

Alexander Gelbukh
Carlos Alberto Reyes-Garcia (Eds.)

LNAI 4293

MICAI 2006: Advances in Artificial Intelligence

5th Mexican International Conference on Artificial Intelligence
Apizaco, Mexico, November 2006
Proceedings



 Springer

Series Editors

Jaime G. Carbonell, Carnegie Mellon University, Pittsburgh, PA, USA
Jörg Siekmann, University of Saarland, Saarbrücken, Germany

Volume Editors

Alexander Gelbukh
Centro de Investigación en Computación
Instituto Politécnico Nacional
Col. Nueva Industrial Vallejo, 07738, DF, Mexico
E-mail: gelbukh@gelbukh.com

Carlos Alberto Reyes-Garcia
Instituto Nacional de Astrofísica, Óptica y Electrónica (INAOE)
Luis Enrique Erro No. 1, Sta. Ma. Tonanzintla, Puebla, 72840, Mexico
E-mail: kargaxxi@inaoep.mx

Library of Congress Control Number: 2006936081

CR Subject Classification (1998): I.2, F.1, I.4, F.4.1

LNCS Sublibrary: SL 7 – Artificial Intelligence

ISSN 0302-9743
ISBN-10 3-540-49026-4 Springer Berlin Heidelberg New York
ISBN-13 978-3-540-49026-5 Springer Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

Springer is a part of Springer Science+Business Media
springer.com

© Springer-Verlag Berlin Heidelberg 2006
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India
Printed on acid-free paper SPIN: 11925231 06/3142 5 4 3 2 1 0

Table of Contents

→ Artificial Intelligence Arrives to the 21st Century 1
Adolfo Guzman-Arenas

Knowledge Representation and Reasoning

Properties of Markovian Subgraphs of a Decomposable Graph 15
Sung-Ho Kim

Pre-conceptual Schema: A Conceptual-Graph-Like Knowledge
Representation for Requirements Elicitation 27
*Carlos Mario Zapata Jaramillo, Alexander Gelbukh,
Fernando Arango Isaza*

A Recognition-Inference Procedure for a Knowledge Representation
Scheme Based on Fuzzy Petri Nets 38
Slobodan Ribarić, Nikola Pavešić

Inference Scheme for Order-Sorted Logic Using Noun Phrases with
Variables as Sorts 49
Masaki Kitano, Seikoh Nishita, Tsutomu Ishikawa

Answer Set General Theorics and Preferences 59
Mauricio Osorio, Claudia Zepeda

A Framework for the E-R Computational Creativity Model 70
Rodrigo García, Pablo Gervás, Raquel Hervás, Rafael Pérez y Pérez

Fuzzy Logic and Fuzzy Control

First-Order Interval Type-1 Non-singleton Type-2 TSK Fuzzy Logic
Systems 81
Gerardo Maximiliano Mendez, Luis Adolfo Leduc

Fuzzy State Estimation of Discrete Event Systems 90
Juan Carlos González-Castolo, Ernesto López-Mellado

Real-Time Adaptive Fuzzy Motivations for Evolutionary Behavior
Learning by a Mobile Robot 101
*Wolfgang Freund, Tomas Arredondo Vidal, César Muñoz,
Nicolás Navarro, Fernando Quirós*

Artificial Intelligence Arrives to the 21st Century

Adolfo Guzman-Arenas

Centro de Investigación en Computación, Instituto Politécnico Nacional, Mexico City
a.guzman@acm.org

Abstract. The paper follows the path that AI has taken since its beginnings until the brink of the third millennium. New areas, such as Agents, have sprout; other subjects (Learning) have diminished. Areas have separated (Vision, Image Processing) and became independent, self-standing. Some areas have acquired formality and rigor (Vision). Important problems (Spelling, Chess) have been solved. Other problems (Disambiguation) are almost solved or about to be solved. Many challenges (Natural language translation) still remain. A few parts of the near future are sketched through predictions: important problems about to be solved, and the relation of specific AI areas with other areas of Computer Science.

1 The Evolution of Artificial Intelligence

Artificial Intelligence (AI) is the branch of computer science that deals with intelligent behavior, learning and adaptation in machines. It is also defined as intelligence exhibited by an artificial (*non-natural, man-made*) entity (<http://en.wikipedia.org/wiki/>). Although there is no “standard breakdown” of AI, traditionally it has been divided in several well defined areas. Initial areas were AI Programming Languages (including Symbolic Manipulation), Theorem Proving (later, Reasoning appears), Vision (including Image Processing, Scene Analysis), Games, Learning, Neural Networks (Perceptrons. Later, with the inclusion of Fuzzy Sets and Genetic Algorithms, expands to Soft Computing), Knowledge Representation (including Semantic Nets; later, Ontologies appear, and within them, Formal Concept Analysis), Robotics, Natural Language Translation (sometimes mixed with Information Retrieval; later, it expands to Natural Language Processing; including Intelligent Text Processing), Constraint Satisfaction, Search (later, Web search appears, as well as Web Processing). Areas appearing later are Distributed AI (including Agents), Expert Systems (including non-procedural languages and systems; later, Diagnosis appears), Qualitative Physics, Lisp Machines. See figure 1.

Sometimes considered as part of AI, but are not included here, are: Philosophical Foundations; Intelligent User Interfaces.

In general, AI has grown in breadth and in depth. Rigor and formalism has been introduced in many of its areas. Various applications have been developed (expert systems; fuzzy systems...). With the availability of a great quantity of texts and information through the Web, Search and Semantic Processing are acquiring vigor. AI has been influenced by concurrent development in other Computer Science areas.

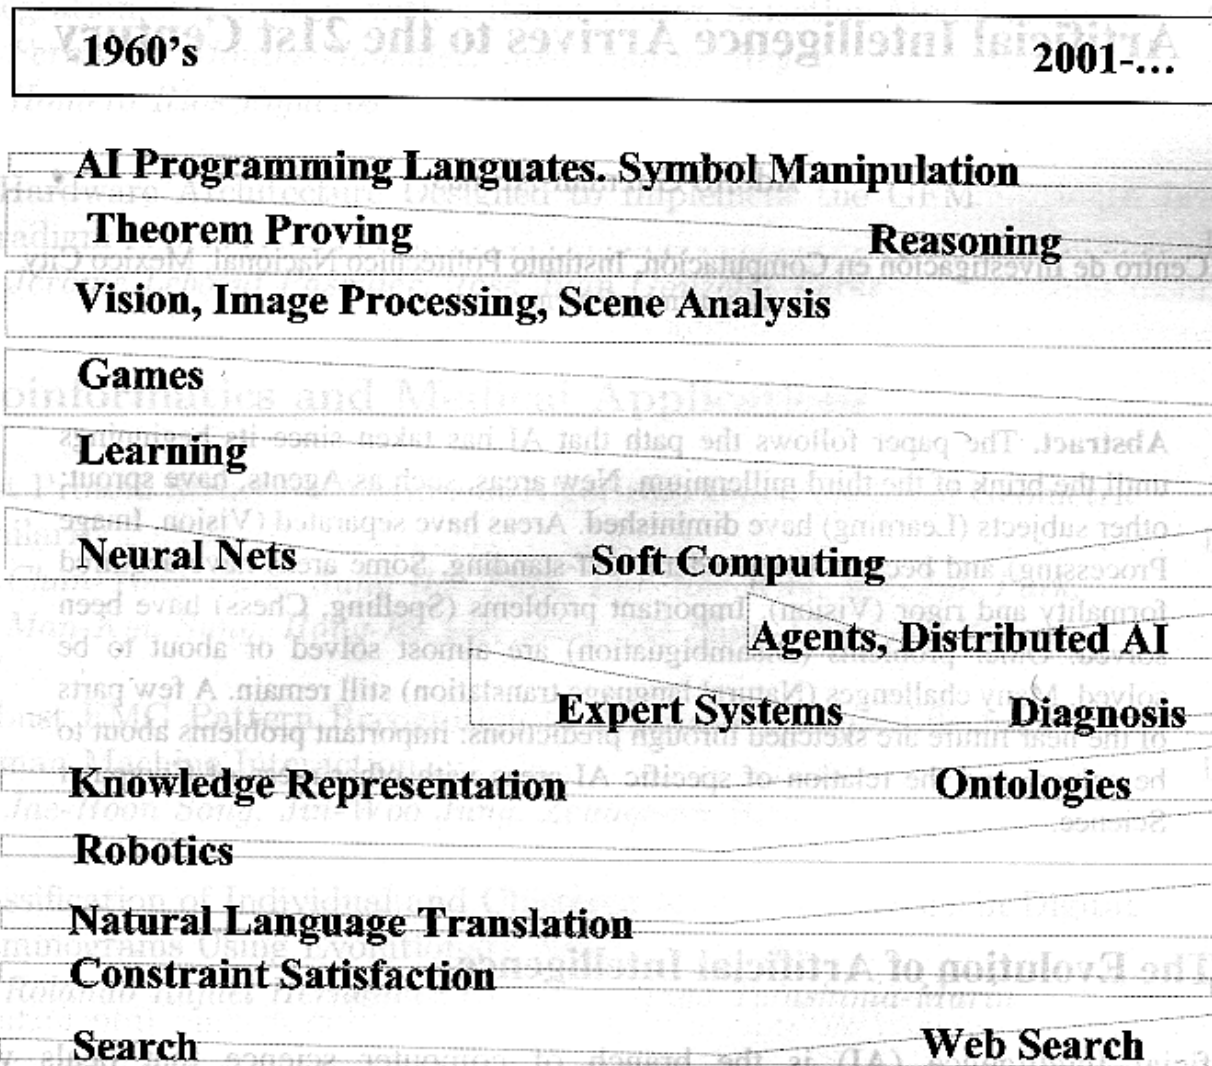


Fig. 1. Path of AI since the early 60's up to day. Some areas have diminished; other have flourished. The text gives more details on each area's development.

2 The Areas of AI

Comments on the development and path of particular areas follow.

2.1 AI Programming Languages

Symbolic (as opposed to numeric) manipulation languages were invented at the beginning of AI: Lisp ([12], Common Lisp or CL, early 90's; CMUCL, a free CL implementation, <http://www.cons.org/cmucl/>) handles lists; Snobol and Comit manipulate strings; Convert [*7] is a pattern matching language suitable for lists; Scheme [22] manipulates lists with *lazy evaluation* or continuations. Lisp, Snobol and Scheme survive until today. Current development of programming languages concentrates in general purpose (not AI) languages. But see §2.9.

2.1.1 Symbolic Manipulation

Early systems were Macsyma [14] and MathLab. The area has gone a long way, with useful commercial products now: MatLab, Maple (www.maplesoft.com). We have

here a mature area where most developments are improvements to specific applications and packages.

2.2 Theorem Proving

Early advances: J. A. Robinson's resolution principle [19], and the demystification of The Frame Problem (John McCarthy).

Advances during 1970's: Non-monotonic reasoning.

Current work: Satisfiability, phase transitions. Spatial reasoning, temporal reasoning (these last two themes belong to Reasoning, found in §2.9). Causality [1]. Description Logics.

Current work is characterized by rigor and scanty applications.

Reasoning was included here initially, due to its dependency at that time on the use of theoretical tools. Later, it has migrated to Knowledge Representations (§2.9), since it depends more and more on the available knowledge.

8% of the International Joint Conference on Artificial Intelligence, 2003 (IJCAI03) sessions (including IAAI03, Innovative Applications of AI; 84 sessions and 200 papers in both) were on Theorem Proving, represented as {8% in IJCAI03}.

2.3 Vision, Image Processing

Early advances: Use of perspective to recognize 3-d solids [18]; decomposition of a scene into bodies [*19].

Some advances in the 1980's: focus of attention, shape from shading, from texture.

The field has grown. At the beginning most advances were empirical discoveries, consisting of methods that work in a few cases. Now, Vision is now based on Mathematics (for instance, Stereo Vision); in Physics (reflections, color of light); on distributed computation (multiple views); on Pattern Recognition; on probabilities. But it currently lacks more Artificial Intelligence influence Takeo Kanade says (Keynote Speech, IJCAI03): we need to reinsert AI into Vision.

Vision and Image Processing are now separating from AI; they have their own Journals (such as Pattern Recognition, Computer Vision and Image Understanding) and scientific meetings (such as CVPR, ICCV). That explains the low numbers of vision papers in IJCAI conferences. 2% of IJCAI03 sessions (including IAAI03) were on Vision, and 2% of IJCAI05 papers (total: 239 papers, excluding poster presentations) were on Vision, too; represented as {2% in IJCAI03; 2% in IJCAI05}.

2.3.1 Image Processing

The processing (deblurring, thinning, sharpening..., but also clustering, classification...) of images through digital algorithms. This is considered the "low part" of Vision, whereas the "high end" is Scene Analysis.

2.3.1.1 Remote Sensing. The processing of pictures taken from Landsat and other satellites. It is somewhat separated from main-stream Vision. Relies mainly on trained classifiers (§2.3.3.1) working on the light spectrum, including infrared. Now, this area is mostly applied. Current research: sensor fusion.

2.3.1.2 Character Recognition. Now: mainly applied work, heuristic-oriented. State of the practice Commercial OCRs have still one-digit error percentages, close to 9%.

2.3.1.3 Text Recognition. Deemphasized since the volume of handwritten and typed text has diminished; most new documents are now born in digital form.

2.3.1.4 Real-Time Text Recognition. Recognition of characters while they are being written. Mainly a solved problem.

2.3.2 Pattern Recognition

The categorization of raw data (it may be a picture, a signal...) and taking an action based on the category of the data. More general than Image Processing, since it handles other kinds of signals. It is not considered a part of AI. It has its own Journals (IEEE Transactions on Pattern Analysis and Machine Intelligence; Pattern Recognition Letters...) and conferences.

Pattern Recognition has matured and uses a formal approach, both in the *syntactical* pattern recognition as well as in the *statistical* pattern recognition.

2.3.2.1 Classification, Clustering. The assignment of classes to a set of objects with given properties. The grouping of objects into subsets (clusters) of similar objects.

Current work: Partially labeled data, Clustering; Learning with Bayesian Networks.

2.3.3 Motion Analysis, Video

Work in This Area Continues. In 1981, Two Powerful Methods to Estimate the Optical Flow were Devised [8, 11].

Optical flow is a technique for estimating object motion in a sequence of video frames. The displacement through time looks as if pixels originate at a given point in the image; this is the point towards which the camera moves. A vector field (vectors placed at each pixel of the image) is described.

2.4 Games

Games were early members of AI; specifically, parlor games: checkers (Samuel in 1957), chess, go, go-moku. Recently, theoretical advances have shown in Market Games (Cf. contest in IJCAI03).

Already solved: Chess by machine has been solved, in the sense that a machine (Deep Blue) is at the level of a World Master. The advance has been possible by much computer power (parallel machines) and large storage (to keep many book games), and to a lesser degree, by advances in AI.

Current tools: Generalized utility (Decision Theory). [2], Nash Equilibria (Cf. IJCAI03).

{4% in IJCAI03}

2.4.1 Game Theory

Invented by John Von Neuman, it has not been widely used in AI Games, except for Market Games.

2.5 Learning

Machine learning develops algorithms that allow computers to learn, that is, to gradually acquire and systematize knowledge from a set of actions, facts or scenarios. An early example of a computer program that learns is Samuel's Checkers playing program. Algorithms for learning can use Statistics (for instance, Bayesian Statistics), Neural Networks (§2.6), Classifiers (supervised learning, §2.3.3.1), Clustering (unsupervised learning), Genetic Programming (§2.6.2), Support Vector Machines... Learning work continues strong. *Current topics*: Kernel Methods; tree learning; ensembles. See for instance IJCAI03.

{10% in IJCAI03; 19% in IJCAI05}

Classification and clustering. These subjects migrated to Pattern Recognition (§2.3.3), listed under Vision and Image Processing (§2.3). They are generally not considered part of AI.

2.6 Neural Networks. Soft Computing

An (artificial) Neural Network (NN) is a group of units (artificial neurons) linked through a network of connections (connectionist approach). The structure of the NN changes with time, as it *learns* or adapts; it is an adaptive system.

First neural networks (perceptrons) were quickly adopted because the *Perceptron Theorem* guaranteed learning (convergence to the sought function) under reasonable conditions. In 1969, the Perceptrons book [13] showed sizable limitations of some of these networks.

Multi-layer perceptrons consist of several layers of neurons, where layer k gets all its inputs from layer $k-1$. The *universal approximation theorem* states that a multi-layer perceptron with just one hidden layer is capable of approximating any continuous function from R_i to R_j , where R_i and R_j are two intervals of real numbers.

Advances: Simulated annealing [9], a learning algorithm for locating a global approximation in a large search space, developed in 1983, quickly revived Neural Networks. *Hopfield networks*, where all connections are symmetric, and *Kohonen self-organizing maps* are also significant advances.

Neural networks are now well understood, they are applicable to a wide variety of problems, although their specificity is below fine-tuned special tools.

2.6.1 Fuzzy Sets

Long time after they were invented in 1965 by Lofti A. Zadeh [25], fuzzy sets were somewhat unpopular with North American scientists, while they were slowly gaining acceptance elsewhere, notably in Japan.

As Zadeh recently has said, up to date more than 5,000 patents exist (4,801 issued in Japan; around 1,700 in U.S.) that use or exploit fuzzy sets. Although fuzzy sets are not in the mainstream of AI, the Portal of Artificial Intelligence in Wikipedia,

http://en.wikipedia.org/wiki/Portal:Artificial_intelligence, considers them “a modern approach to AI”. They are useful tools to master vagueness.

2.6.2 Genetic Algorithms

Also called Genetic Programming or Evolutionary Algorithms, they are inspired in biology, involving mutation, cross-over, reproduction, generations survival of the fittest, and guided search. The more general term is Evolutionary Computation (it tries to optimize combinatorial problems), of which Genetic Algorithms are part of. Sometimes, Swarm Behavior (§2.7) is included in Evolutionary Computation.

Genetic algorithms have moved steadily, and it is a vigorous area useful to do search in large spaces.

2.7 Distributed AI. Agents

Distributed AI studies how to solve complex problems (requiring intelligence) through distributed or parallel computations. An agent is an entity that detects (senses) a portion of its environment, and reacts (acts, produces changes) to them. Usually it has a plan and a set of resources, and it is able to communicate with other agents.

Most distributed solutions with agents require in fact *many* agents. Thus, agents are almost gone; papers have migrated to Multiagents.

Disadvantage: To simulate complex behavior or solve challenging problems, often you have to program many agents of different types (behaviors). This is a heavy task.

Promising tool: Swarm Behavior, Swarm Intelligence. It is a technique that mirrors the collective behavior of groups, hordes, schools and packs of animals that exhibit self-organization. For instance, a tool mimicking ant colonies leaves “chemical traces” in visited places (nodes of a graph); other ants visiting these places detect these pheromones and reinforce (deposit more traces) or weaken the track (go away from those places). Thus, links (tracks) among these nodes are gradually formed, with the help of all ants. Links change with time, some are reinforced, some disappear. Usually, the algorithm converges. *Disadvantage:* it is slow. It reminds me of relaxation labeling of Hummer & Zucker (1983). Links so placed take more into account the global context and are more flexible than links placed once and for all.

{ 15% in IJCAI03, of which 1% is in Swarm Behavior }

2.7.1 Multiagents

A multiagent system is a set of agents that achieve goals by distributed (collective) computation. They cooperate among themselves. *Current tools:* formation of coalitions, Cf. IJCAI03. The field is theoretically dominated, somewhat in early stages; almost no applications. Academically oriented.

{ 14% in IJCAI03; 5% in IJCAI05 }

2.8 Expert Systems, Diagnosis. Non-procedural Systems

Expert Systems or Knowledge-based systems are computer programs that embody a set of rules capturing the knowledge of some human experts. They appeared late in the 60's and grew in the 70's. Applications flourished.

Advantages: the system can be neatly divided into: (a) a set of rules of the domain knowledge (say, Infectious Diseases); (b) a deductive or computing machine, that applies the rules; (c) a data base or set of facts, involving a particular sample: a particular patient, say. The deduction machine (b) is independent of the set of rules, thus programming is reused, while the set (a) of rules can be obtained by interviewing the expert(s). The rules “apply as they see fit”, so that a non-procedural system is obtained.

Disadvantages: Soon two drawbacks emerged: (1) the encompassed knowledge was rather narrow, and the system exhibited fragility at the periphery of its knowledge (the *brittleness* problem): the system does not know that it does not know (and the user is unaware of this, too); (2) too many rules begin interact in undesirable manners.

Due to these and other reasons, work on Expert Systems has diminished, and it now concentrates mainly in applications, for instance, in Diagnosis.

{2% in IJCAI03}

2.9 Knowledge Representation and Reasoning

Knowledge is structured information that represents or generalizes a set of facts in the real world. *Knowledge Representation* is the way in which this knowledge is organized and stored by the computer, to make ready use of it.

Ross Quillian’s semantic nets were one of the early knowledge representations. In 1976, Sowa [20, 21] put forward *Conceptual Graphs*, a way to represent knowledge in patterns of interconnected nodes and arcs. Educators also put forward *Concept Maps* [15] as a tool to communicate knowledge. All these come under the generic term of *Semantic Networks*, the *Ontologies* (§2.9.2) being among the most precise of the modern representation schemes.

In the 80’s, CYC (www.cyc.com) made a brave attempt to construct a common sense ontology. Wordnet (wordnet.princeton.edu) represents a notable contribution from the natural language community. A semantic lexicon for the English languages, it groups words into sets of synonyms (synsets), and contains short definitions, as well as semantic relations among these synsets. It is free. In 2006 it contains 115,000 synsets. There is also a Wordnet for Spanish words.

Probably as a result of the proliferation of documents and information in Internet, work on knowledge representation has rekindled. An approach (§2.11.3.1) is to tag each document, Web page and information source, so as to facilitate their understanding by bots or crawlers (programs that search and read text lying in the Web). A less manual approach is to have a software that extracts knowledge from these sources and stores it in a Knowledge Representation format or language (for instance, Ontolingua [17]) suitable for further useful processing: data integration (§2.9.1), Ontology mapping (§2.9.2.3), Alignment (§2.9.2.4), Ontology Fusion (§2.9.2.5), Reasoning (§2.9.3), etc.

A more shallow form of reasoning is to “ask the right question” to the Web. For instance, in order to find the author of “War and Peace”, you can search the Web for the phrase “The author of War and Peace is...”. This is usually referred to as “Text mining” (§2.11.3.2). Etzioni [5] shows a way to generate (by computer) search phrases from successes with earlier search phrases.

{31% in IJCAI03; 19% in IJCAI05}

2.9.1 Data Integration

It is the combination of data residing in different sources (mainly databases), providing the user with a unified view of these data [10]. Main approaches are: Global-as-View (GAV), where changes in information sources requires revision of a global schema and mapping between the global schema and local schemas; Local-as-View (LAV), which maps a particular query to the schema of each information source; and a hybrid approach [24].

{5% in IJCAI03}

2.9.2 Ontologies

An ontology is a data model representing a domain, a part of the real world. It is used to reason about instances and types in the domain, and the relations between them. Ontologies are precise representations of *shared knowledge*, and usually are formed by nodes (or concepts: types and instances), arcs (the relations among them) and restrictions (logical assertions holding among nodes or relations).

2.9.2.1 Formal Formulations. Formal formulations establish logic restrictions among instances and types. Example: Formal Concept Analysis [7], which defines “concept” as a unary predicate. Problem: everything is a concept, not only those “important concepts”, which I define as those concepts that have a name in a natural language: they are popular enough so that a word has been coined for each.

Defining ontologies with the help of local constraints (say, assertions in some Logics) has the following problem. The restrictions imposed on the instances, types and relations are opaque to (difficult to process by) the software trying to understand the Ontology’s knowledge. This knowledge is stored not only in the nodes and the relations, but also in these restrictions, which are usually written in a notation that the deductive machinery finds difficult to decipher, manipulate and reason about. For instance, it could be difficult to derive new restrictions by processing old restrictions. A way to overcome this, suggested by Doug Lenat, is to express the restrictions *and the software that processes the ontology* in the same notation that other knowledge (such as “Clyde is an elephant”) in the ontology is represented. That is, to represent the restrictions and the software by elements of the ontology (and not in Lisp or in Logical notation). This will render both restrictions and software accessible and open to the deductive machinery. Something like *reflection* in Computer Science. An idea waiting to be implemented.

Another problem with some formal approaches is that almost every assertion has an exception in everyday’s life (Rabbits usually have four legs, but a rabbit may have just three legs and still be a rabbit), so that they have to be expressed, too.

2.9.2.2 Unique Ontology. Common Sense Ontolog. CYC’s idea was to build a single ontology to represent common sense knowledge (a kind of encyclopedia of everyday’s knowledge), and have everybody use it and (perhaps) extend it. This is a worthwhile goal, but the construction of such unique ontology was found to be a challenging task. One problem was simply the size of the effort. Other difficulty was to select the “best view” for representing certain aspects of the real world (emotions,

say). So, it is more practical to recognize that, for a while, multiple ontologies will exist and dominate. Therefore, translation tools (§2.9.2.3) are needed to handle their proliferation; they are also needed to achieve mutual understanding. Eventually, through consensus and standardization, a single ontology will appear.¹

Tools that make a difference: Wikipedia, a real and free encyclopedia of world knowledge (at a level deeper than CYC's intention), with more than a million articles in English and other natural languages. *Also:* Wordnet (§2.9).

2.9.2.3 Mapping one Ontology into Another. Some works [*150, *168] try to map every element (concept) of an ontology into the most similar concept residing in another ontology. These works address the lack of a unique ontology (§2.9.2.2).

2.9.2.4 Alignment. It is the superficial or initial mapping of nodes of an ontology into nodes of another ontology, conflicts being resolved by a user via a link editor [16].

2.9.2.5 Ontology Fusion. To fuse ontologies A and B is to find a new ontology C that contains the knowledge that both A and B contain. Contradictions and inconsistencies must be handled "as best as possible." A Ph. D. thesis in progress [4] tries to achieve fusion in automatic fashion, without intervention of a user.

2.9.3 Reasoning

The derivation of conclusions from certain premises using a given methodology. The two main deductive methods are: deductive reasoning and inductive reasoning.

Current work: Spatial reasoning, temporal reasoning. (Cf. IJCAI03).
{16% in IJCAI03}

2.9.3.1 Case-Based Reasoning. It can be defined as solving a new problem using the solution of a similar past problems. Typically, these are categorized into types or cases. Not a very active field now.

{1% in IJCAI03; 2% in IJCAI05}

2.9.3.2 Belief Revision. It is the change of beliefs to take into account new information. Current work: inconsistency detection, belief updating (Cf. IJCAI03).

2.9.4 Uncertainty

Probably this area will emerge as a new, self-standing part of AI. I have included it into Reasoning and Knowledge Representation because these are the two main problems in dealing with uncertainty: how to reason and compute about it, and how to represent it. See also Fuzzy Sets (§2.6.1).

Previous work: Dempster-Schafer Theory of evidence.

Current work: Inconsistency measurement. Paraconsistent logics. Measuring inconsistency using hierarchies [3]. Probabilistic inference.

{8% in IJCAI05}

¹ Nevertheless, knowledge can not be completely standardized, since each day more sprouts; standardization will always fall behind.

2.10 Robotics

Antecedent: remote manipulators.

Early days: Construction and adaptation of teleoperators, as well as simulations (for instance, simulating the path of a robot). Characterized by Japan's dominance. Slow addition of sensors, mainly vision and touch sensors. Scanty applications.

Recently, robots for rescue missions; for instance, IJCAI contest in 2003.

Current work: SLAM, Simultaneous localization and mapping. Coverage maps.

Robotics is still dominated by engineering designs. It may be an area waiting for applications to strengthen it.

{7% in IJCAI03; 3% in IJCAI05}

2.10.1 Teleoperators, Telemedicine

A teleoperator or remote manipulator is a device that (a) senses its environment, through a camera, perhaps; (b) can make changes to it, such as moving a tool or a piece, and (c) it is controlled by a person (an operator) at some distance from it. The perceptions in (a) go to the operator, which then issues the orders (b). This is different from a robot, in which a computer replaces the human operator.

2.11 Language Translation

Early times: Machine translation of a natural language into another. Failures were due to the inadequacy of the existing hardware and techniques, and to underestimation of the difficulties of the problem.

Later, work was more general than just translation. It was therefore called Natural Language, Intelligent Text Processing, or Natural Language Processing.

Information retrieval (§2.11.2) can be seen as an initial phase of Natural Language Processing.

Solved: To find the topics of themes that a document talks about [*99, *169].

Almost solved: Disambiguation, the assignment of meaning to words according to the context.

Almost solved: constructing a good parser for a natural language.

Still unsolved: translation of general texts in a natural language to another natural language has not been solved until today in a general and reasonable form. Constrained domain translators exist.

Tools that have made a difference: Wordnet (§2.9).

{6% in IJCAI03; 12% in IJCAI05}

2.11.1 Voice Recognition

It has taken distance from AI, and is now more properly considered a part of Signal Processing.

2.11.2 Information Retrieval

Information extraction. Traditionally, not a part of AI. Sometimes, it gets mixed with Search.

{2% in IJCAI03}

2.11.3 Semantic Processing

The work on intelligent text processing is also called Semantic Processing. It is a part of Natural Language.

2.11.3.1 Semantic Web. The Semantic Web was defined as pages with annotations, so that search engines and algorithms could “understand their meaning.” For manual placement of the annotations, SGML and XML (mark up languages) were designed. Unfortunately, the meaning of the names used in these marks are not standard, so a multiplicity of different marks arose. Also, the work needed to mark our own pages in the Web is not trivial. Thus, this approach has been largely abandoned, in favor of processing “raw documents” through Natural Language tools.

2.11.3.2 Text Mining. Since “Data Mining” became popular, the term “Text mining” was coined, but with a different meaning: the intelligent processing of (many) text documents. Hence, it is synonym of Semantic Processing.

Data Mining is defined as the automatic or semi-automatic finding of anomalies, tendencies, deviations and interesting facts in a sea of data. It is not considered part of AI (being more related to Data Base and to Statistics), although some commercial miners (Clementine) perform data mining with the help of neural nets, decision trees and clustering.

2.12 Constraint Satisfaction

Invented 35 years ago [*19] as Constraint Propagation, it survives to date. A technical field, with some applications.

Current work: Stochastic programming of constraints; consistency at the boundary.
{8% in IJCAI03; 13% in IJCAI05}

2.13 Search

With the proliferation of documents in the Web (see also §2.9), finding relevant texts has become important, a revival for Information Retrieval (§2.11). Programs that travel the Web looking for suitable information are called search engines or crawlers. They are combined with suitable text-processing tools (§2.11). One of the goals of “text understanding” is to be able to merge knowledge coming from different sources, for instance by merging ontologies (§2.9.2.5). Another goal is “to find answers by asking the right questions” [5].

Tools that make a difference: crawlers, Google.

{6% in IJCAI03; 8% in IJCAI05}

2.13.1 Planning

Initially, planning was an independent subject. I have merged it into Search, due to its similarity and small number of planning articles.

{5% in IJCAI05}

2.13.2 Search Engines, Semantic Search

A *Search Engine* is a software that finds information inside a computer, a private or local network, or in the Web. Search engines that comb the Web are also called

crawlers. Usually a Search Engine is given a predicate or a test that filters the information and only retrieves relevant data or documents (or their location in the Web). *Semantic Search* refers to the search performed with the help of filters that attend to the semantics or meaning of the information being analyzed.

Search and Semantic Search are very active areas; Semantic Search is considered a branch of AI because it uses Natural Language processing (§2.11), but soon it will emerge as a separate area.

2.14 Qualitative Physics

Also known as Naive Physics. It represents and reasons about the physical world in a common-sense, non-mathematical way. Qualitative Physics arises from the need to share our intuitions about the physical world with our machines [6]. It started and ended in the 1980's.

2.15 Lisp Machines

These were dedicated hardware for efficiently running Lisp programs. Commercial Lisp machines were Symbolics' 3600 (c. 1986), LMI-Lambda and TI-Explorer. Most Lisp machine manufacturers were out of business by the early 90's, due to (a) the appearance of (free) Kyoto Common Lisp running on SUN workstations, and (b) the appearance of cheap PCs that could run Lisp at good speed. This is an example of Gresham's law for special purpose-hardware: if you build a special purpose hardware, it should perform an order of magnitude better (or be an order of magnitude cheaper) than massive available general purpose hardware; otherwise, it will compete at a disadvantage.

2.15.1 Connection Machine

A parallel SIMD processor containing up to 65,536 individual processors (CM-2, CM-3 and CM-5; Thinking Machines, 1987). Not strictly a Lisp Machine, it ran *Lisp (parallel Lisp) [23].

2.15.2 AHR

A Mexican parallel computer built in 1980 [*47], of the MIMD shared-memory type, it had Lisp as its main programming language, and it consisted of up to 64 Z-80A's microprocessors. Subsequently, a Soviet computer [*56] of SIMD type was modified to mimic AHR's behavior. There are no further descendants of AHR.

2.16 Remarks and Conclusions

Artificial Intelligence deals with difficult problems, "problems that require intelligence to be solved." Thus, if AI solves one of these problems, in some sense "it is no longer difficult," hence that domain tends to leave the AI realm to stand in its own feet. Thus, AI will always be faced with "difficult and yet unsolved problems." That seems to be the fate of our discipline.

Advances in AI have been driven by two complementary forces, as in other areas of science. One is the "push" that provide new discoveries, theorems and theories,

such as the Resolution Principle or the invention of fast parallel hardware for chess machines. The other force is the “pull” that provide important practical problems that are still unsolved or only partially solved.

A particular feature of AI researchers that I have observed is that in general they are more inclined to use new tools (even if invented elsewhere), and I believe this produces better (or faster) advances, specially in applied problems.

Has AI produced significant applications? Has it any commercial value, or is just an academic endeavor? The question is posed to me sometimes. Certainly, some relevant applications exist: Expert Systems, visual inspection systems, and many commercial systems (such as those in data mining of large amounts of data) using neural networks and genetic algorithms, or fuzzy sets, to cite a few. More could have been produced by AI, if it were not for the fact that as a domain matures, it abandons AI (as my first remark says).

References

References are incomplete, due to space limitations. Some well-known works are not cited. A reference of the form [*47] is not listed here; it refers to article number 47, found in Adolfo Guzman’s Web page (<http://alum.mit.edu/www/aguzman>).

1. Bochman, A.: Propositional argumentation and causal reasoning. *Proc. IJCAI* (2005) 388-393
2. Chu, F. C., Halpern, J.: Great expectations. Part I: on the customizability of generalized expected utility. Part II: generalized expected utility as a universal decision role. *Proc. IJCAI03* (2003) 291-296 and 297-302
3. Contreras, A.: *Characterization and measurement of logical properties on qualitative values arranged in hierarchies*. Ph. D. thesis in progress. CIC-IPN, Mexico
4. Cuevas, A.: *Ontology Merging using semantic properties*. Ph. D. thesis in progress. CIC-IPN, Mexico
5. Etzioni, O., Cafarella, M., Downey, D., Popescu, A. M., Shaked, T., Soderland, Weld, D. S., Yates, A.: Methods for domain-independent information extraction from the web: An experimental comparison. In *Proceedings of the AAAI Conference* (2004) 391-398
6. Forbus, K. Qualitative physics: past, present and future. In *Exploring Artificial Intelligence*, H. Shrobe, Ed., Morgan Kauffmann, Los Altos, CA (1988) 239-290
7. Bernhard Ganter and Rudolf Wille. *Formal Concept Analysis*. Springer-Berlag. ISBN 3-540-8. Horn, B.K.P., Schunck, B.G.: Determining optical flow. *Artificial Intelligence*, vol 17 (1981) 185-203
8. Kirkpatrick, S., Gelatt, C.D., Vecchi, M.P.: Optimization by Simulated Annealing, *Science*, Vol 220, Number 4598 (1983) 671-680
9. Lenzerini M.: Data integration: a theoretical perspective., Proc. 21st. ACM SIGMOD-SIGACT-SIGART *Symposium on Principles of Database Systems* (2002) 233-246
10. Lucas, B. D., Kanade, T.: An iterative image registration technique with an application to stereo vision. *Proceedings of Imaging understanding workshop*, (1981) 121-130
11. McCarthy, J.: Recursive functions of symbolic expressions and their computation by machine, Part I, *Communications of the ACM*, v.3 n.4 (1960) 184-195
12. Minsky, M. L., Papert, S. A.: *Perceptrons*. MIT Press (1968)
13. Moses, J.: MACSYMA - the fifth year, *ACM SIGSAM Bulletin*, Vol.8 n.3 (1974) 105-110
14. Novak, J. D.: *A Theory of Education*. Cornell University Press, Ithaca, Illinois (1977)

15. Noy, N., Stuckenschmidt, H.: Ontology alignment: An annotated bibliography. In Y. Kalfoglou, M. Schorlemmer, A. Sheth, S. Staab, and M. Uschold, editors, *Semantic Interoperability and Integration*, number 04391 in Dagstuhl Seminar Proceedings. Internationales Begegnungs- und Forschungszentrum (IBFI), Schloss Dagstuhl, Germany, 2005. <<http://drops.dagstuhl.de/opus/volltexte/2005/48>>[date of citation: 2005-01-01]
16. Ontolingua. www.ksl.stanford.edu/software/ontolingua/
17. Roberts, L. G.: Machine perception of three-dimensional solids. In *Optical and Electrooptical Information Processing*, J. Tippett, Ed. MIT Press, Cambridge, Mass(1965) 125-143
18. Robinson, J A.: A Machine-Oriented Logic Based on the Resolution Principle, *Journal of the ACM*, Vol.12 n.1 (1965) 23-41
19. Sowa, J. F.: Conceptual Graphs. *IBM Journal of Research and Development* (1976)
20. Sowa, J. F.: Conceptual Structures. *Information Processing in Mind and Machine*. Addison Wesley (1984)
21. Steele, G., Sussman, G. J.: Revised Report on SCHEME, a Dialect of LISP, Memo 452, Artificial Intelligence Laboratory, M.I.T. (1978)
22. Trew. A., Wilson, G. (eds): *Past, Present, Parallel: A Survey of Available Parallel Computing Systems*. New York: Springer-Verlag (1991) ISBN 0-387-19664-1
23. Xu, L, Embley, D.: Combining the best of global-as-view and local-as-view for data integration. <http://www.deg.byu.edu/papers/PODS.integration.pdf>
24. Zadeh, L. A., *Fuzzy sets*. *Information and Control*, Vol. 8 (1965) 338-353