses and punished for bad ones. The agent uses this sequence of rewards and punishments to form a strategy for operating in its problem space. These three types of learning can be analyzed in terms of decision theory, using concepts like utility. The mathematical analysis of machine learning algorithms and their performance is a branch of theoretical computer science known as computational learning theory.

Within developmental robotics, developmental learning approaches were elaborated for lifelong cumulative acquisition of repertoires of novel skills by a robot, through autonomous self-exploration and social interaction with human teachers, and using guidance mechanisms such as active learning, maturation, motor synergies, and imitation.

Natural Language Processing (Communication)

Natural language processing gives machines the ability to read and understand the languages that humans speak. A sufficiently powerful natural language processing system would enable natural language user interfaces and the acquisition of knowledge directly from human-written sources, such as newswire texts. Some straightforward applications of natural language processing include information retrieval, text mining, question answering and machine translation.



A parse tree represents the syntactic structure of a sentence according to some formal grammar.

A common method of processing and extracting meaning from natural language is through semantic indexing. Increases in processing speeds and the drop in the cost of data storage makes indexing large volumes of abstractions of the user's input much more efficient.

Perception

Machine perception is the ability to use input from sensors (such as cameras, microphones, tactile sensors, sonar and others more exotic) to deduce aspects of the world. Computer vision is the ability to analyze visual input. A few selected subproblems are speech recognition, facial recognition and object recognition.

Motion and Manipulation

The field of robotics is closely related to AI. Intelligence is required for robots to be able to handle such tasks as object manipulation and navigation, with sub-problems of localization (knowing


where you are, or finding out where other things are), mapping (learning what is around you, building a map of the environment), and motion planning (figuring out how to get there) or path planning (going from one point in space to another point, which may involve compliant motion – where the robot moves while maintaining physical contact with an object).

Long-term Goals

Among the long-term goals in the research pertaining to artificial intelligence are: (1) Social intelligence, (2) Creativity, and (3) General intelligence.

Social Intelligence

Affective computing is the study and development of systems and devices that can recognize, interpret, process, and simulate human affects. It is an interdisciplinary field spanning computer sciences, psychology, and cognitive science. While the origins of the field may be traced as far back as to early philosophical inquiries into emotion, the more modern branch of computer science originated with Rosalind Picard's 1995 paper on affective computing. A motivation for the research is the ability to simulate empathy. The machine should interpret the emotional state of humans and adapt its behaviour to them, giving an appropriate response for those emotions.



Kismet, a robot with rudimentary social skills

Emotion and social skills play two roles for an intelligent agent. First, it must be able to predict the actions of others, by understanding their motives and emotional states. (This involves elements of game theory, decision theory, as well as the ability to model human emotions and the perceptual skills to detect emotions.) Also, in an effort to facilitate human-computer interaction, an intelligent machine might want to be able to *display* emotions—even if it does not actually experience them itself—in order to appear sensitive to the emotional dynamics of human interaction.

Creativity

A sub-field of AI addresses creativity both theoretically (from a philosophical and psychological perspective) and practically (via specific implementations of systems that generate outputs that can be considered creative, or systems that identify and assess creativity). Related areas of computational research are Artificial intuition and Artificial thinking.

General Intelligence

Many researchers think that their work will eventually be incorporated into a machine with *general* intelligence (known as strong AI), combining all the skills above and exceeding human abilities at most or all of them. A few believe that anthropomorphic features like artificial consciousness or an artificial brain may be required for such a project.

Many of the problems above may require general intelligence to be considered solved. For example, even a straightforward, specific task like machine translation requires that the machine read and write in both languages (NLP), follow the author's argument (reason), know what is being talked about (knowledge), and faithfully reproduce the author's intention (social intelligence). A problem like machine translation is considered "AI-complete". In order to solve this particular problem, one must solve all the problems.

Approaches

There is no established unifying theory or paradigm that guides AI research. Researchers disagree about many issues. A few of the most long standing questions that have remained unanswered are these: should artificial intelligence simulate natural intelligence by studying psychology or neurology? Or is human biology as irrelevant to AI research as bird biology is to aeronautical engineering? Can intelligent behavior be described using simple, elegant principles (such as logic or optimization)? Or does it necessarily require solving a large number of completely unrelated problems? Can intelligence be reproduced using high-level symbols, similar to words and ideas? Or does it require "sub-symbolic" processing? John Haugeland, who coined the term GOFAI (Good Old-Fashioned Artificial Intelligence), also proposed that AI should more properly be referred to as synthetic intelligence, a term which has since been adopted by some non-GOFAI researchers.

Cybernetics and Brain Simulation

In the 1940s and 1950s, a number of researchers explored the connection between neurology, information theory, and cybernetics. Some of them built machines that used electronic networks to exhibit rudimentary intelligence, such as W. Grey Walter's turtles and the Johns Hopkins Beast. Many of these researchers gathered for meetings of the Teleological Society at Princeton University and the Ratio Club in England. By 1960, this approach was largely abandoned, although elements of it would be revived in the 1980s.

Symbolic

When access to digital computers became possible in the middle 1950s, AI research began to explore the possibility that human intelligence could be reduced to symbol manipulation. The research was centered in three institutions: Carnegie Mellon University, Stanford and MIT, and each one developed its own style of research. John Haugeland named these approaches to AI "good old fashioned AI" or "GOFAI". During the 1960s, symbolic approaches had achieved great success at simulating high-level thinking in small demonstration programs. Approaches based on cybernetics or neural networks were abandoned or pushed into the background. Researchers in the 1960s and the 1970s were convinced that symbolic approaches would eventually succeed in creating a machine with artificial general intelligence and considered this the goal of their field.