A Summary of the History of AI Before Computers

Technical Report UT-CS-07-589

Bruce J. MacLennan*

Department of Computer Science
University of Tennessee, Knoxville
www.cs.utk.edu/~mclennan

January 24, 2007

Abstract

This report addresses the work done on knowledge representation and automated reasoning before the invention of computers, including the work by Leibniz, Boole, and Jevons. We show how pre-computer work in formal logic, artificial languages, and automated reasoning devices contributed to the development of AI. Topics include scholastic logic; the art of memory; Lull's inference mechanisms; Leibniz' design of inferential calculi, knowledge representation methods, and calculating devices; Wilkins' design of a logically-structured language; Boole's investigations of the "laws of thought"; Jevons' construction of the "logical abacus" and "logical piano."

* This report an extended, unedited draft of an article for the *Encyclopedia of Information Science and Technology*, 2nd Ed., ed. by Mehdi Khosrow-Pour, Idea Group.

1 Introduction

The history of artificial intelligence (AI) is commonly supposed to begin with Turing's (1950) discussions of machine intelligence, and to have been defined as a field at the 1956 *Dartmouth Summer Research Project on Artificial Intelligence*. However, the ideas on which AI is based, and in particular those on which *symbolic* AI (see below) is based, have a very long history in the Western intellectual tradition, dating back, in fact, to ancient Greece. It is important for modern researchers to understand this history, for it reflects problematic assumptions about the nature of knowledge and cognition, and which, therefore, can impede the progress of AI.

2 Background

Symbolic AI is the approach to artificial intelligence that has dominated the field throughout most of its history and which remains important. It is based on the *Physical Symbol System Hypothesis*, enunciated by Newell and Simon (1976), which asserts, "A physical symbol system has the necessary and sufficient means for general intelligent action." In effect, it implies that knowledge is represented in the brain by language-like structures or formulas, and that thinking is a computational process that rearranges these structures according to formal rules. This view has also dominated cognitive science, which applies computational concepts to understanding human cognition (Gardner, 1985). (Alternative views will be mentioned at the end of this article.)

3 The Roots of Formal Logic

It is surprising, perhaps, that the original inspiration for symbolic knowledge representation can be found in ancient Greece, in particular in Pythagorean number theory (Burkert, 1972). In ancient Greece, as in many cultures, ancient and modern, pebbles were used for calculation by being moved in grooves in a similar way to the beads on an abacus. Indeed the Latin word for "pebble" is *calculus* and our word "calculate" comes from this manipulation of *calculi* (pebbles). In logic and mathematics we use the word "calculus" for any system of notation in which we can accomplish some purpose by the manipulation of tokens according to formal, game-like, mechanical rules. (For example we have the *differential* and *integral calculi* in mathematics and the *propositional* and *predicate calculi* in logic.) To the extent that the rules are *purely* mechanical, they can, in principle, be carried out by a machine, which is why calculi are important in AI; if a process can be reduced to a calculus, it can be *calculated* by a machine.

The ancient Pythagoreans (Pythagoras: 572–497 BCE) investigated number theory by means of arrangements of pebbles (Burkert, 1972, ch. VI). For example, they observed that certain numbers (1, 4, 9, 16, …) could be arranged into a square shape, and we still call these numbers "squares" today. However, they also investigated triangular numbers, as well as rectangles, pentagons, cubes, pyramids, etc. The pebbles were called "terms" (Grk., *horoi*, Lat., *termini*) — words that refer to boundary stones — and their arrangements were called *schemata* (Lat., *figurae*), terminology that is still used in logic. (This is the reason numbers are sometimes called "figures" and that we "figure things out" with them.) Although they did not prove theorems in the modern sense, they were able to demonstrate the truth of theorems in number theory by means of these arrangements. For example, they were able to show that each square number is a sum of consecutive odd numbers (4 = 3 + 1, 9 = 5 + 3 + 1, 16 = 7 + 5 + 3 + 1, etc.). Thus they discovered calculi could be used for reasoning as well as computation.

According to tradition, Pythagoras was the first to explain consonant musical intervals in terms of numerical ratios (Burkert, 1972, ch. V). For example, a string one-half the length of another string, sounds an octave higher; strings with the ratio 2:3 sound a fifth higher, and so forth. Thus, a subtle perceptual distinction (the consonance of pitches) could be rendered logical and rational by reducing it to numerical ratios (Grk., *logoi*, Lat., *rationes*, terms that also refer to the articulation of thought in words or symbols) (Maziarz & Greenwood, 1968, p. 43). It is an example of the representation of expertise in terms of formal structures; judgments of consonance could be replaced by calculation.

The Pythagoreans believed that everything could be reduced to numbers, and thus made intelligible: *rational*, *logical* (Burkert, 1972, ch. VI; Burnet, 1930, ch. II). Therefore they were committed to the idea that all knowledge could be represented in terms of arrangements of otherwise meaningless tokens, that is, in *formal structures* (and hence, we can conclude, in computer data structures).

Aristotle (384–322 BCE) is known, of course, as the originator of the science of logic, but two of his contributions in this area are especially relevant to AI. First, he began the development of *formal* logic by showing that valid inference could be distinguished from invalid inference of the basis its *form* rather than on the *meaning* of its particular terms (words). For example, the simple syllogism:

> All *M* are *P*.
> All *S* are *M*.
> Therefore, all *S* are *P*.

is valid regardless of what terms are substituted for the variables *S*, *M*, and *P*. (Indeed, Aristotle was the first one to use letters as variables in order to express formal rules of inference.) In other words, Aristotle showed that valid inference is a matter of *syntax* (the grammatical form of an argument) rather than *semantics* (its meaning). This is important because it shows how inference can be carried out by the manipulation of symbols independently of their meaning, which means that, in principle, inference is a kind of computation. In other words, there is a *calculus* of logic.

Aristotle analyzed the 192 possible syllogisms that can be constructed from three propositions, each being in one of four forms ('All *S* are *P*', 'No *S* are *P*', 'Some *S* are *P*', 'Some *S* are not *P*') arranged in three possible "figures" (depending on the position of the "middle" term *M*, which appears in both premises but not the conclusion), and he determined which are valid and which are not.

Aristotle also began the study of *modal logic*, that is, logic in which propositions are not simply true or false, but in which the propositions may be *possible*, *impossible*, *necessary*, or *contingent* (Bocheński, 1970, Pts. II, III; Kneale & Kneale, 1962, chs. II, III). Modal logic and its derivatives (such as *tense logic*, which deals with propositions whose truth values may change in time) are important in AI (Sowa, 1984, pp. 173–187).

Another contribution of Aristotle was to advocate the organization of knowledge into formal deductive structures, in which all the facts of a science were either stated as axioms or were formally derivable from the axioms. The best-known example of this is Euclidean geometry, which was the exemplar of a systematic body of knowledge for over two millennia (Maziarz & Greenwood, 1968, Pt. IV). Similar formal axiomatic structures are used in AI for representing a knowledge domain.

The investigation of logic continued over the following centuries. For example, during the Hellenistic period (third – first centuries BCE) Greek logicians continued the investigation of modal logic and developed tense logics and a logic of "qualitative probability" (in which propositions may be neither true nor false, but "convincing," etc.). They also invented *propositional logic*, for Aristotle's was a *class logic* (see **Terms and Definitions**, below) (Bocheński, 1970, Pts. II, III; Kneale & Kneale, 1962, chs. II, III); both are used in AI.

The Medieval Scholastics (roughly 6th-15th centuries) refined logic into a very precise instrument, although it was still based on natural language (Latin), in contrast to modern symbolic logic. As a consequence they became conscious of the limitations of natural language for exact knowledge representation and strove to compensate for its deficiencies. For example, they knew that the word 'dog' is used differently in the propositions 'a dog is a mammal' and 'dog is a noun'. AI knowledge representation languages have to deal with similar issues (Sowa, 1984, p. 84). In the end, dissatisfaction with natural languages led to an interest in developing artificial languages that were intended to be more "rational" (logical and precise). Behind this was the assumption that there is a universal grammar underlying all natural languages, and that it corresponds to the "language of thought"; therefore an artificial language, as an ideal vehicle for thought, ought to reflect this "deep structure." Similar motivations underlie the development of AI "knowledge representation languages" (see below).

4 Combinatorial Methods

The Middle Ages also saw the development of combinatorial approaches to solving problems (Bocheński, 1970, Pt. III). For example, the letters A, E, I, and O were used to stand for the four different types of propositions mentioned above. Since each Aristotelian syllogism comprises three propositions, the different forms of the syllogism could be represented by the 64 triples, AAA, AAE, AAI, …, OOI, OOO. Since there are three possible positions for the middle term, there are 192 possible Aristotelian syllogisms. The Medieval Scholastics used a combinatorial procedure to generate them all, and then they crossed out the invalid ones. This is an example of a *generate-and-test procedure*, an approach still widely used in AI. For example, a game-playing program might generate all possible moves and then eliminate those that are illegal, lose the game, or are weak. The problem with generate-and-test procedures is *combinatorial explosion*: the number of combinations to be tested increases exponentially with their size.

These combinatorial procedures acquired an increased significance, which contributed to the eventual development of AI, from the *kabbalah*, a Jewish mystical tradition with Pythagorean affinities, which became popular in the Middle Ages (Eco, 1997, ch. 1; Scholem, 1960, p. 167). According to this tradition, the *Torah* reflects the *logos* (rational structure) of the universe. Therefore, since the *Torah* is written in the letters of the Hebrew alphabet, these letters correspond to the elementary categories and archetypal forms underlying the universe. As a consequence the letters of the Hebrew words for things reveal their logical structure to one who knows how to interpret them. Combinatory processes figure prominently in kabbalah, and significant words, especially the "names of God," were permuted in order to reveal hidden wisdom and discover new truths. For this purpose the kabbalists used rotatable wheels and other devices to ensure that they did not miss any combinations of letters, an example of a mechanized generate-and-test procedure.

Similar in spirit to the kabbalah, and perhaps in part inspired by it, was the "Great Art" (*Ars Magna*) of Raymond Lull (also spelled "Llull," 1232–1315) (Johnston, 1987; Llull, 1985; Yates, 1966). He intended it to be a "universal science of all sciences," a systematic method by which knowledge could be discovered and proved. There were several versions of his system, but the most common one made use of nine "divine dignities," or attributes of God, which took different forms in each domain of knowledge, but provided the fundamental categories in each domain. These abstract qualities (Goodness, Magnitude, Duration, etc.) correspond closely to certain kabbalistic names of God. In Lull's Art, as in kabbalah, we see an attempt to isolate the most basic categories that constitute all knowledge and to discover, therefore, an "alphabet of thought." This remains an important goal in contemporary symbolic AI.

A distinctive characteristic of Lull's Art was the extensive use of rotatable wheels to generate combinations of these elementary categories in order to discover and to demonstrate philosophical truths (primarily theological assertions, by which he hoped to convert non-believers to Christianity). Thus the Great Art combines an alphabet of elementary concepts with mechanical procedures for generating their combinations into an automated method of knowledge discovery and proof.

Such, at least, was its goal. In fact, it didn't work, and for the most part it could be used only for "proving" the theological propositions that the operator already believed. Nevertheless, as we'll explain, it inspired many later thinkers to attempt to correct its deficiencies and to construct machines for knowledge discovery and inference, but first it had to be recast into a more logical form.

A step in a more logical direction was made as a part of the educational reforms of Peter Ramus (1515–72), which stressed the organization of knowledge into class hierarchies (an idea rooted in ancient philosophy and stemming ultimately from Aristotle). Ramean trees became very popular, along with other techniques for organizing knowledge into geometrical structures (e.g., ladders, towers, circles). Hierarchical class organization has been important in AI knowledge representation and in object-oriented programming systems, but it has been found to be too restrictive, and so newer systems permit nonhierarchical classification.

Further, Ramean trees were required to be *dichotomies*, that is, hierarchies in which each class was divided into two disjoint, mutually-exclusive subclasses, a *binary* classification (e.g., Fig. 1). These trees, when they resulted from a valid logical analysis, were supposed to reflect the actual structure of reality. Therefore, correct definitions could be read from the binary trees. For example, according to the tree, a man is a rational animal, and an animal is a sensible living being, a plant is an insensible living being, etc. A particular class (man, plant, etc.) corresponds to the binary string that described the path from the root (the *summum genus*, or highest class) to the class in question. This provided a means for using a binary string to represent a class or concept in terms of a logical analysis of its meaning (its *intension*; see below); in modern terminology, the concept is represented by a *binary feature vector*.

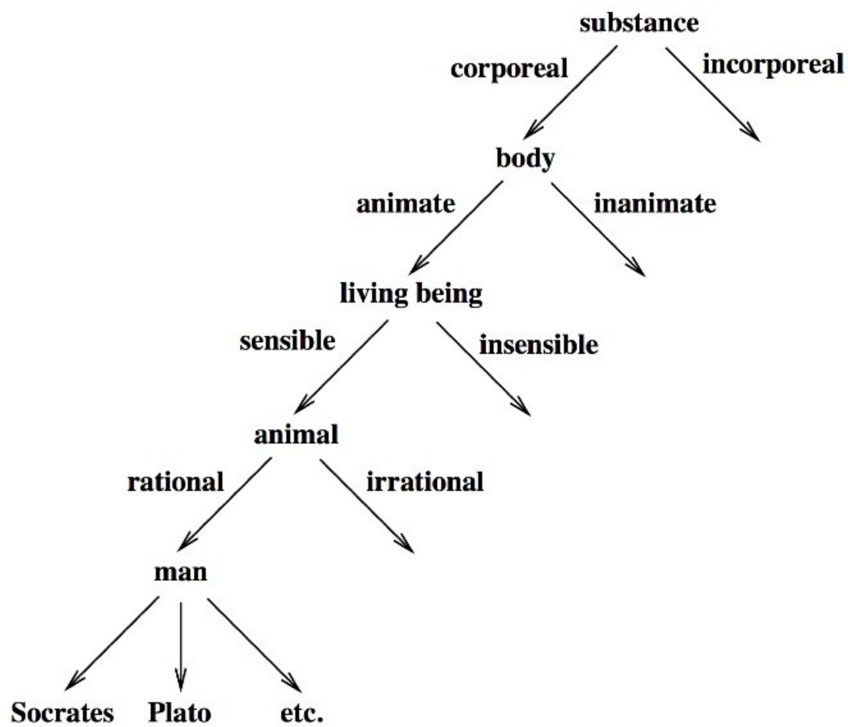

Fig. 1: Fragment of Typical Ramean Tree.

5 Knowledge Representation and Mechanized Inference

As knowledge and inference became more systematized, the idea developed that reasoning, when carefully and methodically executed, was a kind of calculation. One clear exponent of this view was Thomas Hobbes (1588–1679), who said, “By *ratiocination* I mean *computation*” (*Elem. Phil.*, 1.1.1.2). He made clear, however, that the “addition” and “subtraction” of concepts was not the same as addition and subtraction of numbers, for the former is a logical process, the latter, quantitative.

Hobbes also distinguished reasoning from causes to their effects (“forward chaining,” in modern AI terminology) and reasoning from effects to their causes (“backward chaining”). In both cases thought is a kind of mental discourse, which corresponds to a defining assumption of symbolic AI: that there is a *language of thought* (sometimes called “Mentalese”). Words, whether external or in the mind, are tokens, manipulated according to mechanical rules, and correct reasoning is analogized to balancing account books. That is, there is a *calculus* of thought. Furthermore, since Hobbes was a complete materialist, he understood thought as a kind of matter in motion, a strictly mechanical process.

Over the centuries there have been many attempts to design “ideal languages,” that is, artificial languages without the perceived deficiencies of natural languages (Eco, 1997; Large, 1985). As modern science emerged in the seventeenth century, the goal was often to develop a “philosophical language,” that is, a language suitable for philosophical analysis and scientific discourse. One of the most famous of these projects was the Real Character of John Wilkins (1614–72) (Large, 1985; Rossi, 2000; Vickers, 1987, ch. 9). He began by isolating a “universal grammar” that he believed to underlie the particular grammars of all natural languages, and so it is in effect the grammar of “Mentalese” and hence reflects the laws of thought. Inspired by Chinese writing, Wilkins also concluded that the forms of words should reflect their logical analysis (based on a class hierarchy), and he designed a vocabulary and symbolic writing system based on a comprehensive conceptual taxonomy. His language had little direct impact, beyond inspiring the conceptual taxonomy used by *Roget’s Thesaurus*, but symbolic logic and AI knowledge representation languages have similar goals and approaches (and, arguably, similar failings).

Gottfried Wilhelm von Leibniz (1646–1716) made many contributions to philosophy and mathematics, but here we are concerned only with his experiments in knowledge representation and mechanized reasoning (Coudert, 1995; Kneale & Kneale, 1962; Styazhkin, 1969). His work in this area was influenced by kabbalah, Lull, Wilkins’s language, Hobbes, and Chinese writing and philosophy. For example, although he had already invented the binary number system, he later found the binary system in the Chinese *I Ching* (*Book of Changes*) and saw how it reduced all change in the universe to two opposites (yin and yang). This accorded with the kabbalistic and Lullian idea that the world was organized in terms of an alphabet of fundamental ideas (the “divine emanations”) and with his own rationalistic philosophy, which sought the true essences of concepts in a small number of atomic (indivisible) categories.

Leibniz was very impressed by Lull’s Great Art and by Wilkins’ Real Character, but concluded that they would not work, and so he constructed a number of knowledge representation schemes on a more logical plan. In the process he discovered an important relationship between numbers and concepts. According to the Fundamental Theorem of Arithmetic, any positive integer can be decomposed into a unique product of prime numbers (e.g., 12 is the product of 2, 2, and 3), which corresponded to the rationalist idea that any concept could be reduced to a unique conjunction of atomic (elementary) concepts. Therefore, if a prime number were assigned to each atomic concept, then every possible concept would have a unique numerical value. Conversely, if we looked up in a “philosophical dictionary” the number corresponding to any concept, then we could discover its essence, or true definition, by reducing the number into its prime factors. For the sake of an example, suppose that ‘animal’ and ‘rational’ are elementary concepts (Leibniz would not have considered them such) and that 'man is the rational animal' is a correct definition. Further suppose that 2 and 3 are the prime numbers assigned to 'animal' and 'rational' respectively; then 6 would be the number for the concept 'man'. If we did not know the definition of 'man', then we could discover it from its number, for 6 = 2×3, and therefore man is the rational animal.

There are two ways that classes are treated in mathematics and logic, *extensionally* and *intensionally*. The extensional approach, which is the most familiar, is to define a class in terms of its members, its *extension*. Thus the extension of the class 'man' (meaning human) includes Leibniz, Aristotle, Hypatia, and all the rest of us. The other way to define a class is in terms of its *intension*, that is, its essential attributes (although there are various notions of intension). For example, the intension of 'man' could be the attributes 'rational' and 'animal'. Although modern logic and mathematics tend to treat classes extensionally, AI treats them intensionally (i.e., a concept is represented by a "property list") for the simple reason that most concepts have small intensions (e.g., 'rational animal') but infinite extensions, so it is easier to compute with intensions. For the same reasons, Leibniz settled on an intensional representation.

Leibniz agreed with Hobbes' assertion that thought is computation, and worked on a calculus for logical inference. For example, he discovered that propositions of the form 'all *S* are *P*' can be decided computationally if we know the numbers corresponding to *S* and *P*. For if all *S* are *P*, then the essential attributes of *P* are among the essential attributes of *S*; numerically, the prime factors of *P* are among the prime factors of *S*. Therefore, to decide if a proposition 'all *S* are *P*' is true, all we need to do is to look up the numbers for *P* and *S* and see if the number for *P* evenly divides the number for *S*. Leibniz investigated similar computational approaches to deciding propositions of the other forms.

In summary, we can see that Leibniz had all the components of a system of knowledge representation and mechanical inference. In principle, all concepts could be analyzed into a relatively small number of elementary atomic concepts, and each concept could be assigned a unique number on the basis of this analysis. All philosophical questions, then, could be answered rationally and logically by calculation, literally by ratios (*rationes*, *logoi*). Indeed, Leibniz constructed one of the earliest digital calculating machines (1671), the first capable of multiplication and division, and so he had in principle (but not capacity) the means for actual mechanized reasoning.

George Boole (1815–64) is well known, of course, to computer scientists and information technologists as the inventor of Boolean algebra, which is applied to digital circuit design and in many other ways in computer technology. However, his goals were much more far-reaching, and in his *Investigation of the Laws of Thought* he says his goal is "to investigate the fundamental laws of those operations of the mind by which reasoning is performed" and to express them in a calculus (Boole, 1854, p. 1). In common with contemporary logicians, such as Augustus De Morgan (1806–71), he expressed logical operations in an algebraic notation, as opposed to a natural language, thus contributing to the development of *symbolic logic*. He developed an extensional class logic, in which operations on classes correspond to operations on their extensions, that is, on the sets of their members, and so he invented the algebra of sets. In this way he influenced the primarily extensional approach of modern set theory and predicate logic. However, he also showed how the same algebraic operations could be interpreted as a propositional logic, which laid the foundation for Boolean circuit design, later developed by Claude Shannon (1938), the inventor of information theory. Boole stressed the formality of his logic, that is, that its rules of inference depended only on the algebraic properties of the operators (commutativity, associativity, etc.) and not on any interpretation of the terms. Therefore, these operation were not restricted to human thought, but could be implemented by machines, which was accomplished about a decade later by W.S. Jevons.

William Stanley Jevons (1835–82) was a prolific mathematician, scientist, and philosopher, who contributed to statistics, economics, meteorology, and the philosophy of science (Mays & Henry, 1953). However, he was also the first to construct fully functional *logic machines*, capable of automated reasoning (Jevons, 1870; Jevons, 1894, pp. 196–201; Jevons, 1958, pp. 91–96, 104–114), and thus predecessors of AI technology. His system is based, first, on the idea of a *logical alphabet*, which lists all the possible conjunctions of a given set of terms and their complements. For example, a 3-term logical alphabet lists the eight possibilities generated by *A* and non-*A*, *B* and non-*B*, and *C* and non-*C*; in modern terms is corresponds to the eight possible 3-bit strings. (He connects this idea to the Ramean Tree.) Second, he uses an "indirect method" of deduction, which is simply to eliminate from the logical alphabet those combinations that are inconsistent with the premises. Obviously, this is a generate-and-test procedure: list all the possible combinations and remove the impossible conclusions; the result is the broadest conclusion compatible with the premises, which he called the "complete solution."

Jevons' indirect method, like most generate-and-test procedures, is tedious and error-prone to perform manually. Therefore he invented a succession of devices that increasingly automated the process. One of these, called the *logical abacus*, made use of a set of wooden cards, one for each combination in a 2-, 3- or 4-term alphabet. Pins in different positions on a card represented whether a term (e.g., *A*) or its complement (e.g., non-*A*) appeared in the card's combination. The abacus itself was an inclined surface with four ledges capable of holding the cards. The operator used a metal straight-edge to lift all the cards containing a particular term (e.g., non-*B*) from one ledge and to move them to another. By a series of complicated but mechanical procedures, involving removing sets of cards from the evolving solution, or reintroducing them, the operator was able to calculate the complete solution of the problem. (In modern terminology, he was manipulating propositions in *disjunctive normal form*.)

Jevons' most sophisticated logic machine was a completely mechanical device, which he called the *logical piano*. It had a keyboard marked with the terms (*A*, *B*, *C*, *D* and their complements) and with various logical symbols (e.g., equality, inclusive-or, "finis"), which was used for entering a series of logical equations representing the premises of a deduction. Above the keyboard was a (mechanical) display, a kind of spreadsheet, which represented all of the logical combinations consistent with the premises that had been entered so far. Thus the operator could watch the developing mathematical analysis, and even try out hypothetical premises to see how they might affect the conclusion. In 1869 Jevons constructed and demonstrated a 4-term machine and planned the development of a 10-term reasoning engine, which would have required an entire wall to display the 1024 combinations of its logical alphabet. Although the machine performs relatively simple operations on bit strings, Jevons enthused that "After the Finis key has been used the machine represents a mind endowed with powers of thought," and that as each proposition is entered "the machine analyses and digests the meaning of it and becomes charged with the knowledge embodied in that proposition" (Jevons, 1958, pp. 110–111). Thus Jevons invented "AI hype"!

6 Future Trends

In the light of this history, symbolic AI, which has dominated AI research, can be seen as the continuation of a centuries-old tradition concerning the nature of knowledge and inference. This, of course, does not imply that it is the best approach to AI, or conversely that it is not. Although some prominent researchers have declared that the symbolic approach to cognition is "the only game in town," there are alternatives, most notably *connectionism* (or *parallel distributed processing*), which is based on simplified models of neural networks in the brain. This new approach promises to compensate for many of the limitations of the symbolic approach, and also to shed light on cognitive processes in the brains of humans and other animals. Connectionism, however, is beyond the scope of this article.

This article has focused on a few of the principle thinkers who contributed to AI before the era of the computer, but there are many others. A more extensive treatment would discuss some of these other contributors, but other thinkers have not been considered in their relation to AI research. Therefore, much work remains to be done in exploring and explaining the intellectual background of artificial intelligence.

7 Conclusions

We may draw several conclusions from this historical survey. First, symbolic AI is built upon a foundation of philosophical and psychological premises that have been part of Western intellectual history since ancient Greece. Since these assumptions are so deeply embedded in our intellectual background, they easily may be taken for granted and escape adequate scrutiny. However, alternatives, such as connectionism, are being explored. Second, although earlier philosophers discussed the idea that thought is a kind of computation, it was only with the advent of modern computers that there was sufficient computing power to test these theories empirically. As a consequence, experimental AI research in the late twentieth century revealed both the capabilities and limitations of symbolic AI and motivated the search for alternatives. Finally, perhaps the most important conclusion that we can draw is that AI is not an isolated technological discipline, nor simply the applied side of cognitive science, but that it is intimately related to intellectual issues about the mind that have occupied civilization for millennia.

8 References

Barnes, J. (1982). *Aristotle*. Oxford, UK: Oxford University Press.

Bocheński, I.M. (1970). *A history of formal logic*. New York, NY: Chelsea Publishing.

Boole, G. (1854). *An investigation of the laws of thought*. New York, NY: Macmillan.

Burkert, W. (1972). *Lore and science in ancient Pythagoreanism*. Cambridge, MA: Harvard University Press.

Burnet, J. (1930). *Greek philosophy part I: Thales to Plato*. London, UK: Macmillan.

Coudert, A.P. (1995). *Leibniz and the kabbalah*. Dordrecht: Kluwer.

Eco, U. (1997). *The search for the perfect language*, transl. J. Fentress. Cambridge, MA: Blackwell.

Gardner, H. (1985). *The mind's new science: A history of the cognitive revolution*. New York, NY: Basic Books.

Jevons, W.S. (1870). On the mechanical performance of logical inference. *Philosophical Transactions, 160*, 497.

Jevons, W.S. (1894). *Elementary lessons in logic: Deductive and inductive*. New York, NY: Macmillan.

Jevons, W.S. (1958). *The principles of science: A treatise on logic and scientific method*, second ed. New York, NY: Dover.

Johnston, M.D. (1987). *The spiritual logic of Ramon Llull*. Oxford, UK: Clarendon Press.

Kneale, W., & Kneale, M. (1962). *The development of logic*. Oxford, UK: Oxford University Press.

Large, A. (1985). *The artificial language movement*. New York, NY: Basil Blackwell.

Llull, R. (1985). *Selected works of Ramon Llull (1232–1316), vol. I*, edited & transl. by Anthony Bonner. Princeton, NJ: Princeton University Press.

Mays, M., & Henry, D.P. (1953). Jevons and logic. *Mind, 62*, p. 484.

Maziarz, E.A., & Greenwood, T. (1968). *Greek mathematical philosophy*. New York, NY: Frederick Ungar.

Newell, A., & Simon, H.A. (1976). Computer science as empirical inquiry: Symbols and search. *Communications of the Association for Computing Machinery, 19*, 113–126.

Rossi, P. (2000). *Logic and the art of memory: The quest for a universal language*, transl. & intro., S. Clucas. Chicago, IL: University of Chicago Press.

Scholem, G. (1960). *On the kabbalah and its symbolism*. New York, NY: Schocken.

Shannon, C.E. (1938). A symbolic analysis of relay and switching circuits. *Transactions of the AIEE, 57*, 713–723.

Sowa, J.F. (1984). *Conceptual structures: Information processing in mind and machine*. Reading, MA: Addison-Wesley.

Styazhkin, N.I. (1969). *History of mathematical logic from Leibniz to Peano*. Cambridge MA: MIT Press.

Turing, A.M. (1950). Computing machinery and intelligence. *Mind, 59*, 433–60.

Vickers, B. (1987). *English science, Bacon to Newton*. Cambridge, UK: Cambridge University Press.

Yates, F.A. (1966). *The art of memory*. Chicago IL, University of Chicago Press.

9 Terms and Definitions

Calculus: A system of physical symbols and mechanical rules for their manipulation intended to accomplish some purpose, such as calculation, differentiation, integration, or formal inference. In principle, any process that can be accomplished by a calculus can be programmed on a digital computer.

Class Logic: A logic in which the terms refer to classes (sets, concepts, categories) and in which the logical operations and relations deal with classes (union, intersection, equality, inclusion, mutual exclusion, etc.). Contrasted with a *propositional logic* (q.v.). Mathematically, class logic is similar to the algebra of sets.

Disjunctive Normal Form: A special form into which any formula of propositional logic may be put. The form is a disjunction (*inclusive-or*-ings) of conjunctions (*and*-ings) of "literals," each of which is a letter (a *primitive proposition*) or its negation. Putting propositions into disjunctive normal form (DNF) facilitates their manipulation by computer. *Conjunctive normal form* is defined analogously.

Epistemology: The philosophical discipline devoted to the study of knowledge, including its nature and the means by which it may be acquired and validated.

Extension: The set of all individuals to which a general term applies, or that is included in a concept or class. For example, the extension of 'person' is the infinite set of all particular people. Contrasted with *intension* (q.v.).

Formal Logic: A system of logic in which the validity of arguments can be determined from their form and independently of the meaning of the terms in the propositions. That is, validity is a matter of *syntax* (q.v.) rather than *semantics* (q.v.).

Generate-and-Test Procedure: A common method of search, used in AI and other applications, in which possible solutions are generated systematically, and evaluated until a suitable solution is found. For example, a game-playing program might generate possible moves, which are evaluated in terms of their likelihood of leading to a win. The greatest weakness of generate-and-test procedures is *combinatorial explosion*, which refers to the exponential increase of the number of possible solutions of increasing complexity (e.g., the number of moves that a game-playing program looks forward).

Intension: The set of all properties necessarily inhering in a concept, class, or general term. For example, the intension of 'person' includes such properties as 'rational', 'bipedal', and 'featherless'. Sometimes the intension of a term is restricted to just its essential properties (those that are part of its definition), a finite set. Contrasted with *extension* (q.v.). (Note that "intension" is spelled with an "s" and is a different word from "intention.")

Knowledge Representation Language: A formal language, implementable in the data structures of a digital computer, intended to be capable of representing all knowledge, or at least all knowledge in some AI application domain. It is intended as a medium for storing knowledge and for mechanized inference in its domain. A knowledge representation language is the analogue in AI of the “language of thought” (q.v.) in cognitive science.

Language of Thought (“Mentalese”): A hypothesized language-like system in terms of which all human cognition is supposed to take place. Advocates of this hypothesis acknowledge that not all our thinking is *discursive* (by means of an inner dialog), but they argue that the systematic structure of ideas and thinking imply that there must be a language of thought, albeit below the level of conscious access. The language of thought hypothesis partly justifies symbolic AI (q.v.) as a sufficient basis for AI.

Logical Atomism: The view, especially advocated by Bertrand Russell and the early Ludwig Wittgenstein, that meaning can be analyzed into certain atomic (indivisible) units of meaning (atomic facts); and conversely that any intelligible meaning can be expressed in terms of elementary ideas.

Predicate Logic: An extension of *propositional logic* (q.v.) that includes variables referring to the individuals of some domain, quantifiers (“for all” and “for some”) over these variables, and predicates (classes, sets) defined over these individuals.

Propositional Logic: A logic in which the terms refer to propositions (statements that are true or false) and in which the logical operations and relations deal with propositions (conjunction, disjunction, implication, equivalence, etc.). Contrasted with *class logic* and *predicate logic* (q.vv.).

Semantics: Refers to the *meanings* of expressions in a natural or artificial language and to the study of these meanings and their relation to the expressions. Often contrasted with *syntax* (q.v.). Since formal systems, calculi, and symbolic AI systems deal only with the forms of expressions, they can be sensitive to semantics only to the extent that the semantics is explicit in the system’s syntax.

Symbolic AI: An approach to AI based on the manipulation of knowledge represented in language-like (“symbolic”) structures in which all relevant semantics (meaning) is explicit in the syntax (formal structure). The language of thought hypothesis (q.v.) provides part of the justification of the sufficiency of the symbolic approach to AI.

Symbolic Logic: A formal logic with a mathematical notation and algebraic rules of manipulation and inference.

Syntax: Refers primarily to the grammar rules of a language (natural or artificial), that is, to the allowable *forms* of expressions without reference to their meaning (*semantics*, q.v.). In the context of AI, syntax refers to the rules of knowledge representation in terms of data structures and to the computational processes that operate on these structures.