2nd International World Conference on the
Fundamentals of Artificial Intelligence
July 3-7, 1995
(invited talk)

LEARNING AND COGNITION

Ryszard S. Michalski
George Mason University
Fairfax, VA, USA

and

Institute of Computer Science
Polish Academy of Sciences

Abstract

This presentation consists of three interrelated parts:

1) A discussion of the relationships among several concepts fundamental to understanding intelligent behavior, such as intelligence, learning, cognition and inference,
2) A review of the Inferential Theory of Learning that provides a unifying framework for learning processes, and
3) An introduction to research on the theory of guessing that aims at providing a computational foundation for understanding human plausible reasoning and "educated" guessing.

Presented ideas are intended as contributions to the development of the emerging science of learning and inference.

LEARNING AND COGNITION

Ryszard S. Michalski
George Mason University

Extended Summary

It is a commonly held belief that the ability to learn is an indispensable component of an intelligent behavior. To consider such a view as being more than an intuitive opinion, one needs to have an operational definition of intelligence and learning. While several definitions of these concepts have been proposed in the past, they often lack "operationality," by which we mean that they are defined in terms or concepts that themselves need to be defined.

Also, any intellectually satisfactory definition of intelligence should state conditions that are not inherently biological, but would allow any system--biological or not--be viewed as intelligent, as long as it satisfies these conditions. To satisfy the above criteria, the following definition of the intelligent system is proposed.

A system is called intelligent, if it can:

C1. Perceive, that is, is able to obtain information about its environment and its own state (thus, it is equipped with sensors that measure and interpret properties characterizing its environment and its internal state).

C2. Learn, that is, create knowledge from this information (i.e., can classify, organize, abstract and generalize information obtained from the sensors)

C3. Reason, that is, can use its knowledge for achieving its goals (i.e., can access and use knowledge in order to achieve external or internal goals, or to perform associated with them functions, such as self-preservation, danger avoidance, service, problem solving, planning, decision making, object recognition, prediction, etc.

In this definition, knowledge is defined as organized, abstracted and generalized information; information is defined as interpreted data; and data is defined as a collection of symbols. It is also assumed that an intelligent system has goals for which it is designed and which it is pursuing. The goal can be organized into a structure reflecting their importance. For example, people have genetically implemented goals that require them to pursue tasks organized into a hierarchy (e.g., Maslow's hierarchy), such as to seek safety and self-preservation, satisfaction of biological needs, satisfaction of emotional and intellectual needs, and self-actualization.

Intelligence can thus be described by an "equation":

$$Intelligence = Percetpion + Learning + Reasoning$$

Information Gathering + Knowledge Generation + Knowledge Utilization

In the above definition, intelligence is considered as a property that can have a degree, rather than a yes-no property. Specifically, the degree of to which the above three conditions are satisfied by a system determines the degree of its intelligence. Thus, for example, a desk would be viewed as intelligent, if it is equipped with sensors, can create knowledge from the information obtained by them (e.g., knowledge about what height, tilt, shape, etc. of the desk is most desirable or suitable for various people), and then can use that knowledge to automatically adjust its height, tilt, shape, etc., accordingly to the person that seats at it.

The ability to learn is incorporated in the second condition (C2) of the above definition, since learning can be viewed a process of creating knowledge/skill and memorizing it for future use. The input information to a learning process may include any sensory perception, teacher-provided facts and/or knowledge, the learner's prior knowledge, beliefs, feelings, results of learner's reasoning or imagination. Deriving knowledge from the given information and/or knowledge can be viewed as a process of inference.

Thus, learning can be described by an "equation":

Learning = Inference + Memory

When applied to human learning, this definition requires some explanation. Human learning can be of two types, depending on the type of knowledge that is being generated. There are two fundamental types of human knowledge, each being represented, accessed and used differently. There is explicit knowledge (conceptual, declarative) and implicit knowledge (skill, procedural). The terms "explicit" and "implicit" knowledge have been introduced by psychologists, such as Neal Cohen from the University of Illinois, Larry Squire from the University of California at San Diego, and Daniel Schacter from the University of Toronto. Terms "declarative" and "procedural" knowledge have been used mostly by AI researchers to characterize different knowledge structures (order-independent or order-dependent, respectively), regardless of whether they relate to human mind or computer.

The fundamental aspect of human memory organization is that explicit knowledge is stored in the prefrontal cortex, while implicit knowledge is manifested through activation of particular motor or sensory system. In a computer, both declarative and procedural knowledge can be represented using the same memory structures. Moreover, the transfer from one form of knowledge to another can be done in a computer automatically, at least in principle, while such a transfer can not be done automatically by human brain. No matter how well we "know" how to perform a certain skill we cannot do it well (or at all) until we practice.

The view of learning as knowledge creation (in the learner's mind) is the basis for the Inferential Theory of Learning that aims at providing a unifying framework for all learning processes. The theory views learning as a process of traversing knowledge spaces using knowledge operators, called transmutations or transforms (such as generalization, abstraction, similization, prediction, selection, agglomeration, etc.).

The major contribution of the theory is the distinction between knowledge transmutations that change various aspects of knowledge and usually occur as pairs of opposites, and the type of inference (such as deduction, induction or analogy) that are methods for knowledge transform and that characterize knowledge changes along the truth-falsity dimension. In the current version of the theory has introduced 43 named knowledge transmutations, in addition to a range of knowledge derivations that determine one piece of knowledge from another on the basis of some logical or statistical dependency between them.

In contrast to typical machine learning methods, which are "monostrategy," human learning is multistrategy, which means that it uses multiple learning strategies in a goal-oriented fashion. Multistrategy learning may involve different types of inference and/or knowledge representations. Because any type of inference may derive knowledge that is potentially useful and worth remembering, the complete theory of learning must to encompass the theory of inference. Thus, learning and inference are two intertwined processes that are mutually interdependend (e.g., Gaines and Boose, 1990; Michalski, 1990, 1994).

The last part of the presentation reviews recent ideas on the development of a theory of guessing that attempts to explain how people are able to derive useful knowledge from logically incomplete, inconsistent or uncertain premises. This theory is based on the core theory of human plausible inference (Collins and Michalski, 1989; Collins, Burstein and Baker, 1990), the Inferential Theory of Learning (Michalski, 1993, 1994), knowledge representation based on Dynamic Interlaced Hierarchies (Hieb and Michalski, 1993; Alkharouf and Michalski, 1995) and two-tiered knowledge representation that explains how people represent imprecise concepts (Michalski, 1993). The two outlined theories--the Inferential Theory of Learning and the Theory of Guessing--are viewed as a contribution to the development of the emerging science of learning and inference.

BIBLIOGRAPHY

Ajdukiewicz K., Logika Pragmatyczna (in English: Pragmatic Logic), Panstwowe Wydawnictwo Naukowe (State Scientific Publisher), Warsaw 1974.

Alkharouf, N. and Michalski R.S., Multitype Inference in Dynamic Interlaced Hierachies, Reports of Machine Learning and Inference Laboratory, 1995 (to appear).

Boehm-Davis, D., Dontas, K. and Michalski, R.S., "A Validation and Exploration of Structural Aspects of the Collins-Michalski Theory of Plausible Reasoning," Reports of the Machine Learning and Inference Laboratory, MLI 90-5, Center for Artificial Intelligence, George Mason University, January 1990.

Collins, A. and Michalski, R.S., "The Logic of Plausible Reasoning: A Core Theory," Cognitive Science, Vol. 13, pp. 1-49, 1989.

Collins, A., Burstein M and Baker, M, Final Report: Testing and Refining a Core Theory of Plausible Reasoning, BBN Systems and Technologies Corporation, BBN Report No. 7276, 1990.

Hieb, M. and Michalski, R.S., "Knowledge Representation for Multistrategy Task-adaptive Learning: Dynamic Interlaced Hierarchies," Proceedings of the Second International Workshop on Multistrategy Learning (MSL93), Harpers Ferry, WV, Morgan Kaufmann, May 27-29, 1993.

Gaines B.R. and Boose, J.H. (Eds.), Machine Learning and Uncertain Reasoning, Knowledge-based Systems, Vol. 3, Academic Press, 1990.

Johnson-Laird, P., N., Mental Models, Harvard University Press, 1983.

Michalski, R.S., Dontas, K. and Boehm-Davis, D., "Plausible Reasoning: An Outline of Theory and Experiments," Proceedings of the Fourth International Symposium on Methodologies for Intelligent Systems, pp. 17-19, Charlotte, NC, October 1989.

Michalski, R.S., "Multistrategy Constructive Learning: Toward Unified Theory of Learning," Reports of the Machine Learning and Inference Laboratory, MLI 90-1, Center for Artificial Intelligence, George Mason University, Fairfax, VA, January 1990. (Reprinted in Readings in Knowledge Acquisition and Learning: Automating the Construction and Improvement of Expert Systems, B.G. Buchanan and D.C. Wilkins, Morgan Kaufmann, San Mateo, 1993.)

Michalski, R.S., "Beyond Prototypes and Frames: The Two-tiered Concept Representation," in Categories and Concepts: Theoretical Views and Inductive Data Analysis, I. Van Mechelen, J. Hampton, R.S. Michalski and P. Theuns (Eds.), Academic Press, New York, 1993.

Michalski, R.S., "Inferential Theory of Learning: Developing Foundations for Multistrategy Learning," in Machine Learning: A Multistrategy Approach, Vol. IV, R.S. Michalski and G. Tecuci (Eds.), Morgan Kaufmann, San Mateo, CA, 1994.

Newell, A., Unified Theories of Cognition, Harvard Univeristy Press, 1990.

Posner, M. I. (Ed.), Foundations of Cognitive Science, The MIT Press, 1989.

Szaniawski, K., O Nauce, Rozumowaniu i Wartosciach (in English--About Science, Reasoning and Values), Polish Science Publisher PWN, 1994.

Vosniadou, S. and Ortony, A., Similarity and Analogical Reasoning, Cambridge University Press, 1989.

LEARNING AND COGNITION

This presentation discusses the relationship among several fundamental concepts, such as intelligence, learning, cognition and inference, compares briefly their embodiment in brains and machines, and then overviews the development of the science of learning and inference. It is commonly held that the ability to learn is an indispensible component of an intelligent behavior. To consider such a view as being more than an intuitive opinion, one needs to have an operational definion of of learning and intelligence. While many definitions of these concepts have been proposed in the past, they often lack "operationality," because they are defined in terms that themselves need definition. Also, any satisfactory definition of these concepts should state conditions that are not inherently biological, but would allow any system--biological or not-- be viewed as intelligent, as long as it satisfies these conditions. To satisfy this criterion, the following definition of the intelligent system is proposed. A system is called intelligent, if it can: C1. Generate and store information about its environment (i.e., is equipped with senses that measure/perceive some properties of its environment, and with memory for storing information about these properties) C2. Create knowledge from this information (i.e., can classify, organize, abstract and/or generalize information obtained from the senses) C3. Can use this knowledge for achieving its goals (i.e., such as life preservation, object recognition, prediction, planning, reasoning, decision making, etc.) In this definition, knowledge is defined as organized, abstracted and/or generalized information; information is defined as interpreted data, and data as a collection of symbols. The above definition considers "intelligence" as a property that has a degree, rather than a yes-no property. Thus, the degree to which the above three conditions are satisfied by a system determines the degree of intelligence.

Since human brain exhibits the highest degree of inteligence among all known to us systems, and we strive to endow computers with ever increasing degree of intelligence, it might be useful to exress some well-known facts about brains and computers.

Brains are constructed and maintained jointly by genes and by experience. Neurobiological studies have shown that there is a direct correlation between mental events and patterns of nerve impulses in the the brain. The mind can be viewed an abstraction representing the totality of information processing activities in the brain--this is a view widely shared by neuropsychologists and other researchers, including AI researchers. For example, Minsky, one of the founders of AI, said that "Mind is what brain does." Human brain wighs about 3-4 pounds, has about 100 billion neurons, that is is of the same order of magitude as the number of stars in the Milky Way. The secret of intelligence is not, however, in the number of cells. We know that human liver contains about 100 million cells, but 1000 livers will not substitute for the brain. The function of the cells and the stucture of their interconnections seems crucial. Each neuron continually integrates up to about 1000 synaptic units. Some have axons connecting to the neighbors some to distant regions. It is estimated that there are about 10 to power 12-13 connections in the brain. The integration of signals from different neurons is not done in a simple linear manner. While neurons may significantly differ from each other (in shape and sizes), their underlying biological structure (dendrites, axons, synapses, etc.), and underlying electro-chemical processes that occur in them seem to be similar. In contrast, computers consist of highly specialized functional modules with relatively small number of connections. Thus, the fundamental differences between brains and computers are the highly differentiated behavior of basic modules, and vastly diffent degree of connectivity among the modules.

Brains are high-connectivity systems, while computers are relatively low connectivity systems. It is seems to be true, however, that some functions of the brain are performed by some specialized groups of cells. For example, Hubel and Wiesel who worked on visual cortex and Mountcastle who worked on somatosensory cortex independenly observed that neurons of similar function are grouped together in columns (slabs) that extend through the thickness of the cortext. For example, a module that reponds to a line of a particular orientation measures, as perceived by our eye, is about 1/10 of a mm across the cortex and contains more than 100 000 cells. Electrical impulses in the brain, action potentials----measure about 100 milivolts in amplitude and one milisecond in duration. The maximum speed with which an impulse travels through an axon is about 100 m per second (roughly 10 times the speed of the fastest human beings), but less than one milionth of the speed of electrical signals moving through a wire. Neurons are discrete units, performing hybrid type computations-discrete and analog. Computers perform typically discrete computations. Because they are universal computing machines, they can, however, at least in principle, perform every type of computation, and hence there is the AI hypothesis that intelligence can be realized on a computer.

Summary of differences: Specialization of elements (highly specialized modules/chips very less speciallized neurons (especially , after birht) Interconnectivity Due to the above differences, some functions are much better performed by brains and some by computers.

Brains have a great superiority in such areas as interpretation and integration of sensory information, such as visual and speech, the scope and flexibility of knowledge representations they use, the ability to handle and reason with imprecise, partially known and/or highly abstract concepts. On the other hand, computers are vastly superior to people in perfoming numerical calulations, in reliably storing and accessing vasts amount of data and facts, performing some forms of formal reasoning. Computers are also in the process of continouos improvement and growth, which is not the case with our brains, at least mine. The tasks they can not perform today they may be able to perform tomorrow. The question is then if and how computers can perform some of the functions, like learning and cognition, that are the primary domain of brains.

Brains are constructed and maintained jointly by genes and by experience. Neurobiological studies have shown that there is a direct correlation between mental events and patterns of nerve impulses in the the brain. The mind can be viewed an abstraction representing the totality of information processing activities in the brain--this is a view widely shared by neuropsychologists and other researchers, including AI researchers. For example, Minsky, one of the founders of AI, said that "Mind is what brain does."

Human brain wighs about 3-4 pounds, has about 100 billion neurons, that is is of the same order of magitude as the number of stars in the Milky Way. The secret of intelligence is not, however, in the number of cells. We know that human liver contains about 100 million cells, but 1000 livers will not substitute for the brain. The function of the cells and the stucture of their interconnections seems crucial. Each neuron continually integrates up to about 1000 synaptic units. Some have axons connecting to the neighbors some to distant regions. It is estimated that there are about 10 to power 12-13 connections in the brain.

The integration of signals from different neurons is not done in a simple linear manner. While neurons may significantly differ from each other (in shape and sizes), their underlying biological structure (dendrites, axons, synapses, etc.), and underlying electro-chemical processes that occur in them seem to be similar. In contrast, computers consist of highly specialized functional modules with relatively small number of connections. Thus, the fundamental differences between brains and computers are the highly differentiated behavior of basic modules, and vastly different degree of connectivity among the modules.

Brains are high-connectivity systems, while computers are relatively low connectivity systems. It is seems to be true, however, that some functions of the brain are performed by some specialized groups of cells. For example, Hubel and Wiesel who worked on visual cortex and Mountcastle who worked on somatosensory cortex independenly observed that neurons of similar function are grouped together in columns (slabs) that extend through the thickness of the cortext.

For example, a module that reponds to a line of a particular orientation measures, as perceived by our eye, is about 1/10 of a mm across the cortex and contains more than 100 000 cells. Electrical impulses in the brain, action potentials----measure about 100 milivolts in amplitude and one milisecond in duration. The maximum speed with which an impulse travels through an axon is about 100 m per second (roughly 10 times the speed of the fastest human beings), but less than one milionth of the speed of electrical signals moving through a wire. Neurons are discrete units, performing hybrid type computations-discrete and analog. Computers perform typically discrete computations. Because they are universal computing machines, they can, however, at least in principle, perform every type of computation, and hence there is the AI hypothesis that intelligence can be realized on a computer.

Summary of differences: Specialization of elements (highly specialized modules/chips very less speciallized neurons (especially , after birht) Interconnectivity Due to the above differences, some functions are much better performed by brains and some by computers.

Brains have a great superiority in such areas as interpretation and integration of sensory information, such as visual and speech, the scope and flexibility of knowledge representations they use, the ability to handle and reason with imprecise, partially known and/or highly abstract concepts. On the other hand, computers are vastly superior to people in perfoming numerical calulations, in reliably storing and accessing vasts amount of data and facts, performing some forms of formal reasoning.

Computers are also in the process of continouos improvement and growth, which is not the case with our brains, at least mine. The tasks they can not perform today they may be able to perform tomorrow. The question is then if and how computers can perform some of the functions, like learning and cognition, that are the primary domain of brains.