The Challenge of Knowledge Soup

John F. Sowa

VivoMind Intelligence, Inc.

> *Human knowledge is a process of approximation. In the focus of experience, there is comparative clarity. But the discrimination of this clarity leads into the penumbral background. There are always questions left over. The problem is to discriminate exactly what we know vaguely.*
>
> Alfred North Whitehead, *Essays in Science and Philosophy*

Abstract. People have a natural desire to organize, classify, label, and define the things, events, and patterns of their daily lives. But their best-laid plans are overwhelmed by the inevitable change, growth, innovation, progress, evolution, diversity, and entropy. These rapid changes, which create difficulties for people, are far more disruptive for the fragile databases and knowledge bases in computer systems. The term *knowledge soup* better characterizes the fluid, dynamically changing nature of the information that people learn, reason about, act upon, and communicate. This article addresses the complexity of the knowledge soup, the problems it poses for computer systems, and the methods for managing it. The most important requirement for any intelligent system is flexibility in accommodating and making sense of the knowledge soup.

Presented at the Episteme-1 Conference in Goa, India, in December 2004. This version was published in *Research Trends in Science, Technology and Mathematics Education*, edited by J. Ramadas & S. Chunawala, Homi Bhabha Centre, Mumbai, 2006.

1. Issues in Knowledge Representation

The reasoning ability of the human brain is unlike anything implemented in computer systems. A five-dollar pocket calculator can outperform any human on long division, but many tasks that are easy for people and other animals are surprisingly difficult for computers. Robots can assemble precisely machined parts with far greater accuracy than any human, but no robot can build a bird nest from scattered twigs and straw or wash irregularly shaped pots, pans, and dishes the way people do. For recognizing irregular patterns, the perceptual abilities of birds and mammals surpass the fastest supercomputers. The rules of chess are defined with mathematical precision, but the computers of the 1960s were not fast enough to analyze chess patterns at the level of a novice. Not until 1997 did the world chess champion lose to a supercomputer supplemented with special hardware designed to represent chess patterns. The rules and moves of the oriental game of Go are even simpler than chess, but no computer can play Go beyond the novice level. The difference between chess and Go lies in the nature of the patterns: chess combinations can be analyzed in depth by the brute force of a supercomputer, but Go requires the ability to perceive visual patterns formed by dozens of stones placed on a 19×19 board.

The nature of the knowledge stored in people's heads has major implications for both education and artificial intelligence. Both fields organize knowledge in teachable modules that are axiomatized in logic, presented in textbooks, and stored in well structured databases and knowledge bases. A systematic organization makes knowledge easier to teach and to implement in computer systems. But as every student learns upon entering the workforce, “book learning” is limited by the inevitable complexities, exceptions, and ambiguities of engineering, business, politics, and life. Although precise definitions and specifications are essential for solving problems in mathematics, science, and engineering, most problems aren't well defined. As Hamlet observed, “There are more things in heaven and earth, Horatio, than are dreamt of in your philosophy.”

The knowledge soup poses a major challenge to any system of organizing knowledge for ease of learning by people or ease of programming in computers. Section 2 of this article surveys attempts to develop such systems. Section 3 discusses the inevitable exceptions and disruptions that cause well organized systems to degenerate into knowledge soup. As a framework for accommodating the complexity, managing it, and even taking advantage of it, Section 4 presents some issues in cognitive science and the semiotics of Charles Sanders Peirce. His insights into both the power and the limitations of logic suggest methods for addressing the challenge and designing more adaptable and ultimately more human-like systems. The concluding Section 5 puts the issues in perspective and proposes directions for future research.

2. Attempts to Organize and Formalize Knowledge

For over two thousand years, Aristotle's categories and his system of syllogisms for reasoning about the categories were the most highly developed system of logic and ontology. The syllogisms are rules of reasoning based on four sentence patterns, each of which relates one category in the subject to another category in the predicate:

1. *Universal affirmative*. Every employee is human.
2. *Particular affirmative*. Some employees are customers.
3. *Universal negative*. No employee is a competitor.
4. *Particular negative*. Some customers are not employees.

The two affirmative patterns are the basis for inheriting properties from more general categories to more specialized ones. The two negative patterns state constraints that rule out combinations that are not meaningful or permissible.

In the third century AD, Porphyry drew the first known tree diagram for organizing Aristotle's categories according to the method of definition by *genus* and *differentiae*. Figure 1 shows a version translated from the *Summulae Logicales* by Peter of Spain (1239). It shows that the category Body is defined as the category Substance with the differentia material, and Human is defined as Animal with the differentia rational. By following the path from the top, the category Human would inherit all the differentiae along the way: rational, sensitive, animate, material Substance.

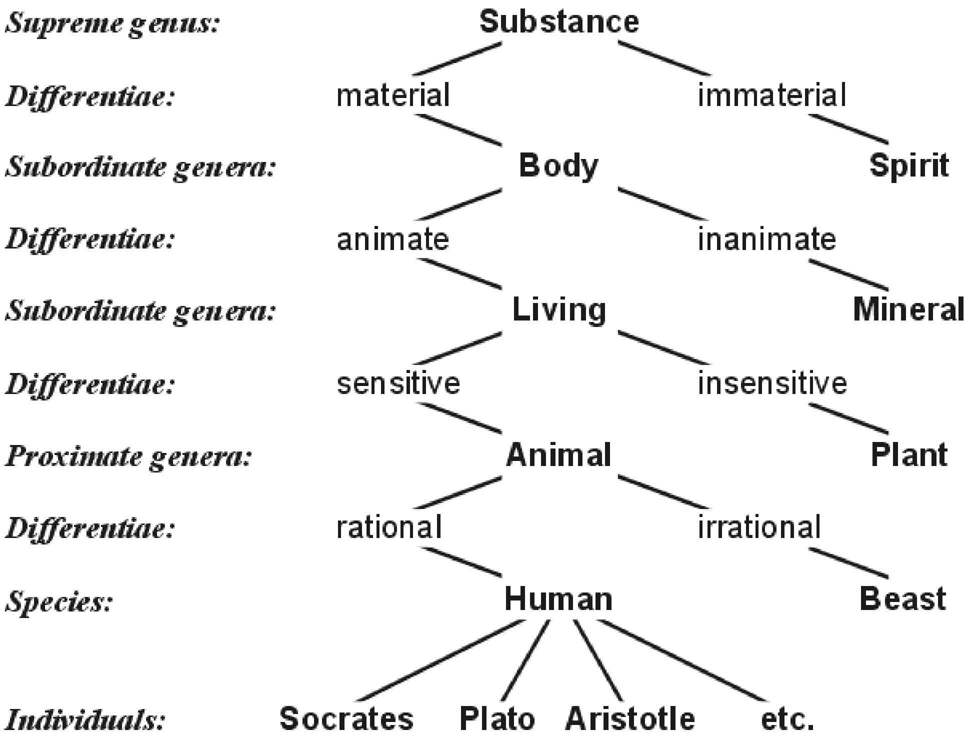

Figure 1: Tree of Porphyry

Similar tree diagrams are widely used today to represent hierarchies of concept types in modern knowledge representation languages. Although Aristotle's syllogisms are the oldest system of formal logic, they form the core of modern *description logics* (DLs), such as OWL, which are often used for defining ontologies. OWL and other DLs add important features, such as numeric-valued functions, to Aristotle's monadic predicates. For many applications, however, the Aristotelian subset of logic serves as the framework that supports all the rest.

In the 17th century, Latin was losing its status as the common literary and scientific language of Europe. To avoid fragmentation in a Babel of mutually unintelligible languages, various schemes were proposed for a universal language based on Aristotle's logic and categories. Scientists as renowned as Descartes, Mersenne, Boyle, Newton, and Leibniz devoted some attention to the project (Knowlson 1975). The idea was even satirized by Jonathan Swift as one of the projects at the grand academy of Laputa in *Gulliver's Travels*. In a scheme that resembled the sentence-generating machine in Laputa, Leibniz (1666) hoped to automate Aristotle's syllogisms by encoding the categories as integers:

> The only way to rectify our reasonings is to make them as tangible as those of the Mathematicians, so that we can find our error at a glance, and when there are disputes among persons, we can simply say: Let us calculate, without further ado, in order to see who is right.

Leibniz used prime numbers to encode primitive concepts and products of primes to encode compound concepts: if 2 represents Substance, 3 material, and 5 immaterial, the product 2×3 would represent Body and 2×5 would represent Spirit. A sentence of the form *Every human is animate* would be true if the number for *human* is divisible by the number for *animate*. This method works well for reasoning about affirmative propositions, but Leibniz never found a satisfactory method for handling negation. Although he abandoned his early work on the project, Leibniz (1705) still believed in the importance of developing a hierarchy of categories:

> The art of ranking things in genera and species is of no small importance and very much assists our judgment as well as our memory. You know how much it matters in botany, not to mention animals and other substances, or again moral and notional entities as some call them. Order largely depends on it, and many good authors write in such a way that their whole account could be divided and subdivided according to a procedure related to genera and species. This helps one not merely to retain things, but also to find them. And those who have laid out all sorts of notions under certain headings or categories have done something very useful.

In his *Critique of Pure Reason*, Immanuel Kant adopted Aristotle's logic, but he proposed a new table of twelve categories, which he claimed were more fundamental than Aristotle's. Although he started with a new choice of categories, Kant was no more successful than Leibniz in completing the grand scheme:

> If one has the original and primitive concepts, *it is easy* to add the derivative and subsidiary, and thus give a complete picture of the family tree of the pure understanding. Since at present, I am concerned not with the completeness of the system, but only with the principles to be followed, I leave this supplementary work for another occasion. It can *easily* be carried out with the aid of the ontological manuals.

Note the added italics: whenever a philosopher or a mathematician says that something is easy, that is a sure sign of difficulty. No one ever completed Kant's “supplementary work.”

With the advent of computers, the production and dissemination of information was accelerated, but ironically, communication became more difficult. When product catalogs were printed on paper, an engineer could compare products from different vendors, even though they used different formats and terminology. But when everything is computerized, customer and vendor systems cannot interoperate unless their formats are identical. This problem was recognized as soon as the first database systems were interconnected in the 1970s. To enable data sharing by multiple applications, a *three-schema approach*, illustrated in Figure 2, was proposed as a standard for relating a common semantics to multiple formats (Tsichritzis & Klug 1978).

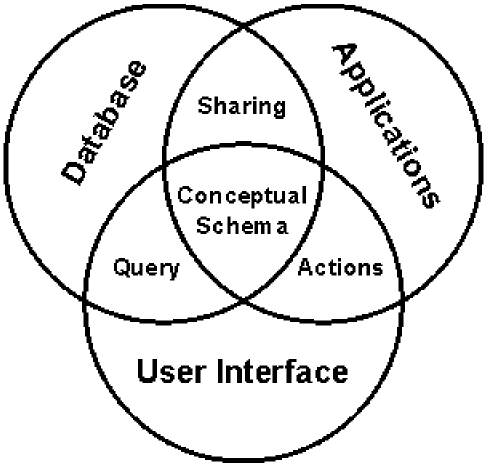

Figure 2: The ANSI/SPARC three-schema approach

The three overlapping circles in Figure 2 represent the database, the application programs, and the user interface. At the center, the *conceptual schema* defines the ontology of the concepts as the users think of them and talk about them. The *physical schema* describes the internal formats of the data stored in the database, and the *external schema* defines the view of the data presented to the application programs. The contents of a database are uninterpreted character strings, such as `"Tom Smith"` and `"85437"`. The conceptual schema specifies the *metadata* for interpreting the data as facts, such as “Employee Tom Smith has employee number 85437.” It also states constraints and business rules, such as “Every employee must have a unique employee number.”

The ANSI/SPARC report was intended as a basis for interoperable computer systems. All database vendors adopted the three-schema terminology, but they implemented it in incompatible ways. Over the next twenty years, various groups attempted to define standards for the conceptual schema and its mappings to databases and programming languages. Unfortunately, none of the vendors had a strong incentive to make their formats compatible with their competitors'. A few reports were produced, but no standards.

Meanwhile, the artificial intelligence community developed several large dictionaries and ontologies. Three of the largest were Cyc (Lenat 1995), WordNet (Miller 1995), and the Electronic Dictionary Research project (Yokoi 1995). In terms of the sheer amount of knowledge represented, Cyc is the largest with the most detailed axioms, and WordNet is the smallest and least detailed. WordNet, however, is the most widely used, largely because it is freely available over the Internet. For some purposes, the lack of detail in WordNet makes it more flexible, since it imposes fewer restrictions on how the definitions can be used.

The Cyc project, founded 1984 by Doug Lenat, illustrates the frustrations faced by AI researchers. The name comes from the stressed syllable of the word *encyclopedia* because its original goal was to encode the knowledge in the *Columbia Desk Encyclopedia*. As the project continued, the developers realized that the information in a typical encyclopedia is what people typically do not know. Much more important is the implicit knowledge everybody knows, but few people verbalize — what is often called *common sense*. The

Cyc developers, however, seriously underestimated the amount of common knowledge required to understand a typical newspaper. After 20 years of elapsed time and 700 person-years of work at a cost of 70 million dollars, the Cyc project had encoded 600,000 concept types, defined by two million axioms, and organized in 6,000 microtheories.

As people mature, they seem to learn faster by building on their previously learned background knowledge: university students learn more information more quickly than high-school students, who in turn learn faster than elementary-school students. For that reason, Lenat hoped that a large knowledge base would enable Cyc to acquire new knowledge at an ever increasing rate. Unfortunately, Project Halo, a study funded by Paul Allen, the cofounder of Microsoft, suggested that the effort required to encode new knowledge in Cyc is about the same as in other systems with much smaller knowledge bases (Friedland et al. 2004).

For Project Halo, three groups were asked to represent the knowledge in a chemistry textbook: Cycorp, Ontoprise, and SRI International. The researchers in each group translated the selected pages from English, mathematics, and chemical formulas to the formats of their system. After the three groups had tested and debugged the new knowledge base, they were given questions from a freshman-level chemistry exam, which were based on the information in those pages. Each system was required to answer the questions (as translated to its own formats) and generate English-like explanations of the answers. Following are the results:

- The scores ranged from 40% to 47% correct.
- The average cost to encode the information was about $10,000 per page from the textbook.
- Despite its large knowledge base, Cyc had the lowest score, and it did not have any advantage over the other systems in the cost of encoding knowledge.

Although the three systems were able to generate English-like explanations, none of them was able to read and understand the original English from the textbook or the English questions on the exam.

Cyc represents the culmination of the Aristotelian approach. Its hierarchy of 600,000 concept types is the largest extension of the Tree of Porphyry ever implemented, and its automated reasoning is the fulfillment of Leibniz's dream. Although Cyc's categories are not based on Kant's, they are the closest realization of the “supplementary work” that Kant envisioned. Unfortunately, “ranking things in genera and species,” as Leibniz said, has not proved to be profitable. The Cyc project has survived for over twenty years, but only with the infusion of large amounts of research funds. A few commercial applications of the Cyc technology have been modestly successful, but they have not generated enough revenue to support ongoing research and development, let alone provide any return on the original investment.

3. Knowledge Soup

The major obstacle that Cyc or any similar project must address is the complexity of the knowledge soup — the heterogeneous, often inconsistent mixture that people have in their heads. Some of it may be represented in symbolic or propositional form, but much, if not most of it is stored in image-like forms. The soup may contain many small chunks, corresponding to the typical rules and facts in Cyc, and it may also contain large chunks that correspond to Cyc's microtheories. As in Cyc, the chunks should be internally consistent, but they may be inconsistent with one another.

The complexity does not arise from the way the human brain works or the way that natural languages express information. As Whitehead observed, it results from left-over questions lurking in "the penumbral background" and the difficulty of recognizing which ones are relevant to the "focus of experience":

- *Overgeneralizations.* Birds fly. But what about penguins? A day-old chick? A bird with a broken wing? A stuffed bird? A sleeping bird? A bird in a cage?
- *Abnormal conditions.* If you have a car, you can drive from New York to Boston. But what if the battery is dead? Your license has expired? There is a major snowstorm?
- *Incomplete definitions.* An oil well is a hole drilled in the ground that produces oil. But what about a dry hole? A hole that has been capped? A hole that used to produce oil? Are three holes linked to a single pipe one oil well or three?
- *Conflicting defaults.* Quakers are pacifists, and Republicans are not. But what about Richard Nixon, who was both a Quaker and a Republican? Was he or was he not a pacifist?
- *Unanticipated applications.* The parts of the human body are described in anatomy books. But is hair a part of the body? Hair implants? A wig? A wig made from a person's own hair? A hair in a braid that has broken off from its root? Fingernails? Plastic fingernail extender? A skin graft? Artificial skin used for emergency patches? A band-aid? A bone implant? An artificial implant in a bone? A heart transplant? An artificial heart? An artificial leg? Teeth? Fillings in the teeth? A porcelain crown? False teeth? Braces? A corneal transplant? Contact lenses? Eyeglasses? A tattoo? Make-up? Clothes?

These exceptions and borderline cases result from the nature of the world, not from any defect in natural language. A logical predicate like bodyPart(x) would solve nothing, since hair and fingernails would raise the same questions as the English phrase *body part*. The discrepancy results from a mismatch of the continuous world with the discrete words of language or predicates of logic: all languages consist of discrete symbols organized in discrete syntactic patterns; the real world, however, contains an endless variety of things, events, forms, substances, gradations, changes, and continuous flows with imperceptible transitions from one to another. No language based on discrete words or symbols can ever capture the full complexity of a continuous system.

Even for knowledge that can be represented in discrete symbols, the enormous amount of detail makes it practically impossible to keep two independently designed databases or knowledge bases consistent. Most banks, for example, offer similar services, such as checking accounts, savings accounts, loans, and mortgages. Banks have always been able to interoperate in transferring funds from accounts in one to accounts in another. Yet when two banks merge, they never merge their accounts. Instead, they adopt one of two strategies:

1. Keep running both databases indefinitely, or
2. Close some or all accounts of one bank, and open new accounts in the database of the other bank.

There are too many poorly understood and incompletely documented details for any bank to risk catastrophic failure by merging independently developed databases.

The continuous gradations and open-ended range of exceptions make it impossible to give complete, precise definitions for any concepts that are learned through experience. Kant (1800) observed that artificial concepts invented by some person for some specific purpose are the only ones that can be defined completely:

> Since the synthesis of empirical concepts is not arbitrary but based on experience, and as such can never be complete (for in experience ever new characteristics of the concept can be discovered), empirical concepts cannot be defined.
>
> Thus only arbitrarily made concepts can be defined synthetically. Such definitions... could also be called *declarations*, since in them one declares one's thoughts or renders account of what one understands by a word. This is the case with *mathematicians*.

Kant's observation has been repeated with variations by philosophers from Heraclitus to the present. Two of the more recent statements are the principles of *family resemblance* by Ludwig Wittgenstein (1953) and *open texture* by Friedrich Waismann (1952):

- *Family resemblance*. Empirical concepts cannot be defined by a fixed set of necessary and sufficient conditions. Instead, they can only be defined by giving a series of examples and saying “These things and everything that resembles them are instances of the concept.”
- *Open texture*. For any proposed definition of empirical concepts, new instances will arise that “obviously” belong to the category but are excluded by the definition.

These principles imply that all classifications are approximations. For any collection of concepts, new examples will inevitably arise that don't quite fit any of the existing categories. But deductive reasoning requires precise definitions, clearly stated axioms, and formal rules of inference.

Kant, Wittgenstein, and Waismann were philosophers who fully understood the power and limitations of logic. The mathematician and philosopher Alfred North Whitehead, who was the senior author of the *Principia Mathematica*, one of the most comprehensive treatises on logic ever written, was even more explicit about its limitations. Following are some quotations from his last book, *Modes of Thought*:

- “Both in science and in logic, you have only to develop your argument sufficiently, and sooner or later you are bound to arrive at a contradiction, either internally within the argument, or externally in its reference to fact.”
- “The topic of every science is an abstraction from the full concrete happenings of nature. But every abstraction neglects the influx of the factors omitted into the factors retained.”
- “The premises are conceived in the simplicity of their individual isolation. But there can be no logical test for the possibility that deductive procedure, leading to the elaboration of compositions, may introduce into relevance considerations from which the primitive notions of the topic have been abstracted.”

Whitehead certainly believed that logic is important, but he also realized that it is only part of any comprehensive system of learning, reasoning, and acting in and upon the world. He summarized his position in one sentence: “We must be systematic, but we should keep our systems open.” Logic is an excellent means for reasoning about well-defined knowledge, but by itself, logic cannot make poorly defined terms precise or determine if any relevant information is missing.

The complexities of the knowledge soup are just as troublesome whether knowledge is represented in logic or in natural languages. When the Académie Française attempted to legislate the vocabulary and definitions of the French language, their efforts were undermined by uncontrollable developments: rapid growth of slang that is never sanctioned by the authorities, and wholesale borrowing of words from English, the world's fastest growing language. In Japan, the pace of innovation and borrowing has been so rapid that the older generation of Japanese can no longer read their daily newspapers.

4. Cognitive Processing

Over the past two million years, evolutionary processes added the human ability to think in discrete words and syntactic patterns to an ape-like ability to integrate continuous geometrical information from the visual, tactile, auditory, and motor regions of the brain. The ape brain itself took about a hundred times longer to evolve from a fish-like stage, which in turn took several times longer to evolve from a worm-like ganglion. Figure 3 illustrates the evolutionary stages.

The cognitive systems of the animals at each level of Figure 3 build on and extend the capabilities of the earlier levels. The worms at the top have rudimentary sensory and motor mechanisms connected by ganglia with a small number of neurons. A neural net that connects stimulus to response with just a few intermediate layers might be an adequate model. The fish brain is tiny compared to the mammals, but it already has a complex structure that receives inputs from highly differentiated sensory mechanisms and sends outputs to just as differentiated muscular mechanisms, which support both delicate control and high-speed propulsion. Exactly how those mechanisms work is not known, but the neural evidence suggests a division into perceptual mechanisms for interpreting inputs and motor mechanisms for controlling action. Between perception and action there must also be some sort of cognitive processing that combines and relates new information from the senses with memories of earlier perceptions.

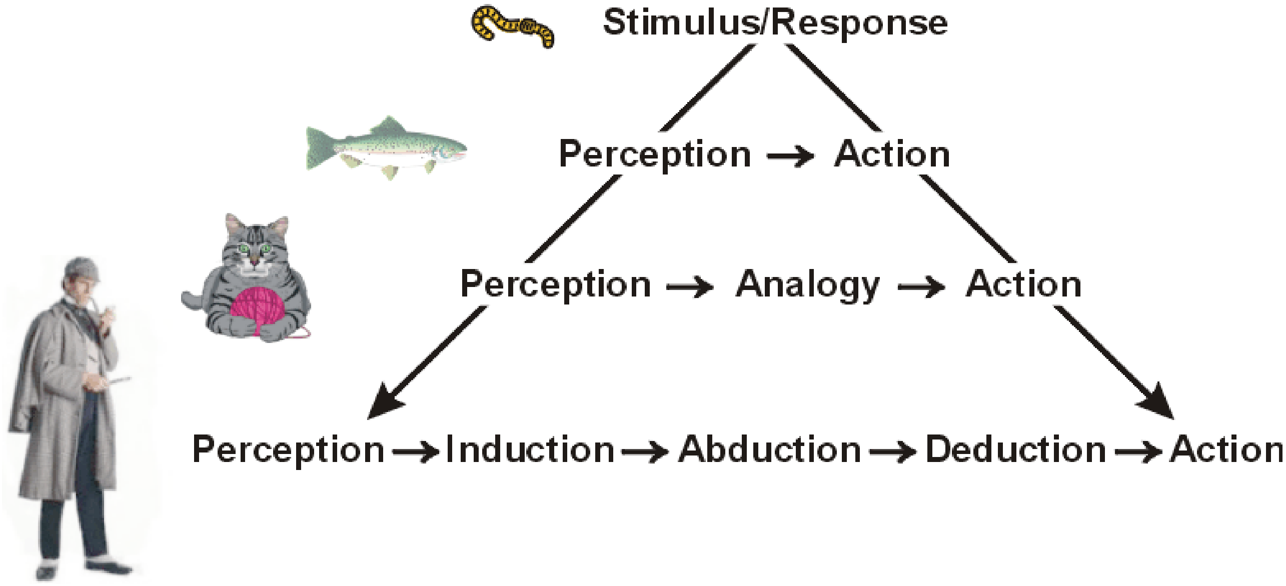

Figure 3: Evolution of Cognition

At the next level, mammals have a cerebral cortex with distinct *projection areas* for each of the sensory and motor systems. If the fish brain is already capable of sophisticated perception and motor control, the larger cortex must add something more. Figure 3 labels it *analogy* and symbolizes it by a cat playing with a ball of yarn that serves as a mouse analog. The human level is illustrated by a typical human, Sherlock Holmes, who is famous for his skills at induction, abduction, and deduction. Those reasoning skills, which Peirce analyzed in detail, may be characterized as specialized ways of using analogies, but they work seamlessly with the more primitive abilities.

Without language and logic, human society and civilization would still be at the level of the apes. Those animals are extremely intelligent, but they lack the ability to recognize and generate symbols. The neurophysiologist and anthropologist Terence Deacon (1997) argued that the slow evolution of the brain and vocal tract toward modern forms indicates that early hominids already had a rudimentary language, which gave individuals with larger brains and better speaking ability a competitive advantage. To distinguish human and animal communication, Deacon used Peirce's semiotic categories of *icon, index,* and *symbol* to classify the kinds of signs they could recognize and produce. He found that higher mammals easily recognize the first two kinds, icons and indexes, but only after lengthy training could a few talented chimpanzees learn to recognize symbolic expressions. Deacon concluded that if chimpanzees could make the semiotic transition from indexes to symbols, early hominids could. The evidence suggests that the transition to symbolic communication occurred about two million years ago with *Homo Habilis*. Once that transition had been made, language was possible, and the use of language promoted the co-evolution of both language and brain.

Logic is much broader than the mathematical version invented by Boole, Peirce, and Frege. It includes Aristotle's syllogisms, which are a stylized form of the reasoning expressed in ordinary language, but the only reasoning method supported by syllogisms is deduction, which is also the primary method used in Cyc and the Semantic Web. Peirce recognized that deduction is important, but he maintained that two other methods — induction and abduction — are equally important:

1. *Deduction.* Apply a general principle to infer some fact.

   ```
   Given:  Every bird flies.  Tweety is a bird.
   Infer:  Tweety flies.
   ```

2. *Induction.* Assume a general principle that subsumes many facts.

   ```
   Given:  Tweety, Polly, and Hooty are birds.  Fred is bat.
           Tweety, Polly, and Hooty fly.  Fred flies.
   Assume: Every bird flies.
   ```

3. *Abduction.* Guess a new hypothesis that explains some fact.

   ```
   Given:  Every bird flies.  Tweety flies.
   Guess:  Tweety is a bird.
   ```

These three methods of reasoning depend on the ability to use symbols. In deduction, the general term *every bird* is replaced by the name of a specific bird *Tweety*. Induction generalizes a property of multiple individuals — Tweety, Polly, and Hooty — to the category Bird, which subsumes all the instances. Abduction guesses the new proposition *Tweety is a bird* to explain one or more observations. According to Deacon's hypothesis that symbols are uniquely human, these three reasoning methods could not be used by nonhuman mammals, not even the apes.

According to Peirce (1902), “Besides these three types of reasoning there is a fourth, analogy, which combines the characters of the three, yet cannot be adequately represented as composite.” Analogy is more primitive than logic because it does not require language or symbols. Its only prerequisite is *stimulus generalization* — the ability to recognize similar patterns of stimuli as signs of similar objects or events. In Peirce's terms, logical reasoning requires symbols, but analogical reasoning could also be performed on image-like signs called *icons*.

Analogical reasoning is general enough to derive the kinds of conclusions typical of a logic-based system that uses induction to derive rules followed by deduction to apply the rules. In AI systems, that method is called *case-based reasoning* (Riesbeck & Schank 1989), but the principle was first stated by Ibn Taymiyya in his comparison of Aristotle's logic to the analogies used in legal reasoning (Hallaq 1993).

Ibn Taymiyya admitted that deduction in mathematics is certain. But in any empirical subject, general axioms can only be derived by induction, and induction must be guided by the same principles of evidence and relevance used in analogy. Figure 4 illustrates his argument: Deduction proceeds from a *theory* containing general axioms. But those axioms must have earlier been derived by induction with the same criteria used for analogy. The only difference is that induction produces a theory as intermediate result, which is then used in a subsequent process of deduction. By using analogy directly, legal reasoning dispenses with the intermediate theory and goes straight from cases to conclusion. If the theory and the analogy are based on the same evidence, they must lead to the same conclusions.

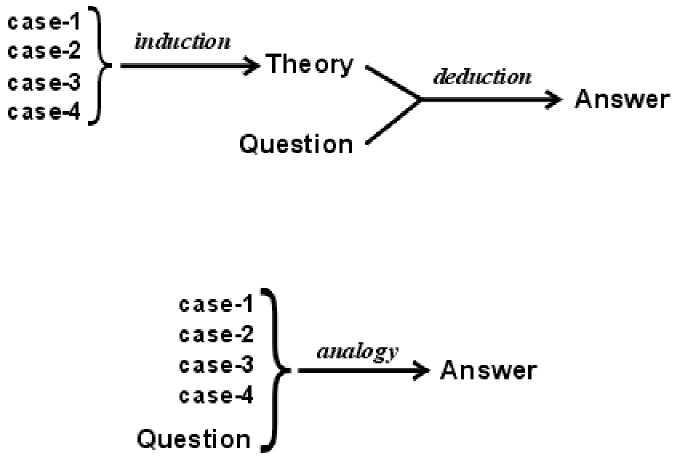

Figure 4: Comparison of logical and analogical reasoning

Note that the theory in Figure 4 requires some language or system of symbols for stating the axioms, but case-based reasoning (CBR) can be applied directly to image-like icons without any requirement for symbols as intermediaries. In legal reasoning, all the cases are described in symbols, but animals without language could apply analogical reasoning to cases recorded in imagery. Temple Grandin, an autistic woman who despite the odds managed to earn a PhD, gave a vivid description of her use of imagery (Grandin & Johnson 2004). She believes her image-based reasoning methods are of the same nature as the reasoning methods used by other mammals. She noted that her own reasoning tended to be very concrete, and she found it impossible to understand abstractions that could not be related to concrete images. She was able to do mathematics, but she thought of numbers as images, which she could visualize and manipulate.

Whether the medium consists of words or images, the methods of CBR are the same: start with a question or goal Q about some current problem or situation P. By the methods of analogy, previously experienced cases that resemble P are recalled from long-term memory. When the cases in Figure 4 are recalled, they must be ranked according to their similarity or *semantic distance* to the current situation P. The case with greatest similarity to P (i.e., the smallest semantic distance) is considered the most relevant and the most likely to provide a suitable answer to the question Q. When a relevant case has been found, the aspect of the case that provides the information requested by Q is the predicted answer. If two or more cases have nearly equal relevance, they may or may not predict the same answer. If they do, that answer can be accepted with a high degree of confidence. If not, the answer is a disjunction of alternatives: Q_1 or Q_2. Psychologically, none of these operations depend on consciousness; in fact, the methods of measuring similarity or relevance are necessarily preconscious because they determine which images from long-term memory are introduced to consciousness.

As an application of CBR, the VivoMind Analogy Engine (Sowa & Majumdar 2003) was used to evaluate free-form answers to algebra word problems, which ranged in length from a short phrase to one or two sentences. Such texts are difficult to analyze correctly by computer, and even good teachers are not 100% correct in their evaluations. Two previous attempts had failed to provide a satisfactory solution:

- *Logic based.* One approach used a deductive method as described in Section 2. However, it required a sophisticated ontology and methods of knowledge representation that were unfamiliar to most high-school teachers.
- *Statistical.* Another was based on statistical methods commonly used for information retrieval. However, it could not distinguish correct answers from incorrect answers because they often used the same words, but in different word order.

The VivoMind method was short and direct: translate the English sentences or phrases to conceptual graphs (CGs); use VAE to compare them to previously evaluated answers; and report the teacher's evaluation for the answer that gave the closest match.

In terms of Figure 4, the cases were the student answers that had previously been evaluated by some teacher. For each evaluated answer, there was a teacher's response of the form "correct," "incorrect," or "partially correct with the following information missing." For any new case P, the question Q was a request for an evaluation. Unlike the logic-based or statistical approaches, the CBR method had an excellent success rate on cases for which VAE found a match with a small semantic distance. If VAE could not find a close match for some case or if two or more answers were equally close, VAE would send the case to a teacher, who would write a new evaluation. Then the new case with its associated evaluation would enlarge the set for future evaluations. After 30 or 40 cases had accumulated, a teacher's assistance was rarely needed to evaluate new cases. This example illustrates several points about analogical reasoning in general and case-based reasoning in particular:

1. Every case is stored in memory as a pattern of signs, but those signs need not be symbols. External stimuli are usually iconic or image-like; internal signs or feelings may be differentiated by their origin and strength, but they are indexical, not symbolic.
2. Some measure of semantic distance is necessary for determining which previous case P is the most relevant to the current goal or question Q.
3. Measuring semantic distance requires some basis for determining the relevance or preference for one pattern of signs over another, but the signs need not be symbolic. Semantic distance may be determined by similarities of icons or indexical references.
4. A hierarchical classification, similar to the Tree of Porphyry, would be useful for determining semantic distance, but the elements of the hierarchy need not be concepts, words, or other symbols. A hierarchy of images could be determined by the perceptual mechanisms of stimulus generalization.
5. For any measure of semantic distance, the structure or geometry of the patterns is at least as important as the elements that occur in the patterns. For symbolic

patterns, such as linguistic sentences and stories, the relational structure and the types of concepts are equally relevant.

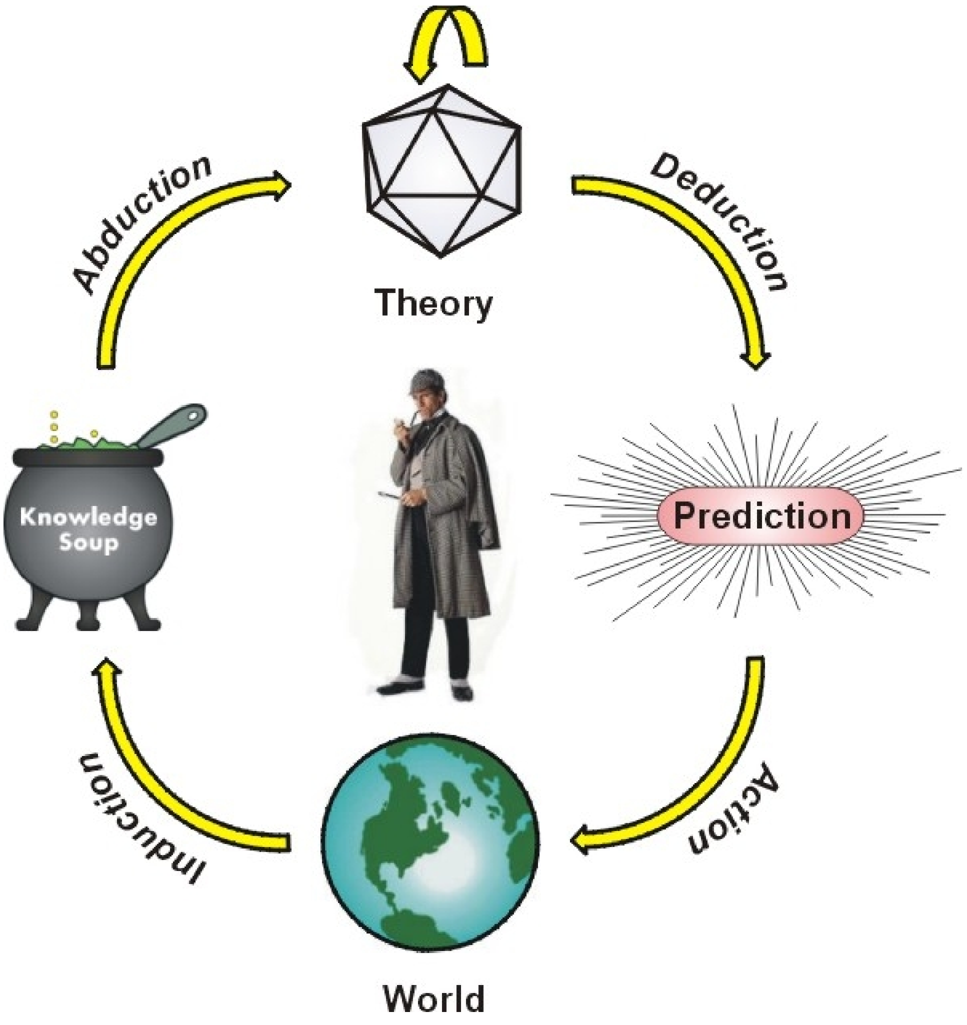

Figure 5: The Cycle of Pragmatism

A measure of semantic distance is essential for analogy, but it can also improve the logical methods of induction, deduction, and abduction. To illustrate the relationships, Figure 5 shows an agent who repeatedly carries out the stages of induction, abduction, deduction, and action. The arrow of induction indicates the accumulation of previously useful patterns in the knowledge soup. The crystal at the top symbolizes the elegant, but fragile theories that are constructed from chunks extracted from the soup by abduction. The arrow above the crystal indicates the process of belief revision, which uses repeated abductions to modify the theories. At the right is a prediction derived from a theory by deduction. That prediction leads to actions whose observable effects may confirm or refute the theory. Those observations are the basis for new inductions, and the cycle continues.

Of the three methods of logic, abduction is the only one that can introduce a truly novel idea. In Peirce's system, abduction is the replacement for Descartes's innate ideas and for Kant's synthetic *a priori* judgments. Abduction is the process of using some measure of semantic distance to select relevant chunks from the knowledge soup and assemble them in a novel combination. It can be performed at various levels of complexity:

- *Reuse*. Do an associative search for a previously used rule, pattern, or theory that can be applied to the current problem.
- *Revise*. Find a theory or fragment of a theory that approximately matches the problem at hand and apply the belief revision operators to adapt it to the current situation.
- *Combine*. Search for scattered fragments or *chunks* of knowledge and perform repeated steps of belief revision to combine them or modify them to form a theory appropriate to the current situation. If two theories are consistent with one another, their combination can be defined by a conjunction of the axioms of both. If their conjunction is inconsistent, several different consistent combinations may be extracted by choosing different axioms to delete in order to avoid contradictions.

All these processes may be used iteratively. After a hypothesis is formed by abduction, its implications must be tested against reality. If its implications are not confirmed, the hypothesis must be revised in another stage of abduction. In Peirce's logic of pragmatism, the free creations of thought generated by abduction are constrained at the two “gates” of perception and action:

> The elements of every concept enter into logical thought at the gate of perception and make their exit at the gate of purposive action; and whatever cannot show its passports at both those two gates is to be arrested as unauthorized by reason. (EP 2.241).

Note Peirce's word *elements*: abduction does not create totally new elements, but it can reassemble previously observed elements in novel combinations. Each combination defines a new concept, whose full meaning is determined by the totality of purposive actions it implies. As Peirce said, meanings grow as new information is received, new implications are derived, and new actions become possible.

In summary, deduction can be very useful when a theory is available, as in mathematics, science, and engineering. But analogy can be used when no theory exists, as in law, medicine, business, and everyday life. Even when logic is used, the methods of induction and abduction on the left side of Figure 5 are necessary for learning new knowledge and organizing it into the systematic theories required for deduction. Figure 5 suggests why Cyc or any other purely deductive system will always be limited: it addresses only the upper left part of the cycle. Today, computerized theorem provers are better at deduction than most human beings, but deduction is only 25% of the cycle. Instead of automating Sherlock Holmes, Cyc and other deductive systems require people with his level of expertise to write axioms at a cost of $10,000 to encode one page from a textbook. The methods of induction, abduction, and analogy are key to designing more robust systems, and good measures of semantic distance are key to analogy as well as all three methods of logic.

5. Directions for Future Research

Language understanding requires pattern recognition at every level from phonology and syntax to semantics and pragmatics. The sound patterns learned in infancy enable an adult to understand his or her native language in a noisy room while selectively attending to one of several simultaneous conversations. Although most syntactic patterns can be programmed as grammar rules, the enormous flexibility and novel collocations of ordinary language depend on semantic patterns, background knowledge, extralinguistic context, and even the speaker's and listener's familiarity with each other's interests, preferences, and habits. Even when the syntax and semantics of a sentence is correctly parsed, the listener must recognize the speaker's intentions and expectations and their implications in terms of the current situation. Any and every aspect of human knowledge and experience may be needed to understand any given sentence, and it must be accessible from long-term memory in just a few milliseconds. Furthermore, much if not most of the knowledge may be stored in nonlinguistic patterns of images derived from any sensory modality.

In a few short years, children learn to associate linguistic patterns with background knowledge in ways that no computer can match. The following sentence was spoken by Laura Limber at age 34 months:

> *When I was a little girl, I could go "geek, geek" like that;*
> *but now I can go "This is a chair."*

John Limber (1973) recorded Laura's utterances throughout her early years and analyzed her progress from simple to complex sentences. In this short passage, she combined subordinate and coordinate clauses, past tense contrasted with present, the modal auxiliaries *can* and *could*, the quotations "geek, geek" and "This is a chair," metalanguage about her own linguistic abilities, and parallel stylistic structure. Admittedly, Laura was a precocious child who probably benefited from the extra attention of her father's psychological studies, but children in all cultures master the most complex grammars with success rates that put computer learning systems to shame. The challenge of simulating that ability led the computer scientist Alan Perlis to remark "A year spent in artificial intelligence is enough to make one believe in God" (1982).

The despair expressed by Perlis afflicted many AI researchers. Terry Winograd, for example, called his first book *Understanding Natural Language* (1972) and his second book *Language as a Cognitive Process: Volume I, Syntax* (1983). But he abandoned the projected second volume on semantics when he realized that no existing semantic theory could explain how anyone, human or computer, could understand language. With his third book, coauthored with the philosopher Fernando Flores, Winograd (1986) switched to his later work on the design of human-computer interfaces. Winograd's shift in priorities is typical of much of the AI research over the past twenty years. Instead of language understanding, many people have turned to the simpler problems of text mining, information retrieval, and designing user interfaces.

One of the projects that is still pursuing a purely deductive approach is the Semantic Web (Berners-Lee et al. 2001). Its goals are similar to the goals of the ANSI/SPARC approach illustrated in Figure 2: support interoperability among independently developed computer applications. But instead of interoperability among multiple applications that use the same database, the Semantic Web is addressing the bigger problem of supporting interoperability among applications scattered anywhere across the World Wide Web. Unfortunately, current proposals for the Semantic Web are based on a subset of the Cyc logic with an ontology that is still a fraction of the size of Cyc. The only hope for any project with the scope of the Semantic Web is to address the full pragmatic cycle in Figure 5, with emphasis on the methods of induction and abduction on the left side of the cycle. With anything less, it's unlikely that thousands of ontologists scattered across the WWW using a weaker version of logic will be more successful than Cyc's tightly managed group located at a single site.

The central problem for any kind of reasoning is finding relevant information when needed. As Leibniz observed, a well-organized hierarchy "helps one not merely to retain things, but also to find them." But as this article has shown, the complexities of the world cause such systems to break down into knowledge soup. Humans and other animals have been successful in dealing with that soup because their brains have high-speed associative memories for finding relevant information as needed. To support such a memory in computer systems, Majumdar and Sowa (forthcoming) developed a method of *knowledge signatures*, which uses continuous numeric algorithms for encoding knowledge and finding relevant knowledge with the speed of an associative memory. This method enabled the VivoMind Analogy Engine to find analogies in time proportional to (N log N) instead of the older N-cubed algorithms. Such a method for finding relevant information can make a dramatic improvement in every method of reasoning, including induction, deduction, abduction, and analogy. More research is needed to develop tools and techniques to take advantage of these algorithms and incorporate them in practical systems that can manage the knowledge soup and use it effectively. Such systems should be available within the next few years, and they promise to revolutionize the way knowledge systems are designed and used.

As a closing thought, the following quotation by the poet Robert Frost develops the theme of the opening quotation by Whitehead:

> *I've often said that every poem solves something for me in life. I go so far as to say that every poem is a momentary stay against the confusion of the world.... We rise out of disorder into order. And the poems I make are little bits of order.*
>
> Robert Frost, *A Lover's Quarrel with the World*

In fact, the word *poem* comes from the Greek *poiein*, which means to make or create. This quotation could be applied to the creations in any field of human endeavor just by replacing every occurrence of *poem* with the name of the characteristic product: *theory* for the scientist, *design* for the engineer, *building* for the architect, or *menu* for the chef. Each person creates new bits of order by reassembling and reorganizing the chunks in the knowledge soup. Frost's metaphor of a lover's quarrel is an apt characterization of the human condition and our mental soup of memories, thoughts, insights, and the inevitable misunderstandings. Any attempt to simulate human intelligence or to give people more flexible, more human-like tools must meet the challenge.

References

Berners-Lee, Tim, James Hendler, & Ora Lassila (2001) "The Semantic Web," *Scientific American*, May 2001.

Friedland, Noah S., Paul G. Allen, Gavin Matthews, Michael Witbrock, David Baxter, Jon Curtis, Blake Shepard, Pierluigi Miraglia, Jürgen Angele, Steffen Staab, Eddie Moench, Henrik Oppermann, Dirk Wenke, David Israel, Vinay Chaudhri, Bruce Porter, Ken Barker, James Fan, Shaw Yi Chaw, Peter Yeh, Dan Tecuci, & Peter Clark (2004) "Project Halo: towards a digital Aristotle," *AI Magazine* **25:4**, 29-48.

Frost, Robert (1963) *A Lover's Quarrel with the World* (filmed interview), WGBH Educational Foundation, Boston.

Grandin, Temple, & Catherine Johnson (2004) *Animals in Translation: Using the Mysteries of Autism to Decode Animal Behavior*, Scribner, New York.

Hallaq, Wael B. (1993) *Ibn Taymiyya Against the Greek Logicians*, Clarendon Press, Oxford.

Kant, Immanuel (1787) *Kritik der reinen Vernunft*, translated by N. Kemp Smith as *Critique of Pure Reason*, St. Martin's Press, New York.

Kant, Immanuel (1800) *Logik: Ein Handbuch zu Vorlesungen*, translated as *Logic* by R. S. Hartmann & W. Schwarz, Dover Publications, New York.

Knowlson, James (1975) *Universal Language Schemes in England and France, 1600-1800*, University of Toronto Press, Toronto.

Leibniz, Gottfried Wilhelm (1666) "Dissertatio de arte combinatoria," *Leibnizens mathematische Schriften* **5**, Georg Olms, Hildesheim.

Leibniz, Gottfried Wilhelm (1705) *Nouveaux Essais concernant l'Entendement humain*, translated as *New Essays on Human Understanding*, P. Remnant & J. Bennet, Cambridge University Press, Cambridge, 1981.

Lenat, Douglas B. (1995) "Cyc: A large-scale investment in knowledge infrastructure," *Communications of the ACM* **38:11**, 33-38.

Limber, John (1973) "The genesis of complex sentences," in T. Moore, ed., *Cognitive Development and the Acquisition of Language*, Academic Press, New York, 169-186. http://pubpages.unh.edu/~jel/JLimber/Genesis_complex_sentences.pdf

Majumdar, Arun K., & John F. Sowa (forthcoming) "Cognitive memory based on knowledge signatures."

Miller, George A. (1995) "WordNet: A lexical database for English," *Communications of the ACM* **38:11**, 39-41.

Peirce, Charles S. (1902) *Logic, Considered as Semeiotic*, MS L75, edited by Joseph Ransdell, http://members.door.net/arisbe/menu/library/bycsp/l75/l75.htm

Peirce, Charles Sanders (EP) *The Essential Peirce*, ed. by N. Houser, C. Kloesel, and members of the Peirce Edition Project, 2 vols., Indiana University Press, Bloomington, 1991-1998.

Perlis Alan J. (1982) "Epigrams in Programming," *SIGPLAN Notices*, Sept. 1982, ACM. http://www.cs.yale.edu/homes/perlis-alan/quotes.html

Peter of Spain or Petrus Hispanus (circa 1239) *Summulae Logicales*, edited by I. M. Bocheński, Marietti, Turin, 1947.

Riesbeck, Christopher K. & Roger C. Schank, (1989) *Inside Case-Based Reasoning,* Erlbaum, Hillsdale, NJ.

Sowa, John F. (2000) *Knowledge Representation: Logical, Philosophical, and Computational Foundations*, Brooks/Cole Publishing Co., Pacific Grove, CA.

Sowa, John F., & Arun K. Majumdar (2003) "Analogical reasoning," in A. de Moor, W. Lex, & B. Ganter, eds., *Conceptual Structures for Knowledge Creation and Communication*, Proceedings of ICCS 2003, LNAI 2746, Springer-Verlag, Berlin, pp. 16-36.

Tsichritzis, Dionysios C., & Anthony Klug, eds. (1978) "The ANSI/X3/SPARC DBMS framework," *Information Systems* **3**, 173-191.

Waismann, F. (1952) "Verifiability," in A. Flew, ed., Logic and Language, first series, Basil Blackwell, Oxford.

Wittgenstein, Ludwig (1953) Philosophical Investigations, Basil Blackwell, Oxford.

Whitehead, Alfred North (1937) *Essays in Science and Philosophy*, Philosophical Library, New York.

Whitehead, Alfred North (1938) *Modes of Thought*, Macmillan Co., New York.

Winograd, Terry (1972) *Understanding Natural Language*, Academic Press, New York.

Winograd, Terry (1983) *Language as a Cognitive Process: Volume 1, Syntax*, Addison-Wesley, Reading, MA.

Winograd, Terry, & Fernando Flores (1986) *Understanding Computers and Cognition: A New Foundation for Design*, Ablex, Norwood, NJ.

Wittgenstein, Ludwig (1953) *Philosophical Investigations*, Basil Blackwell, Oxford.

Yokoi, Toshio (1995) "The EDR electronic dictionary," *Communications of the ACM* **38:11**, 42-44.