STUDY ON THE INTELLIGENT AGENT'S NOTIONAL ARCHITECTURE OF ARTIFICIAL INTELLIGENCE

Chamkour Singh[1], Rupali[2]
[1,2]Guru Kashi University, Talwandi Sabo

ABSTRACT

Artificial Intelligence (AI) is a branch of science and engineering concerned with the theory and practice of creating systems that display the we identify with intelligence in human behaviour. Starting with a brief history of artificial intelligence, this paper provides a broad overview of this multidisciplinary field, organized around the main modules of an intelligent agent's notional architecture (knowledge representation; problem solving and planning; knowledge acquisition and learning; natural language, speech, and vision; action processing and robotics), which are highlighted both the main areas of artificial intelligence research, development and application.

Keywords: Artificial Intelligence, cybernetics, Engineering

INTRODUCTION

Artificial Intelligence (AI) is the Science and Engineering domain concerned with the theory and practice of developing systems that exhibit the characteristics we associate with intelligence in human behavior, such as perception, natural language processing, problem solving and planning, learning and adaptation, and acting on the environment. Its main scientific goal is understanding the principles that enable intelligent behavior in humans, animals, and artificial agents. This scientific goal directly supports several engineering goals, such as, developing intelligent agents, formalizing knowledge and mechanizing reasoning in all areas of human endeavor, making working with computers as easy as working with people, and developing human-machine systems that exploit the complementariness of human and automated reasoning.

Artificial Intelligence is a very broad interdisciplinary field which has roots in and intersects with many domains, not only all the computing disciplines, but also mathematics, linguistics, psychology, neuroscience, mechanical engineering, statistics, economics, control theory and cybernetics, philosophy, and many others. It has adopted many concepts and methods from these domains, but it has also contributed back.

Artificial intelligence researchers investigate powerful techniques in their quest for realizing intelligent behavior. But these techniques are pervasive and are no longer considered AI when they reach mainstream use. Examples include time-sharing, symbolic programming languages (e.g., Lisp, Prolog, and Scheme), symbolic mathematics systems (e.g., Mathematica), graphical user interfaces, computer games, object-oriented programming, the personal computer, email, hypertext, and even the software agents. While this tends to diminish the merits of AI, the field is continuously producing new results and, due to its current level of maturity and the increased availability of cheap computational power, it is a key technology in many of today's novel applications.

The next section provides a brief history of the evolution of Artificial Intelligence. This is followed by short presentations of its main areas of research which correspond to the agent modules from Fig. 1.

I. BRIEF HISTORY OF ARTIFICIAL INTELLIGENCE

Artificial intelligence is as old as computer science since from the very beginning computer science researchers were interested in developing intelligent computer systems [3]. The name “artificial intelligence” was proposed by John McCarthy when he and other AI influential figures (Marvin Minsky, Allen Newell, Herbert Simon, a.o.) organized a summer workshop at Dartmouth in 1956.

These successes have generated much enthusiasm and the expectation that AI will soon create machines that think, learn, and create at levels surpassing even human intelligence. However, attempts to apply the developed methods to complex real-world problems have consistently ended in spectacular failures. A famous example is the automatic translation of the phrase “the spirit is willing but the flesh is weak” into Russian, and then back to English, as “the vodka is good but the meat is rotten” [1]. This has led to an AI winter when previously generous funding for AI research was significantly reduced.

Progress in various areas of AI has led to a renewed interest in developing agents that integrate multiple cognitive functions. This, in turn, has led to an understanding that various approaches and methods developed in the isolated subfields of AI (natural language processing, knowledge representation, problem solving and planning, machine learning, robotics, computer vision, etc.) need to be interoperable to both facilitate and take advantage of their integration. This has also led to anunderstanding that the symbolic and subsymbolic approaches to AI are not competing but complementary, and both may be needed in an agent. The result was the development of agent architectures, such as ACT [2], SOAR [3], and Disciple [4], the development of agents for different types of applications (including agents for WWW, search and recommender agents), robots, and multi-agent systems (for instance an intelligent house).

Another aspect of reintegration and interoperability is that algorithms developed in one area are used to improve another area. An example is the use of probabilistic reasoning and machine learning in statistical natural language processing.

II. KNOWLEDGE REPRESENTATION

An intelligent agent has an internal representation of its external environment which allows it to reason about the environment by manipulating the elements of the representation. For each relevant aspect of the environment, such as an object, a relation between objects, a class of objects, a law, or an action, there is an expression in the agent’s knowledge base which represents that aspect. For example, Fig. 2 shows one way to represent the situation shown in its upper-right side. The upper part of Fig. 2 is a hierarchical representation of the objects and their relationships (ontology). Under it is a rule to be used for reasoning about these objects. This mapping between real entities and their representations allows the agent to reason about the environment by manipulating its internal representations and creating new ones.

This simple example illustrates an important architectural characteristic of an intelligent agent, the separation between knowledge and control by separate modules for the knowledge base and the problem solving engine. While the knowledge base contains the data structures that represent the entities from the environment, the inference engine implements general methods of solving input problems based on the knowledge from the knowledge base, as will be discussed in the next section.

When designing the knowledge representation for an intelligent agent, one has to consider four important characteristics [5]. The first is the representational adequacy which characterizes the ability to represent the knowledge needed in a certain application domain. The second is the inferential adequacy which denotes the ability to represent the inferential procedures needed to manipulate the representational structures to inferred new knowledge. The third is the problem solving efficiency characterizing the ability to represent efficient problem solving procedures. Finally, is the learning efficiency characterizing the ability to acquire and learn new knowledge and to integrate it within the agent’s knowledge structures, as well as to modify the existing knowledge structures to better represent the application domain.

III. PROBLEM SOLVING AND PLANNING

Artificial intelligence has developed general methods for theorem proving, problem solving and planning, such as, resolution, state space search, adversarial search, problem reduction, constraint satisfaction, and case-based reasoning. One important characteristic of these methods is the use of heuristic information that guides the search for solutions in large problem spaces.

While heuristics never

guarantee optimal solutions, of even finding a solution, useful heuristics lead to solutions that are good enough most of the time.
In state space search, a problem P is represented as an initial state I, a set O of operators (each transforming a state into a successor state), and a set G of goal states. A solution of the problem P is a finite sequence of applications of operators, such as (O4, O5, O1, O3, O2), that change the initial state into one of the goal states. Consider, for example, a robot that can manipulate the objects. We may ask this robot to bring us the book. The robot needs to find a sequence of actions that transforms the initial state I shown in Fig. 2 into a state G where we have the book in our hands, such as: pick-up cup1, place cup1 on table1, pick-up book1, etc. The definitions of all the actions that the robot can perform (e.g., pick-up, place, etc.), with their applicability conditions and effects on the state of the world, are represented in the knowledge base of the robot. The actual algorithm that applies these operators in order to build the search tree in Fig. 4 is part of the inference engine. The actual tree is built in the Reasoning area.

It is this computational complexity that explains why only in 1997 was an automated agent (Deep Blue of IBM) able to defeat Gary Kasparov, the reigning world champion. The program ran on a very powerful parallel computer generating up to 30 billion positions per move to explore about 14 moves in advance. It contained a database of about 4000 open positions, 700,000 grandmaster games, a large number of end-game solutions, coupled with a heuristic evaluation function based on about 8000 features.

In general, game playing agents are better than humans in games where they can search much of the game space (such as Othello). But they are much weaker in games where the search space is very large, such as Go.

Case-based Reasoning is a form of problem solving by analogy in which a new problem is solved by recognizing its similarity to a previously solved problem (which could be classifying the disease of a patient, planning a meal, or designing a circuit), then transferring and adjusting the solution to the new problem.

Another general problem solving method that has been employed in expert systems for a wide variety of tasks, including planning, design, critiquing, symbolic integration, and intelligence analysis, is problem reduction2, 35. In this approach a problem is solved by successively reducing it top-down to simpler problems, finding the solutions of the simplest problems, and combining these solutions, from bottom-up, to obtain the solution of the initial problem.

This top-down problem reduction continues until one reaches problems which have known solutions. Then these solutions are successively combined, from bottom-up, to obtain the solutions of the upper-level problems, and of the top-level problem. In the illustration from Fig. 5, these solutions are probabilistic (e.g., "It is almost certain that the people the United States desire the United States to be a global leader in wind power within the next decade.") and are combined using operators such as min, max, average, or weighted sum.

An important characteristic of the problem reduction method is that it shows very clearly the reasoning logic, making it suitable for developing knowledge-based agents that assist can experts and non-experts in problem-solving, and can teach expert problem solving to students.

IV. KNOWLEDGE ACQUISITION AND LEARNING

Much of the power of an intelligent agent derives from the knowledge in its knowledge base. A main goal of the knowledge acquisition and machine learning research is precisely to enable an agent to acquire or learn this knowledge from a user, from input data, or from agent's own problem solving experience. This results in improving the competence of the agent in solving a broader class of problems, and in making fewer mistakes in problem solving. It may also result in improving the efficiency of the agent in solving the problems faster and with less memory.

Due to the high complexity of learning, much of the research has focused on the basic task of concept learning, such as learning the concept "cup," or the concept "person who will default on bank loan." In essence, concept learning consists in finding a classification function which distinguishes between the entities that are instances of the concepts from those that are not. Many of the developed learning strategies can be characterized as empirical inductive learning from examples, which consists of learning the definition of a concept by comparing positive and negative examples of the concept in terms of their similarities and differences, and inductively creating a generalized description of the similarities of the positive examples [6], [7]. Some methods are based on the information theory to learn the concept in the form of a decision tree [8] that is used to classify the objects. Other methods represent the learned concept as a neural network, whose output unit determines whether the entity at its input units belongs or not to the concept. Learning in a neural network consists in continuously classifying known examples and updating the weights associated with the connection between the units, to improve the recognition accuracy [9]. Support vector classifiers map the positive and the negative examples of the concept, nonlinearly, into a higher-dimensional feature space via a kernel function, and construct a separating hyperplane there with maximum margin which yields a nonlinear decision boundary in the input space [10]. Bayesian classifiers determine the most likely hypothesis or concept by using the Bayes' rule $P(H|E^*)=P(E^*|H)\cdot P(H)/P(E^*)$ that computes the posterior probability of the hypothesis H based on its prior probability and the observed evidence. This type of learning proved to be very effective in applications where prior probabilities can be computed, and is extensively used for statistical natural language processing [9].

These learning methods may be used to extend the ontology of an agent with new concepts or facts, or to learn and refine its reasoning rules. Many of these methods are complementary in terms of the input from which the agent learns, the a priori knowledge the agent needs in order to learn, the inferences made during learning, what is actually learned, and the effect of learning on agent's performance. For instance, in the case of empirical inductive learning from examples, in which the primary type of

\inference is induction, the input may consist of many (positive and/or negative) examples of some concept C, the knowledge base usually contains only a small amount of knowledge related to the input, and the goal is to learn a description of the concept C in the form of an inductive generalization of the positive examples which does not cover the negative examples. This description extends or refines the knowledge base and improves the competence of the agent in solving a larger class of problems and in making fewer mistakes.

In the case of explanation-based learning, in which the primary type of inference is deduction, the input may consist of only one example of a concept C, the knowledge base should contain complete knowledge about the input, and the goal is to learn an operational description of C in the form of a deductive generalization of the input example. This description is a reorganization of some knowledge pieces from the knowledge base and improves the problem solving efficiency of the agent.

V. NATURAL LANGUAGE, SPEECH, AND VISION

The perceptual processing module in Fig. 1 summarizes agent's capabilities to process natural language, speech, and visual inputs. All are very easy for humans and very difficult for automated agents.

When an agent receives input in natural language, it has to understand it, that is, to build an internal representation of its meaning, which can then be used by the problem solving engine. This process, however, is very difficult for several reasons. Natural language is ambiguous at all levels: morphology, syntax, semantics, and discourse [11], [12]. Just by hearing a common word such as "run" we cannot say whether it is a noun or a verb. The WordNet semantic dictionary [13] gives 16 senses for the noun interpretation and 41 senses for the verb interpretation. What does the word "diamond" mean? Does it mean the mineral consisting of nearly pure carbon in crystalline form? Does it mean a gem or other piece cut from this mineral? Does it mean a lozenge-shaped plane figure (♦)? Does it mean the playing field in Baseball? Therefore the meaning of a word needs to be interpreted in the context of its surrounding words. But sentences themselves may be ambiguous. What does "Visiting relatives can be boring" mean? Does it mean that the act of visiting relatives can be boring? Maybe it means that the relatives who visit us can be boring. Consider also the possible meanings of the following sentence: "She told the man that she hated to run alone." Therefore the meanings of individual sentences need themselves to be interpreted in the context of the paragraphs that contain them. This is also necessary because of additional complexities of natural language, such as, the use of paraphrases (where the same meaning may be expressed by different sentences), ellipses (the use of sentences that appear ill-formed because they are incomplete, which requires the extraction of the missing parts from previous

sentences), and references (where entities are referred by pronouns, such as "it" or "they," without giving their names). But even considering larger paragraphs may not be enough for understanding their meaning, unless the agent has a large amount of knowledge about the domain of discourse.

VI. ACTION PROCESSING AND ROBOTICS

The action processing module in Fig. 1 corresponds to the agent's actions upon that environment aimed at realizing the goals or tasks for which it was designed. Such an action could be the generation of an answer to a question, the solution of an input problem, the manipulation of an object, or the navigation to a new position. Since most of these actions have already been addressed in the above sections, here we only address object manipulation and navigation, which are the main concern of the robotics area [15].

One may distinguish between three main types of robots: (1) manipulators which are robotic arms attached to their workspace, such as those used in car assembly and painting; (2) mobile robots with wheels, legs, or wings, used to move objects around, such as the unmanned air vehicles for surveillance, crop-spraying, or military operations, or the autonomous underwater vehicles; and (3) mobile manipulators that combine mobility with manipulation to accomplish more complex tasks. A challenging problem in robotics is localization and mapping, which consists in finding out where things are and building a map of the environment. Another challenging problem is path planning from one point in space to another point, which may involve compliant motion, where the robot moves while maintaining physical contact with an object (e.g., an obstacle, a box it pushes, or a screw it inserts).

Robots have many applications in industry (e.g., for part assembly or painting), agriculture (e.g., as special machines), transportation (e.g., autonomous vehicles), health care (e.g., as devices for surgery), hazardous environments (e.g., for clearing minefields or cleaning up nuclear waste), space exploration, and entertainment. They can provide personal services (e.g., vacuum cleaning), or can act as human augmentation devices (e.g., by providing additional force to facilitate walking or arm movement).

VII. CONCLUSION

The main goal of Artificial Intelligence is to develop computational agents that exhibit the characteristics we associate with intelligence in human behavior. Such an agent has an internal representation of its external environment which is at their basis of its reasoning abilities. In general, an agent solves complex real-world problems by using large amounts of knowledge and heuristic methods.

It is highly desirable that the agent's knowledge and reasoning are understandable to humans, and the agent is able to explain its behavior, what decisions it is making, and why. The agent may reason with data items that are more or less in contradiction with one another, and may provide some solution without having all the relevant data. The agent should be able to communicate with its users, ideally in natural language, and it may continuously learn.

Humans are slow, sloppy, forgetful, implicit, and subjective. But they have common sense and intuition, and may find creative solutions in new situations. By contrast, agents are fast, rigorous, precise, explicit, and objective. But they lack common sense and the ability to deal with novel situations. Humans and agents may thus engage in mixed-initiative reasoning that takes advantage of their complementary strengths and reasoning styles. As such, intelligent agents enable us to do our tasks better, and help us in coping with the increasing challenges of globalization and the rapid evolution toward the knowledge economies.

REFERENCES

[1]. Anderson JR. The Architecture of Cognition. Harvard University Press, 1983.

[2]. Buchanan BG, Sutherland GL, Feigenbaum EA, Heuristic DENDRAL: A program for generating explanatory hypotheses in organic chemistry. In Meltzer B, Michie D, Swan M, ed. Machine Intelligence 1969, 4:209-254.

[3]. Campbell MS, Hoane AJ, Hsu F-H. Deep Blue. Artificial Intelligence Journal 2002, 134:57-83. [4]. Cohen LJ. An Introduction to the Philosophy of Induction and Probability. Clarendon Press, Oxford, 1989.

[5]. Mitchell TM. Machine Learning. McGraw-Hill, 1997.

[6]. Mitchell TM. Version Spaces: An Approach to Concept Learning. Doctoral Dissertation. Stanford University, 1978.

[7]. Nilsson NJ. Problem Solving Methods in Artificial Intelligence. McGraw-Hill, New York, NY, 1971.

[8]. OWL 2 Web Ontology Language, http://www.w3.org/TR/owl2-overview/ (accessed August 14 2011)

[9]. Pearl J. Causality: Models, Reasoning, and Inference. Cambridge University Press, New York, 2009.

[10]. Tecuci G. Building Intelligent Agents: An Apprenticeship Multistrategy Learning Theory, Methodology, Tool and Case Studies, Academic Press, San Diego, 1998.

[11]. Tecuci G. DISCIPLE: A Theory, Methodology and System for Learning Expert Knowledge. Thèse de Docteur en Science. University of Paris-South, 1988.

[12]. Tecuci G. Plausible justification trees: a framework for the deep and dynamic integration of learningToward Knowledge Societies, http://unesdoc.unesco.org/images/0014/001418/141843e.pdf

[13]. Tufiş D, Ion R, Ide N. Fine-grained word sense disambiguation based on parallel corpora, word alignment, word clustering and aligned wordnets. In Proceedings of the 20th International Conference on Computational Linguistics, COLING2004, Geneva, 2004 1312-1318.