

Artificial Intelligence and Machine Learning for Business for Non-Engineers

CRC Press
Taylor & Francis Group
6000 Broken Sound Parkway NW, Suite 300
Boca Raton, FL 33487-2742

© 2020 by Taylor & Francis Group, LLC
CRC Press is an imprint of Taylor & Francis Group, an Informa business

No claim to original U.S. Government works

Printed on acid-free paper

International Standard Book Number-13: 978-0-367-36574-5 (Paperback)

This book contains information obtained from authentic and highly regarded sources. Reasonable efforts have been made to publish reliable data and information, but the author and publisher cannot assume responsibility for the validity of all materials or the consequences of their use. The authors and publishers have attempted to trace the copyright holders of all material reproduced in this publication and apologize to copyright holders if permission to publish in this form has not been obtained. If any copyright material has not been acknowledged please write and let us know so we may rectify in any future reprint.

Except as permitted under U.S. Copyright Law, no part of this book may be reprinted, reproduced, transmitted, or utilized in any form by any electronic, mechanical, or other means, now known or hereafter invented, including photocopying, microfilming, and recording, or in any information storage or retrieval system, without written permission from the publishers.

For permission to photocopy or use material electronically from this work, please access www.copyright.com (<http://www.copyright.com/>) or contact the Copyright Clearance Center, Inc. (CCC), 222 Rosewood Drive, Danvers, MA 01923, 978-750-8400. CCC is a not-for-profit organization that provides licenses and registration for a variety of users. For organizations that have been granted a photocopy license by the CCC, a separate system of payment has been arranged.

Trademark Notice: Product or corporate names may be trademarks or registered trademarks, and are used only for identification and explanation without intent to infringe.

Visit the Taylor & Francis Web site at
<http://www.taylorandfrancis.com>

and the CRC Press Web site at
<http://www.crcpress.com>

Contents

[Preface](#)

[Acknowledgment](#)

[Editors](#)

[Contributors](#)

[1 Introduction to Artificial Intelligence and Machine Learning](#)

FRANK M. GROOM

[2 The Basic Elements of Artificial Intelligence](#)

FRANK M. GROOM

[3 Artificial Intelligence and the Cloud](#)

RYAN SUTER

[4 Artificial Intelligence Security Threats](#)

RYAN SUTER

[5 Artificial Intelligence and IT Management](#)

DAVE PICKENS

[6 Artificial Intelligence in Marketing](#)

JOE CIUFFO

[7 Artificial Intelligence in Cloud Marketing](#)

LAUREN DONAHUE AND FATEMEH HAJIZADEH

[8 Artificial Intelligence in Healthcare](#)

[RASHIDA PEETE, KASIA MAJOWSKI, LAURA LAUER,
AND ANNIE JAY](#)

[9 Artificial Intelligence and Medical Devices](#)

DUSHYANTH MARATT

10 Artificial Intelligence and Public Policy

SOPHIE GUETZKO

11 Artificial Intelligence in Education

MATTHEW LIEVERTZ

Index

Preface

Technology and machinery have been ever-changing subjects since the first shavings were taken off of a rock to create a sharp-enough tool to kill and skin primitive food supplies. This eventually paved the way for simple machines and technology to ease the daily lives of everyone – with each new advance impacting the way work was done previously. Fast forward to the future and technology is now capable of processing massive amounts of data in a split second, analyzing billions of bits of information to tug out the digital needle in a haystack. All of this leads the world into a new realm that seems to be on the tip of every company and innovator’s spear as they gear up to attack a new economy: artificial intelligence and learning machines.

This book will serve as an introduction to the field and technology that make up artificial intelligence and learning machines (AI and LM). It will delve deeper into specific fields that are optimized by the use and integration of AI and LM into their daily used systems in a non-technical manner to allow for the reader to gain knowledge of the emerging systems without the need to learn the algorithms. Some of these fields will include security, marketing, education, healthcare, IT management, and medical devices. This book will also inform individuals without an academic background in technology further insight into the often-daunting world of artificial intelligence. Knowledge gained in this book will support a knowledge base when discussing the world of AI and LM and help decision makers understand the critical components of the technologies. Further knowledge will be required to explore and achieve a greater breadth of understanding in the world of AI and LM; these tools will need to be part of everyone’s lexicon in the near future as they will be part of every new technology.

Acknowledgment

The authors would like to thank Derick Brady and Skylar McArthur for their management, research, and patience in putting this edited book together. Without their help, this book would not have been possible.

Editors



Frank M. Groom is a Professor of Information and Communication Sciences at the Center for Information and Communication Sciences at Ball State University. He has conducted research into high-bandwidth networking and the storage and transmission of multimedia objects. He has most recently concentrated his research into Multi-Protocol Label Switching (MPLS)-driven fiber networks, intelligent agents, network-based data deployment, cloud computing, big data storage and analysis, network security, and artificial intelligence and machine learning. Most recently, he has

delivered a network and data security class and his 2016 textbook of the same name as well as an Artificial Intelligence and Machine Learning class. Both have been filmed for online delivery. He has conducted a number of national research projects using surveys, focus groups, personal interviews, and student research culminating in two of his earlier published books. Furthermore, he has conducted many specialized statistical research studies for AT&T, McDonalds Corp., and Nth Dimension Software.

Furthermore, Dr. Groom annually conducts a graduate Research Methods course for Ball State graduate students where he teaches many of the methods he has employed in his own research conducted both in industry and at the university level studying both big data problems as well as smaller situations.

Dr. Groom is currently examining the detailed infrastructure and operations of Cloud Data Centers. He is visiting a number of such centers, examining the process by which they distribute processing and storage across racks and arrays of inexpensive components. He furthermore is investigating the approaches Cloud Centers take to elevate control and operations to minimally managed control centers by means of software-defined approaches. He is now extending that investigation into how carriers such as AT&T and Verizon are transforming the control and components of their national networks as software-defined structures with provisioning of virtual network slices of services from customers, similar to the cloud computing approach. Dr. Groom has presented networking and data processing courses to major American corporations, among whom are PricewaterhouseCoopers; IBM; AT&T and its various units including Bell Labs; Motorola; Digital Equipment Corp. (now HP); Unisys; Ford Motor; Hillenbrand Industries; and McDonalds. IBM sponsored him to present courses on chip engineering and wireless networking to Motorola's engineering teams in Hong Kong and Macao in 1993. AT&T has twice sponsored him to present advanced data processing and networking courses to the graduate students and faculty of Beijing University of Posts and Telecommunications (BUPT) and the People's Republic of China Government Office of Telecommunications. He was honored with being the plenary speaker and having two of his papers presented at the Plenary Session of the 1996 International Conference on Information Infrastructure (ICII'96) in Beijing, China, and a paper presented as the Plenary Session for the World Broadband 2000 conference in Tokyo. Furthermore, in 1996, 1998, and 2006, he presented papers on ATM networking, multimedia, and VoIP at the leading French Graduate School of Telecommunications (Ecole Nationale Supérieure des Telecommunications - ENST). In addition to publishing over 120 technical papers concerning networking, systems design, corporate re-engineering, and object-oriented storage, he has published 12 books and contributed chapters in many other books.

Dr. Groom has a Ph.D. in Management Information Systems from the University of Wisconsin-Milwaukee, all but dissertation (ABD) in Statistics from NYU, MBA from the University of Hartford, and AB in Mathematics from the College of the Holy Cross. He was Division Manager in charge of the Information Systems and Data Processing Division of Wisconsin Bell and is the retired Senior Director of Information Systems for Ameritech (now once again part of AT&T).



Stephan S. Jones joined the Center for Information and Communication Sciences faculty in August 1998. He came to Ball State University (BSU) from completing his doctoral studies at Bowling Green State University where he served as the Dean of Continuing Education, developing a distance-learning program for the College of Technology's undergraduate Technology Education program. He was instrumental in bringing the new program on board because of his technical background and extensive research in the distance-learning field.

Prior to coming to higher education, Dr. Jones spent over sixteen and a half years in the communication technology industry. He owned his own teleconnect, providing high-end commercial voice and data networks to a broad range of end users. He provided all the engineering and technical support for his organization that grew to over twenty employees and two and a half million dollars per year revenue. Selling his portion of the organization in December 1994, he worked briefly for Panasonic Communications and Systems Company as a District Sales Manager providing application engineering and product support to distributors in a five-state area prior to starting doctoral studies.

Since coming to BSU, Dr. Jones has been engaged in development efforts to provide up-to-date research opportunities outside of the classroom on areas associated with

broadband delivery systems, unified communications, healthcare information technologies, cloud infrastructure, and cybersecurity, and has appeared as a presenter at international, national, and regional technical and economic development gatherings. He has been the Co-Director of the Center's Applied Research Institute, Director of the Australia Centre, and actively involved in the University's community. He has written/edited eighteen books, authored numerous book chapters, and presented at conferences and seminars delivering content related to simplifying the complexity of information and communication technologies. He served ten years as the Director of the largest graduate STEM program at Ball State University, being charged with external funding development, student career development and placement, pursuit of new curriculum ideas, graduate student recruiting, and the out-of-classroom learning experience of the Student Social Learning Program.

Contributors

Joe Ciuffo

Product Marketing

Genesys

San Francisco, California

Lauren Donahue

Solution Architect

Lev+

Indianapolis, Indiana

Sophie Guetzko

Infrastructure Consultant

Accenture

Chicago, Illinois

Fatemeh Hajizadeh

Solution Architect

Lev+

Indianapolis, Indiana

Annie Jay

Healthcare Consulting Practice

Impact Advisors

Naperville, Illinois

Laura Lauer

Healthcare Consulting Practice

Impact Advisors

Naperville, Illinois

Matthew Lievertz

Software Development

Zylo

Indianapolis, Indiana

Kasia Majowski

Healthcare Consulting Practice

Impact Advisors

Naperville, Illinois

Dushyanth Maratt

Sales

ITC Infotech

Ellicott City, Maryland

Rashida Peete

Healthcare Consulting Practice

Impact Advisors

Naperville, Illinois

Dave Pickens

Infrastructure and Architecture Analysis

Pace Harmon

Indianapolis, Indiana

Ryan Suter

Business Analyst

enVista

Carmel, Indiana

Chapter 1

Introduction to Artificial Intelligence and Machine Learning

Frank M. Groom

Ball State University

Contents

References

As the promise of Artificial Intelligence (AI) becomes the hot topic of the information technology world, Silicon Valley companies are allocating sizable resources toward developing and trialing their various products. Meanwhile giants such as Google are buying up promising AI companies such as AlphaGo which was created by the London-based company, DeepMind. AlphaGo then proceeded to beat the Go world champion Ke Jie (Russell, 2017)¹. Meanwhile Ford, Tesla, and others are in the late development process of producing an automobile that drives itself.

These newsworthy incidents have heightened the interest of technocrats while simultaneously alerting the doomsayers. Apparently both are correct. AI and its partner, Machine Learning, are holding out enormous potential in all areas of business, communication, health, transportation, warehousing and delivery, on-line wholesale and retail, manufacturing, assembly, military, and even the last bastion government services. The benefits of automated decision-making and precision handling of objects have been gradually altering many areas of business from Amazon's Fulfillment Centers to Mercedes, BMW, and Tesla's virtually completely automated car assembly factories (McKinsey, 2017, 2018).

Although many believe there are limits to where AI can penetrate business operations, it increasingly appears that no job is safe from automation assistance and in some cases complete replacement (Bomey, 2017). In the 1990s, American businesses began the move of assembly and manufacturing work to the countries of Southeast Asia where labor rates were no more than 1%–3% of American rates, and as Labor Unions lost power, jobs were moved to the US Southern states and then on to Mexico and Southeast Asia. However, labor rates in those new assembly and manufacturing countries quickly began their own rise. Now China seeks to be the dominant player in offering Artificial Intelligent products to the world and is beginning to install such software in their own factories.

Initially it was thought that only the routine daily repetitive activities of white-collar work could be automated. But now with the demonstrated decision-making capabilities of IBM's Deep Blue and Watson, and Google's AlphaGo, there is no reason to expect decision-making jobs are the exception. Even the areas that seemed impenetrable, such as writing and journalism, are being invaded. As *The New York Times* has noted, AI programs are widely being employed to sift through millions of items for financial news postings daily and summarizing them and even subsequently writing standard articles for direct posting to various news outlets and newspapers.

So where will the new jobs arise. Obviously the initial expectation is that there will be a dramatic demand for program coders, particularly those with the interest, intelligence, and coding skill to address AI and the enhancement of its capabilities by Machine Learning. The training of Machine Learning through thousands of actual examples augmented by continued learning from the Machine Learning Module's actual live experiences has appealed to many potential programming hopefuls. Particularly attracted are the many game players and those who have devoted themselves to writing, sharing, and publishing their own games or variations on existing games. And the plentiful number of coding workshops, coding camps, and self-help on-line tutorials have set the stage for additional numbers who have interest to enter the programming community. Some will join the high-tech development community; some will become entrepreneurs and develop their own AI products. Others will join the business community as purchasers, analysts, implements, testers, and systems managers of the AI product line (Berbazzani, 2018).

However, there will be some limitations. With the emergence of Microsoft's DeepCoder which searches through stacks of programmed routines and assembles programs from these libraries of code, Matej Balog at the University of Cambridge, along with Microsoft Research's Alexander L. Gaunt, Marc Brockschmidt, Sebastian Nowozin, and Daniel Tarlow, has begun the inevitable process of building AI coding programs that are tasked with assembling and writing, and required code on demand without the assistance of human coders.

Others say that human communication is the great barrier, but increasingly the initial experiments with automated help desks and other automated assistance have shown where such services can be applied as an option to deliver fast assistance where the wait for an available person can be lengthy and when available can be slow to understand, and even slower to determine how to fix the offered problem (Chui et al., 2016).

Decision-making will be the last bastion of human workers, but when the Hedge Fund Managers in Stamford Connecticut. It around while the programmed trading of most traffic on the stock exchanges under code provided by the mathematicians, physicists, engineers who inhabit the work; of the "Quants", if our finances are underpinned by AI coded stock, bond, and futures trading, our day to day management decision-making seems rather mundane.

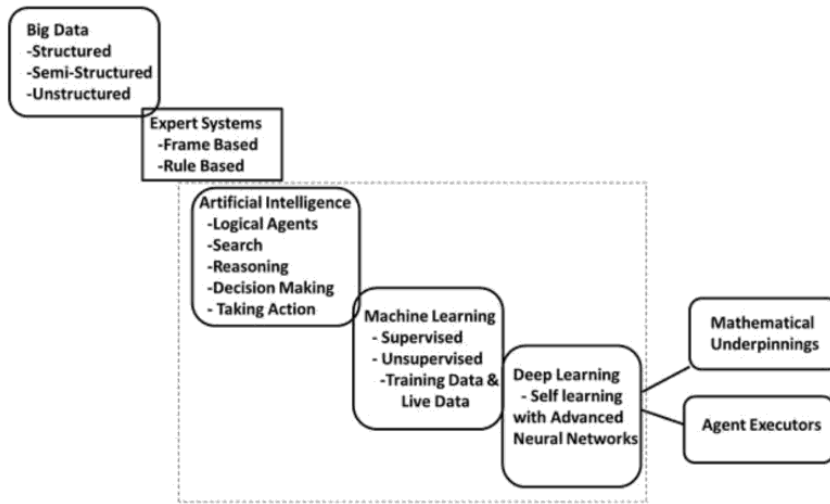


Figure 1.1 The major components that encompass Artificial Intelligence and Machine Learning.

So our book intends to first present you with some of the basic elements of AI and Machine Learning, and then to show how it is currently being employed in a number of major industries in the United States. Figure 1.1 indicates the broad area covered by those who currently discuss AI. We discuss the general topic of deterministic systems which store the knowledge of experts and are triggered into action based upon a component of input that can be translated into an expert's action. These were initially developed in the late 1980s as expert systems with much hope for utilization, but soon faded in the 1990s. Since then, beginning in the late 1990s and early 2000s, AI emerged as a topic of much interest in companies such as Google, particularly to solve their enormous search problems, and IBM, who saw potential to employ such programs on their large computers. These projects quickly advanced toward systems that could be trained with large amounts of data whose identity was known. Neuroscience added additional knowledge of wiring of neurons in human and animal brains (up to 100 billion neurons in humans, each with up to a thousand branches of axons that interconnect these neurons). From this insight, neural networks became the component of Machine Learning which can be taught with millions of live cases, and then begin learning on its own as it experiences live data in the situation where they are deployed. (Deep Learning then incorporated all the advancements resulting from creating increasing numbers of layers of intelligent nodes that could learn from their experiences (Aggarwal, 2018). Figure 1.1 portrays some of the components that comprise our current experiments with AI and Machine Learning.

Our book is organized with some initial chapters on the components and operation of AI and Machine Learning. These chapters follow the structure of the two books that are treated as the "bibles" in these areas: Stuart Russell and Peter Norvig's *Artificial Intelligence: A Modern Approach*, (Russell & Norvig, 2010) which is the standard textbook in virtually all universities offering AI courses, and Ian Goodfellow, Yoshua Bengio, and Aaron Courville's *Deep Learning* (Goodfellow et al., 2017), which is quickly assuming a similar stature covering Machine Learning. Other sources include Charu Aggarwal's *Neural Networks and Deep Learning* (Aggarwal, 2018), Eugene Charniak's *Introduction to Deep Learning* (Charniak, 2018), and Jeff Heaton's three-volume *Artificial Intelligence for Humans* (Heaton, 2015). Furthermore, numerous papers on these subjects are available over the Internet including Vishal Maini and Samer Sabri's very lucid *Machine Learning for Humans* (Maini & Sabri, 2017).

The initial basics chapters are followed by a series of chapters that address how and where AI is being employed by specific companies and industries. We have attempted to cover a broad spectrum of American businesses in order to demonstrate the importance and rapid deployment of resources to utilize such technologies and systems in their daily operation.

References

- Aggarwal, Charu, *Neural Networks and Deep Learning*, Springer, Cham, Switzerland, 2018.
- Berbazzani, Sophia, 10 Jobs Artificial Intelligence Will Replace (and 10 Jobs That Are Safe), *HubSpot*, Nov 7, 2018.
- Bomey, Nathan, Special Report: Automation Puts Jobs in Peril, *USA Today*, Feb 6, 2017.
- Charniak, Eugene, *Introduction to Deep Learning*, MIT Press, Cambridge, MA, 2018.
- Chui, Michael, James Manyika, and Mehdi Miremadi, *Where Machines Could Replace Humans - And Where They Can't*, McKinsey Quarterly, July 2016, <https://www.mckinsey.com/business-functions/mckinsey-digital/our-insights/where-machines-could-replace-humans-and-where-they-cant-yet>.
- Heaton, Jeff, *Artificial Intelligence for Humans, Vol 3 Deep Learning and Neural Networks*, Heaton Research, Inc., Chesterfield, MO, 2015.
- Goodfellow, Ian, Yoshua Bengio, and Aaron Courville, *Deep Learning*, MIT Press, Cambridge, MA, 2016.
- Maini, Vishal and Samer Sabri, *Machine Learning for Humans*, August 19, 2017. <https://medium.com/machine-learning-for-humans/why-machine-learning-matters-6164fafdf12>.
- McKinsey, *Jobs Lost, Jobs Gained: Workforce Transitions in a Time of Automation*, McKinsey Global Institute, Dec 2017, <https://www.mckinsey.com/~media/mckinsey/featured%20insights/future%20of%20organizations/what%20the%20future%20of%20work%20will%20mean%20for%20jobs%20skills%20and%20wages/mgi-jobs-lost-jobs-gained-report-december-6-2017.ashx>.

McKinsey, *Notes from the AI Frontier - Modeling the Impact of AI on the World Economy*, McKinsey Global Institute, Sep 2018.

Russell, Jon, Google's AlphGo AI Wins a Three-Match Series Against the World's Best Go Player, *TechCrunch*, May 5, 2017.
<https://techcrunch.com/2017/05/24/alphago-beats-planets-best-human-go-player-ke-jie/>.

Russell, Stuart and Peter Norvig, *Artificial Intelligence: A Modern Approach*, 3rd ed., Pearson, Hoboken, NJ, 2010.

Chapter 2

The Basic Elements of Artificial Intelligence

Frank M. Groom

Ball State University

Contents

Introduction

Representation

 Representing a Problem

 Representation of Information

 Storing Knowledge

 Representing Knowledge

 Frames

 Trees

 Objects

Formulating Problems and Solutions

Storing Knowledge and Information

 File Structures

 Storage Hardware

Searching for Knowledge and Actions

 Breadth-First Search

 Depth-First Search

 Informed Search

 Sophisticated Search

 Heuristic Search

 Local Search Algorithms

 Adversarial Search and Games

Problem-Solving, Decision-Making, and Taking Action

Reasoning

Agents Employing Reasoning and Making Decisions

Introduction

As businesses recognize the breakthrough potential of Artificial Intelligence (AI) in the past decade, they have discovered that the theoretical possibilities are actually achievable with current software and hardware. From a software perspective, we have seen the IBM Deep Blue overcome world chess champion Garry Kasparov in 1997 after Kasparov won the first match in 1996. However, in 2017 Google's AI program, AlphaGo, beat Ke Jie who is arguably the best player of the most complicated game ever devised. AlphaGo was developed by DeepMind, the AI arm of Google's parent, Alphabet. Most Silicon Valley companies are placing large financial and personnel resources into not only AI research, but actually focusing on software and products that can be quickly brought to market.

Our newspapers delight in describing the successes and occasional tragedies emerging from the introduction of self-driving and assisted driving cars. The number of sensors, controllers, and communication modules offers a test-bed for a large number of AI components which the automotive industry is willing to finance. These automated components which the automotive industry is experimenting with will provide information which all technology companies will use to determine the level of control that localized and remote algorithmic software can exercise.

As these live experiments and deep research are employed, the decade of 2019–2029 is emerging as the anticipated time frame within which companies and customers could expect products that significantly rely on AI algorithms would become available and evidence some level of self-growth by means of programmed self-actuated Machine Learning which will augment the capability of AI control.

AI resides in an agent that interacts with its targeted environment by means of sensors which it either contains or is connected to, and actuators that trigger action in the environment (Heaton, 2013, 2014).

Other information used in a more delayed fashion can reside in an external solid-state disk (SSD) or disk storage which is external and mostly distant from the actual agent enacting an intelligent action.

Representation

Representing a Problem

The first step in dealing with a problem involves establishing a set of well-defined goals each with a specification of a situation and what from that situation would cause

an agent to recognize it by means of its sensors, and to perform a set of actions as a consequence of such recognition. Following the establishment of goals, problem formulation involves creating a definitive link between sensor recognition and action implementation. The set of sensor stimuli, and links to a particular action set, is considered require a set of reasoning steps whose result choses one of the available senses to action links to invoke at that particular instant as the most appropriate problem formulation.

Representation of Information

All information and programs within a computer exist in the digital form (0 and 1s). Of principal importance is how much information is stored as an accessible unit, and how fast and how frequently it can be delivered to the AI agent actuator. In traditional data processing, information can be stored seconds away from the agent employing it for action. With AI enactment, the agent activator usually responds to some recognized stimulation within the time frame of nanoseconds. This requires that the employed information utilized by the actuator reside within the agent or attached to that agent. Remote data storage can host versions of that information, and prepare and download periodic updates to the remote AI agent.

Storing Knowledge

Once gathered, the “intelligence” component of AI must be stored. On the one hand, it can be stored inside the agent that actively makes use of it to address an issue or solve a problem. On the other hand, it can be stored in local but connected to the agent. Then again, it can be stored in a remote location and accessed or downloaded when needed for taking an action decision.

The intelligence can be used by a purely active agent that is either embedded in the local environment, or located close to it, or placed in a remote location. Furthermore, this agent can be a blind reactor to a stimulus or a weak or strong method-based (algorithm executing) problem solver.

Representing Knowledge

Many of the commonly employed approaches for organizing knowledge are based upon the general principles of knowledge representation used in Expert Systems. These include frame structures, scripts, objects, and conceptual graphs and semantic nets.

Frames

A classic approach for storing knowledge is using an algorithmic execution engine having a frame with a structure that contains the following components:

- A **Frame ID** to access the frame and use its enclosed information and execute its algorithm.
- The **Relationship of the Frame** to other frames – frequently implemented as a pointer.
- The **Description** of the enclosed knowledge and agent.
- The **Method** or procedure algorithm that makes use of the encoded knowledge to solve a problem or react to a stimulus and recognized situation.

Trees

Another classic storage of expert information is a tree structure. Information is identified by an ID which might logically appear as Layer 1-Key, Layer 2-Key, Layer 3-Key, and finally a Bottom Leaf-Key. A file of information organized in tree structure is accessed by the first part of the ID Key, where only matching information are selected and the others are avoided, and then it moves to the next ID Key of the branch, making decisions at each branch based upon matching the subsequent components of the ID set until it reaches the bottom level where the desired information lies. The information is obtained as an integrated set whether it be text, numeral, picture, sound, map, or other specialized characters.

Objects

An object structure is formed as a logical tree. At the top of the tree exists a Class component that contains all information shared by a number of similar Object Instances and is identified with the first characters/numbers of the overall Object ID (OID). The later part of the ID is specific to an individual instance of that object. Think of a file of professors as part of the class of professors in a school of a university. All shared information of the school's professors is stored at the school level. Shared department information might reside in a Sub-Class level extended off the school class level. Then at the bottom would exist individual clusters of information associated with Object Instances for each individual professor, containing all the details specific only to that professor (Adshead, n.d.).

Formulating Problems and Solutions

The first step in solving problems is to formulate a goal. This goal is based upon the nature of the presented problem, the situation of that problem, the elements involved, and the capability of the solving AI component. The goal, once stated, allows one to