

The Role of Logic and Ontology In Language and Reasoning

John F. Sowa

Abstract. Natural languages have words for all the operators of first-order logic, modal logic, and many logics that have yet to be invented. They also have words and phrases for everything that anyone has ever discovered, assumed, or imagined. Aristotle invented formal logic as a tool (*organon*) for analyzing and reasoning about the ontologies implicit in language. Yet some linguists and logicians took a major leap beyond Aristotle: they claimed that there exists a special kind of logic at the foundation of all NLs, and the discovery of that logic would be the key to harnessing their power and implementing them in computer systems. Projects in artificial intelligence developed large systems based on complex versions of logic, yet those systems are fragile and limited in comparison to the robust and immensely expressive natural languages. Formal logics are too inflexible to be the foundation for language; instead, logic and ontology are abstractions from language. This reversal turns many theories about language upside down, and it has profound implications for the design of automated systems for reasoning and language understanding. This article analyzes these issues in terms of Peirce's semiotics and Wittgenstein's language games. The resulting analysis leads to a more dynamic, flexible, and extensible basis for ontology and its use in formal and informal reasoning.

This article is a slightly revised preprint of Chapter 11 in *Theory and Applications of Ontology: Philosophical Perspectives*, edited by R. Poli & J. Seibt, Berlin: Springer, pp. 231-263.

1. The Search for Foundations

Natural languages are the most sophisticated systems of communication ever developed. Formal ontologies are valuable for applications in science, engineering, and business, but they have been difficult to generalize beyond narrowly defined *microtheories* for specialized domains. For language understanding, formal systems have only been successful in narrow domains, such as weather reports and airline reservations. Frege, Russell, and the Vienna Circle tried to make formal logic the universal language of science, but that attempt failed. Only the final results of any research can be stated formally, never the vague hunches, intuitive explorations, and heated debates that are necessary for any creative advance. Scientists and engineers criticized formal methods with the pithy slogans "Physicists don't do axioms" and "All models are wrong, but some are useful." Even Aristotle, who invented the first formal logic, admitted that his syllogisms and categories are important for stating the results of research, but that informal methods are necessary for gathering and interpreting empirical evidence.

Aristotle's logic and categories still serve as a paradigm for the ontologies used in modern computer systems, but his grand synthesis began to break down in the 16th century. Aristotle's physics and cosmology were demolished by the work of Copernicus, Galileo, Kepler, and Newton. In philosophy, the skeptical tradition of antiquity was revived by the publication in 1562 of a new edition of the works of Sextus Empiricus, whose attacks on Aristotle were popularized by the essays of Michel de Montaigne. In responding to the skeptics, Descartes began his search for certainty from the standpoint of universal doubt, but he merely reinforced the corrosive effects of skepticism. The British empiricists responded with new approaches to epistemology, which culminated in Hume's devastating criticisms of the foundations of science itself. Two responses to Hume helped to restore the legitimacy of science: Thomas Reid's critical common sense and Immanuel Kant's three major *Critiques*. Kant (1787)

adopted Aristotle's logic as the basis for his new system of categories, which he claimed would be sufficient for defining all other concepts:

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. (A:82, B:108)

Two centuries later, Kant's "easy" task is still unfinished. His *Opus postumum* records the struggles of the last decade of his life when Kant tried to make a transition from his *a priori* metaphysics to the experimental evidence of physics. Förster (2000) wrote "although Kant began this manuscript in order to solve a comparatively minor problem within his philosophy, his reflections soon forced him to readdress virtually all the key problems of his critical philosophy: the objective validity of the categories, the dynamical theory of matter, the nature of space and time, the refutation of idealism, the theory of the self and its agency, the question of living organisms, the doctrine of practical postulates and the idea of God, the unity of theoretical and practical reason, and, finally, the idea of transcendental philosophy itself."

Unlike Aristotle, who used logic as a tool for analyzing language, Kant assumed that logic is a prerequisite, not only for language, but for all rational thought. Richard Montague (1970) pushed Kant's assumption to an extreme: "I reject the contention that an important theoretical difference exists between formal and natural languages." That assumption, acknowledged or not, motivated much of the research in artificial intelligence and formal linguistics. The resulting systems are theoretically impressive, but they cannot learn and use ordinary language with the ease and flexibility of a three-year-old child. But if logic is inadequate, what other foundation could support linguistics and AI? What kind of semantics could represent the highly technical formalisms of science, the colloquial speech of everyday life, and the requirements for sharing and reasoning with the knowledge scattered among millions of computers across the Internet? How would the research and development change under the assumption that logic is a derivative from language, not a prerequisite for it?

One major change would be a shift in emphasis from the rigid views of Frege, Russell, and Carnap to the more flexible philosophies of Peirce, Whitehead, and the later Wittgenstein. As logicians, those two groups were equally competent, but the former considered logic to be superior to language, while the latter recognized the limitations of logic and the power of language. Peirce, in particular, was a pioneer in logic, but he included logic within the broader field of semiotics, or as he spelled it, *semeiotic*. That broader view relates the precise formalisms of grammar and logic to the more primitive, yet more flexible mechanisms of perception, action, and learning. Language is based on vocal signs and patterns of signs, whose more stable forms are classified as vocabulary and grammar. But instead of starting with formal precision, the first signs are vague, ambiguous, and uncertain. Precision is a rare state that never occurs in the early stages of learning, and absolute precision is unattainable in any semiotic system that represents the real world. Grammar, logic, and ontology describe stable patterns of signs or *invariants* under transformations of perspective. Those stable patterns, which are fossilized in formal theories, develop as each individual interacts with the world and other creatures in it.

Although formal logic can be studied independently of natural language semantics, no formal ontology that has any practical application can ever be developed and used without acknowledging its intimate connection with NL semantics. An ontology for medical informatics, for example, must be related to medical publications, to a physician's diagnoses, and to the discussions among general practitioners, specialists, nurses, patients, and the programmers who develop the software they use. All these people are constantly thinking and using NL semantics, not the formal axioms of some theory. Frege (1879) hoped "to break the domination of the word over the human spirit by laying bare the misconceptions

that through the use of language often almost unavoidably arise concerning the relations between concepts.” Wittgenstein agreed that language can be misleading, but he denied that an artificial language could be better. At best, it would be a different *language game* (Sprachspiel).

These philosophical observations explain why large knowledge bases such as Cyc (Lenat 1995) have failed to achieve true artificial intelligence. An inference engine attached to a large collection of formally defined facts and axioms can prove theorems more efficiently than most people, but it lacks the flexibility of a child in learning new information and adapting old information to new situations (Sowa 2005). Two computer scientists who devoted their careers to different aspects of AI have concluded that the goal of a fixed formal ontology of everything is both unattainable and misguided. Alan Bundy, who developed formal methods for theorem proving and problem solving, proposed *ontology evolution* as a method for systematically relating smaller domain ontologies and adapting them to specific problems (Bundy & McNeill 2006; Bundy 2007). Yorick Wilks, who developed informal methods of *preference semantics* for natural language processing, maintained that the lexical resources used in language analysis and interpretation are sharply distinct from and should not be confused with formal ontologies (Wilks 2006, 2008a,b). These two views can be reconciled by using linguistic information as the basis for indexing and relating an open-ended variety of task-oriented ontologies.

Instead of a static ontology, this article develops a dynamic approach that can relate the often vague and shifting meanings of ordinary words to the formal ontologies needed for computer applications. Section 2 surveys semiotic theories from Aristotle to Saussure with an emphasis on Peirce’s contributions. Section 3 summarizes 20th-century theories of language, the more recent recognition of their limitations, and the relevance of Peirce and Wittgenstein. Section 4 reviews a computational approach to language developed by Margaret Masterman, a former student of Wittgenstein’s. Section 5 presents a method of mapping language to logic based on Wittgenstein’s notions of *Satzsystem* and *Beweissystem*. Section 6 shows the underlying relationships between the formal methods of deduction, induction, and abduction and the informal methods of analogy. The concluding Section 7 develops a foundation for ontology that supports both formal and informal methods of reasoning. The formal theories necessary for precise reasoning are embedded in a framework that can accommodate the open-ended flexibility of natural languages.

2. A Semiotic Foundation for Ontology

Without vagueness at the foundation, the words and syntax of our stone-age ancestors could never have been adapted to every innovation of modern civilization. The evolution of language over millions of years is recapitulated in a dozen years as an infant grows to a teenager and in another dozen years to a PhD. The growth in precision of the second dozen years, although valuable, is almost trivial in comparison to the achievements of the first dozen — or even the first three. The key to understanding how language works lies in the first three years from infancy to early childhood — alternatively, in the past three million years from *Australopithecus* to *Homo sapiens*. The signs of language are rooted in the prelinguistic signs of the infant and the ape, which in turn have evolved from more primitive signs generated and recognized by every living thing from bacteria to humans. The signs of every ontology, formal or informal, are derived from those same roots.

Aristotle began the study of signs by *psyches* at every level from the vegetative to the rational. His treatise *On Interpretation*, one of the most influential books on language ever written, begins by relating language to the psyche, while avoiding *psychologism*. The following passage relates language to internal processes, whose existence is not in doubt, but whose nature is unknown. As he said, the psyche is a different subject.

First we must determine what are noun (*onoma*) and verb (*rhêma*); and after that, what are negation (*apophasis*), assertion (*kataphasis*), proposition (*apophansis*), and sentence (*logos*). Those in speech (*phonê*) are symbols (*symbola*) of affections (*pathêmata*) in the psyche, and those written (*graphomena*) are symbols of those in speech. As letters (*grammata*), so are speech sounds not the same for everyone. But they are signs (*sêmeia*) primarily of the affections in the psyche, which are the same for everyone, and so are the objects (*pragmata*) of which they are likenesses (*homoiômata*). On these matters we speak in the treatise on the psyche, for it is a different subject. (16a1)

In this short passage, Aristotle introduced ideas that have been adopted, ignored, revised, rejected, and dissected over the centuries. By using two different words for sign, Aristotle recognized two distinct ways of signifying: he adopted *sêmeion* for a natural sign and *symbolon* for a conventional sign. With the word *sêmeion*, which was used for symptoms of a disease, Aristotle implied that the verbal sign is primarily a natural sign of the mental affection or concept and secondarily a symbol of the object it refers to. That triad of sign, sense, and object constitutes the *meaning triangle*, which Ogden and Richards (1923) drew explicitly.

From the 12th to the 16th centuries, the Scholastics developed semiotics in great depth and subtlety. They observed the distinction between natural and conventional signs, but they applied the generic term *signum* to signs of all kinds. For the other two vertices of the meaning triangle, they coined the terms *significatio* for the sense of a sign and *suppositio* for its intended object. They recognized that the supposition of a sign might not exist: one example would be a mythical beast or chimera, but more often the intended object would be the result of a future action. They also recognized that the supposition could be another sign and coined the terms *prima intentio* for a sign whose supposition is a physical entity and *secunda intentio* for a sign whose supposition is another sign (Figure 1).




Figure 1. Scholastic meaning triangles

For the signification at the top of a triangle, a common Latin term was the translation *passio animae* of Aristotle's *pathêma* or a mental term such as *intellectus conceptio*. The sign at the lower left could be a word, an image, or a concept. The supposition at the lower right could be an existing thing, a spoken sign, or an imagined sign of something sought, feared, planned, or otherwise intended. The Scholastic semiotics with a Tarski-style semantics by Ockham (1323) provided a more structured and robust foundation for language and logic than the loose associations proposed a few centuries later by Locke, Condillac, and others. Franz Brentano (1874) had studied Scholastic logic, and he adopted the word *intentio* as the basis for his theory of *intentionality*, which he defined as the directedness (*Gerichtetheit*) of thought toward some object, real or imagined.

Not all linguists and logicians recognized the sign relation as triadic. For his *sémiologie*, Saussure (1916) defined a sign as a dyadic relation that excludes the object at the lower right of the meaning triangle: "The linguistic sign unites, not a thing and a name, but a concept and a sound-image. The

latter is not the material sound, a purely physical thing, but the psychological imprint of the sound, the impression that it makes on our senses.” Tarski, Quine, and many other logicians ignored the top of the triangle and focused on the dyadic link between the sign and object. The dyadic versions by Saussure and Tarski have complementary weaknesses: Tarski had a clear criterion for truth, but no recognition of intention; Saussure’s dyads permitted multiple levels of interpretation, but they did not relate words and sentences to objects and events. A complete theory of meaning must recognize the full triad.

Frege had a triadic definition of sign with the labels *Zeichen*, *Sinn*, and *Bedeutung*, usually translated *sign*, *sense*, and *reference*, which correspond to the Latin *signum*, *significatio*, and *suppositio*. In comparing Frege and Husserl, Mohanty (1982:57) observed that they had similar definitions and similar difficulties with context-dependent indexicals, such as *I*, *you*, *this*, or *that*:

Husserl refers to indexicals and their like as threatening “to plunge all our hard-won distinctions back into confusion.” These are what he calls “essentially occasional” expressions, in their case “it is essential to orient actual meaning to the occasion, the speaker, and the situation.”

Those “hard-won distinctions,” as Mohanty explained, included “objective, self-subsistent, ideal meanings.” Ordinary words, such as *green* or *tree*, could have an objective sense, independent of speaker or context. But the meaning of the word *I* would vary with each speaker, and the meaning of *this* would vary at each occurrence. Frege’s triangle, which assigned a single objective sense to each word, could handle *green*, but not *I*. Husserl suggested that sense could be a function of both the word and its use, but that solution conflicted with Frege’s goal of context-independent meaning.

In the *Tractatus*, Wittgenstein (1921) observed Frege’s restrictions, but in the preface to his second book, he called them “grave errors.” In his later philosophy, Wittgenstein focused on usage in context, especially in the social activities that give language its meaning. Syntax is valuable for precise, concise, fluent speech, but with enough context, the syntax can be garbled without destroying comprehension. For an infant, a foreigner, or an instant-message addict, context is more important than syntax.

Peirce handled context with a more general definition of sign, which includes the triangles of Aristotle, the Scholastics, and Frege as special cases. In the 1860s, he discovered a metalevel principle for generating triads when he and his father were studying Kant. To derive his twelve categories, Kant started with four major groups: *Quantity*, *Quality*, *Relation*, and *Modality*, each of which he divided in three. Relation, for example, produced the triad of *Inherence*, *Causality*, and *Community*. While searching for a deeper principle underlying Kant’s categories, Peirce noticed that Inherence could be defined by a monadic predicate that characterizes an entity by what it has in itself, independent of anything else; Causality requires a dyadic relation that characterizes some reaction between two entities; and Community requires a triadic relation that relates an entity (1) to a community (2) for some purpose (3). He generalized that observation to his trichotomy of *Firstness*, *Secondness*, and *Thirdness*.

For his theory of signs, Peirce adopted the trichotomy as the unifying theme that relates signs of all kinds to language and the world. Every sign is a triad that relates a perceptible *mark* (1), to another sign called its *interpretant* (2), which determines an existing or intended *object* (3). Following is one of Peirce’s most often quoted definitions:

A sign, or *representamen*, is something which stands to somebody for something in some respect or capacity. It addresses somebody, that is, creates in the mind of that person an equivalent sign, or perhaps a more developed sign. That sign which it creates I call the *interpretant* of the first sign. The sign stands for something, its *object*. It stands for that object, not in all respects, but in reference to a sort of idea, which I have sometimes called the *ground* of the representamen. (CP 2.228)

A pattern of green and yellow in the lawn, for example, is a mark, and the interpretant is some type, such as Plant, Weed, Flower, SaladGreen, or Dandelion. The guiding idea that determines the interpretant depends on the context and the intentions of the observer. The interpretant determines the word the observer chooses to express the experience. As Peirce noted, a listener who is an expert in the subject matter can sometimes derive a richer interpretant than the speaker. Mohanty (1982:58) remarked “Not unlike Frege, Husserl would rather eliminate such fluctuations from scientific discourse, but both are forced to recognize their recalcitrant character for their theories and indispensability for natural languages.” Fortunately for science, theoreticians like Einstein and Bohr could often derive more meaning from scientific language than the authors intended.

Unlike Aristotle’s categories, which represent types of *existence*, Peirce’s *phenomenological* categories represent aspects of how something is perceived, conceived, or described. Any phenomenon can be described in all three ways. An object, for example, can be recognized as a dog, cat, man, or dandelion by directly observable properties of the individual (Firstness), but it cannot be recognized as a pet, stray, owner, or salad green without some evidence of an external relationship (Secondness). The corresponding Thirdness involves intentions that are indicated by signs, such as a contract or a habitual pattern of behavior.

Like nouns, verbs can also be classified by what aspect they describe: a directly observable event (Firstness); a causally related effect (Secondness); or a mediating intention (Thirdness). The next three sentences describe the same event in each of those ways:

1. Brutus *stabbed* Caesar.
2. Brutus *killed* Caesar.
3. Brutus *murdered* Caesar.

An act of stabbing can be recognized at the instant it happens. That is a classification by Firstness, since no other events or mental attitudes are involved. But an act of stabbing cannot be called killing unless a second event of dying occurs. Murder involves Thirdness because the stabbing (1) is related to the dying (2) by the intention (3). Determining whether an act of stabbing that resulted in killing should be considered a murder depends on subtle clues, whose interpretation may require a judge, a jury, and a lengthy trial.

| | 1. Quality | 2. Indexicality | 3. Mediation |
|----------------------|------------------------------------------------------------------------------|-------------------------------------------------------------------------------------|------------------------------------------------------------------------------|
| 1. Material | <p>Qualisign <i>A quality which is a sign.</i></p> | <p>Sinsign <i>An actual existent thing or event which is a sign.</i></p> | <p>Legisign <i>A law which is a sign.</i></p> |
| 2. Relational | <p>Icon <i>Refers by virtue of some similarity to object.</i></p> | <p>Index <i>Refers by virtue of being affected by object.</i></p> | <p>Symbol <i>Refers by virtue of some law or association.</i></p> |
| 3. Formal | <p>Rheme <i>A sign of qualitative possibility.</i></p> | <p>Dicent Sign <i>A sign of actual existence.</i></p> | <p>Argument <i>A sign of law.</i></p> |

Figure 2. Peirce’s triple trichotomy

On the surface, Peirce's triads seem similar to Aristotle's and Frege's. The difference, however, is that each of the three terms — the sign, the interpretant, and the object — can be further analyzed by the same metalevel principle. By analyzing the method by which the sign determines its object, Peirce (1867) derived the triad of *icon*, *index* and *symbol*: an icon refers by some similarity to the object; an index refers by a physical effect or connection; and a symbol refers by a habit or conventional association. Figure 2 shows this *relational* triad in the middle row.

After thinking about signs for another thirty years, Peirce discovered deeper patterns of relationships. He realized that the relational triad is based on Secondness — the relationship between a sign and its object. He therefore searched for two other triads. The signs in the first or *material* triad signify by the nature of the sign itself; those in the third or *formal* triad signify by a formal rule that associates sign and object. The labels at the top of Figure 2 indicate how the sign directs attention to the object: by some quality of the sign itself, by some causal or pointing effect, or by some mediating law, habit, or convention. The following examples illustrate nine types of signs:

1. *Qualisign* (material quality). A ringing sound as an uninterpreted sensation.
2. *Sinsign* (material indexicality). A ringing sound that is recognized as coming from a telephone.
3. *Legisign* (material mediation). The convention that a ringing telephone means someone is trying to call.
4. *Icon* (relational quality). An image that resembles a telephone when used to indicate a telephone.
5. *Index* (relational indexicality). A finger pointing toward a telephone.
6. *Symbol* (relational mediation). A ringing sound on the radio that is used to suggest a telephone call.
7. *Rheme* (formal quality). A word, such as *telephone*, which can represent any telephone, real or imagined.
8. *Dicent Sign* (formal indexicality). A sentence that asserts an actual existence of some object or event: "You have a phone call from your mother."
9. *Argument* (formal mediation). A sequence of dicent signs that expresses a lawlike connection: "It may be an emergency. Therefore, you should answer the phone."

Peirce coined the term *indexical* for the words or *deictics* that have the effect of a pointing finger or other kind of index. Instead of being troublesome exceptions, as they were for Frege and Husserl, indexicals become an integral part of a systematic framework.

The nine categories in Figure 2 are more finely differentiated than most definitions of signs, and they cover a broader range of phenomena. Anything that exists can be a sign of itself (*sinsign*), if it is interpreted by an observer. But Peirce (1911:33) did not limit his definition to human minds or even to signs that exist in our universe:

A sign, then, is anything whatsoever — whether an Actual or a May-be or a Would-be — which affects a mind, its Interpreter, and draws that interpreter's attention to some Object (whether Actual, May-be, or Would-be) which *has already* come within the sphere of his experience.

The mind or quasi-mind that interprets a sign need not be human. In various examples, Peirce mentioned dogs, parrots, and bees. A dog, like its owner, could experience ringing as a *qualisign* and recognize it as a sound from a particular source (*sinsign*). By a kind of Pavlovian conditioning, a dog could be taught a *legisign* to answer a specially-designed telephone. But an intelligent dog might

discover for itself that a ringing phone is an index of its owner's habit (another legisign) of running to answer it. In fact, there is anecdotal evidence that some parrots imitate a phone ring in order to summon their owners. Higher animals typically recognize icons and indexes, and some might recognize symbols. A language of some kind is a prerequisite for signs at the formal level of rhemes, dicent signs, and arguments. Whether dolphins or trained apes have a language adequate to express such signs is still an open question.

By building on and extending the semiotics of Aristotle and the Scholastics, Peirce avoided the dangers of psychologism. As he said, "Thought is not necessarily connected with a brain" (CP 4.551), but every thought is a sign, and every sign depends on some mind or quasi-mind. In the following definition, Peirce emphasized the independence of signs on any particular implementation:

I define a sign as something, A, which brings something, B, its interpretant, into the same sort of correspondence with something, C, its object, as that in which it itself stands to C. In this definition I make no more reference to anything like the human mind than I do when I define a line as the place within which a particle lies during a lapse of time. (1902:235)

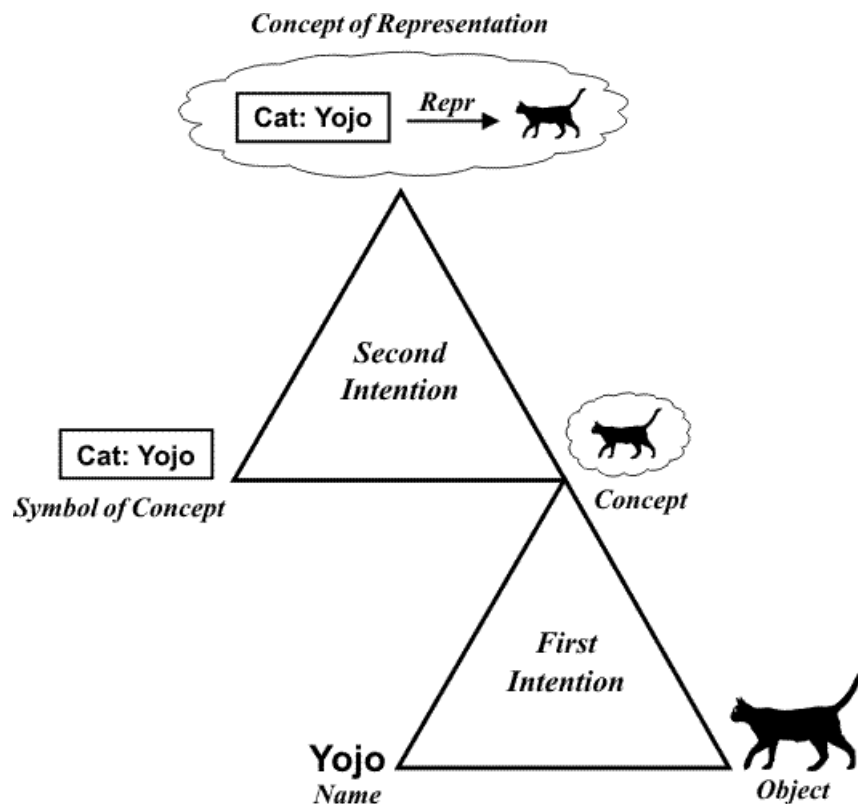


Figure 3. Meaning triangles for the concept of representation

As an example, Figure 3 illustrates the concept of representation by means of two meaning triangles. The first-intentional triangle at the bottom shows that the name Yojo refers to a cat illustrated by an image at the bottom right. The peak of that triangle is a concept illustrated by the same image enclosed in a balloon. The second-intentional triangle at the top, shows that the symbol **Cat: Yojo** refers to the same concept that is shared with the peak of the first-intentional triangle. The uppermost balloon illustrates a concept of representation that relates the symbol **Cat: Yojo** to the concept of the same cat. To explain how language is learned and used, this article includes many second-intentional statements that express such concepts.

Of the five signs shown in Figure 3, three have physical marks and two are mental concepts, which are shown in balloons. Although the mental concepts are internal, no introspection is needed to infer their existence. The triangles are based on semiotic principles, which, as Peirce said, would hold for “any scientific intelligence” — human, nonhuman, extraterrestrial, or even artificial. Humans and apes can understand one another because they have similar bodies and live in similar environments. Dolphins, however, are intelligent mammals with an utterly different method of communication about a radically different environment. But if they are capable of thinking about their signaling system, diagrams such as Figure 3 would characterize their thoughts. The same principles would hold for beings from another galaxy whose biology had no resemblance to anything on earth. The meaning triangles are as formal as mathematics, which is independent of psychology, biology, and physics.

As these examples show, Peirce’s theory of signs provides a more nuanced basis for analysis than the all-or-nothing question of whether animals have language. Unlike the static meaning triangles of Aristotle or Frege, the most important aspect of Peirce’s triangles is their dynamic nature: any node can spawn another triad to show three different perspectives on the entity represented by that node. During the course of a conversation, the motives of the participants lead the thread of topics from triangle to triangle. Understanding the interconnections of themes and motives requires a map of those triangles.

3. Twentieth-Century Theories of Language

In the *Tractatus*, Wittgenstein began the practice of treating natural language as a version of logic, but he rejected that view in his second book. Chomsky (1957) revived the formal approach by defining a language as the set of sentences generated by a formal grammar. Besides a formal grammar, Montague (1970) added a *formal semantics* that was much richer than Wittgenstein’s early version. That kind of formalization was ideal for the artificial languages of logic and computer science, but many linguists rejected it as a distortion of natural language. Some logicians were also skeptical about attempts to formalize the notoriously ambiguous natural languages. Peter Geach, a logician and former student of Wittgenstein’s, derided Montague’s system as “Hollywood semantics.”

The logical positivists renamed the traditional language arts — grammar, logic, and rhetoric — with the terms *syntax*, *semantics*, and *pragmatics*. Of the three, syntax is a relatively compact subject, although there are many theories about how to represent it. Semantics and pragmatics, however, raise controversial issues about the meaning of “meaning” and the endless variety of uses for language:

1. Syntax defines grammar by some sort of rules or patterns. Around the fifth century BC, Panini produced one of the most detailed grammars of all time, with nearly 4000 rules for the syntax of Sanskrit. Bharati et al. (1995) claimed that Panini-style grammars have important advantages for highly inflected languages with a free word order, such as the languages of modern India. In the early part of the 20th century, Ajdukiewicz (1935) invented rule formats for *categorial grammar*, Tesnière (1959) invented *dependency grammar*, and Post (1943) developed *production rules*, which are used to specify *phrase-structure* grammars. Harris (1951) combined a phrase-structure base with *transformation rules*, which Chomsky (1957) adopted and elaborated in several variations. Finally, Chomsky (1995) settled on a *minimalist program*: express grammar as a set of constraints that determine the simplest mapping from a *conceptual-intentional* representation to the spoken forms of a language. Yet the early hopes for formal grammars that could define all and only the grammatical sentences of a language were never realized. For a more flexible and extensible method of accommodating novel patterns, Fillmore (1988) and Goldberg (1995) developed *construction grammars* as an open-ended system of form-meaning pairs.

2. Semantics, loosely speaking, is the study of meaning. But the meaning triangle (Figure 1) has three sides, and different studies typically emphasize one side or another: the link from words to the concepts they express; the link from words and sentences to objects and truth values; or the link from concepts to percepts of objects and actions upon them. In different books, Aristotle addressed all three aspects: expression (*lexis*), reason (*logos*), and thought (*dianoia*). Modern approaches also address them in different books, usually by different authors:
- *Lexical semantics*, according to Cruse (1986), is a “contextual approach,” which derives “information about a word’s meaning from its relations with actual and potential linguistic contexts.” That definition corresponds to the left side of the meaning triangle, which omits the connection between words and the objects they refer to. It is compatible with Saussure’s definition of language (*langue*) as “the whole set of linguistic habits, which allow an individual to understand and be understood” (1916). Lexicographers analyze a corpus of contextual citations and catalog the linguistic habits in lexicons, thesauri, and terminologies.
 - *Formal semantics* studies the logical properties of words and sentences and relates them to objects and configurations of objects. The first logic-based systems were designed as computer implementations (Bohnert & Backer 1967; Woods 1968; Winograd 1972), but Montague’s theories were more influential among philosophers and logicians. Other formalisms include *discourse representation theory* (Kamp & Reyle 1993) and *situation semantics* (Barwise & Perry 1983). Yet despite 40 years of sustained research, none of the implementations can translate one page from an ordinary textbook to any version of logic. Lexical semantics covers a broader range of language than the formal versions, and it addresses more aspects of syntax and vocabulary that affect meaning. But unlike the logic-based theories, lexical semantics does not define a mapping from language to objects or a method of reasoning about them.
 - *Cognitive semantics* studies the concepts and patterns of concepts that relate language to perception and action. Locke’s associations influenced many 19th-century psychologists, but Kant’s *schemata* led to more structured theories by Selz (1913) and Bartlett (1932). Other versions included *Gestalt theory* (Wertheimer 1925), *activity theory* (Vygotsky 1934), and *cognitive maps* (Tolman 1948). The earliest computer implementations, called *semantic networks*, were designed for machine translation; among the first were the *correlational nets* by Ceccato (1961). Other highly influential computational versions include *conceptual dependencies* by Schank (1975), *chunks* by Newell and Simon (1972), who cited Selz as an inspiration, and *frames* by Minsky (1975), who cited Bartlett. Robotics applications use concepts and cognitive maps to relate a robot’s language interface to its sensory and motor mechanisms. Among linguists, Lakoff (1987), Langacker (1999), Talmy (2000), and Wierzbicka (1996) devoted their careers to analyzing cross-linguistic cognitive patterns and their relationship to extralinguistic objects and activities. The term *conceptual structure* is commonly used for those patterns, both in linguistics (Jackendoff 1983) and in artificial intelligence (Sowa 1976, 1984).
3. Pragmatics or rhetoric analyzes the use of language for some purpose. Like semantics, pragmatics can be studied from different perspectives: the structure of a text or discourse; the intentions of the speaker or author; or the function of language in its social context. Unlike the single semantic triad, the intentions of two or more participants in a social setting can entangle the pragmatic triad with more triads and subtriads, as in Figures 2 and 3. The plots of literary and historical narratives illustrate the complexity that can develop from a clash of multiple perspectives and motivations. Following is a brief summary:

- *Discourse, narrative, and argumentation* depend on structural patterns that organize strings of sentences into a coherent whole that may range in length from a paragraph to a book. Aristotle began the systematic study of such patterns in his books *Rhetoric* and *Poetics*. Propp (1928) classified the thematic patterns of folktales. Lord (1960) and Parry pioneered the study of oral epics, related them to the classical epics, and demonstrated the formulaic patterning at every level from short phrases to global themes. With their theory of *scripts*, Schank and Abelson (1977) developed methods for representing thematic patterns and analyzing them in computer programs. Despite many centuries of analysis by literary critics, the technology for automatically recognizing these patterns or generating them from a knowledge base lags far behind the work on syntax and semantics.
- *Motivation*, which depends on feelings and emotion, determines the direction of discourse and the choice of patterns to accomplish the desired effect. Aristotle considered desire (*orexis*) the ultimate source of direction, but he distinguished three kinds of desire: appetite (*epithymia*), passion (*thymos*), and will (*boulēsis*). He classified appetite and passion as feelings shared with beasts and will as the result of rational thought. In classifying emotions, Arieti (1978) proposed three categories, which resemble a Peircean triad: First-order or *protoemotions* — tension, appetite, fear, rage, and satisfaction — are feelings that arise from basic mechanisms, such as hormones. Second-order emotions associate protoemotions with imagined objects and events: anxiety, anger, wishing, and security. Third-order emotions are mediating processes that relate feelings of all kinds to past experiences and future expectations: love, hate, joy, and sadness. Without emotions, an intelligent system, no matter how logical, would have no reason to do or say anything. To design computational mechanisms that could provide some direction, Minsky (2006) proposed an *emotion machine* that would integrate intelligence with motivation.
- *Social interaction* is the ultimate ground of communication, and language develops in social activity. One of the most significant insights into language development and use was Wittgenstein's view of language games as an integral part of behavior. It had a strong influence on many later developments, including *speech acts* (Austin 1962), *conversational implicatures* (Grice 1975), and *relevance* (Sperber & Wilson 1986). One of the linguists influenced by Wittgenstein was Halliday (1978, 1999), whose *systemic-functional* approach was implemented in Winograd's early logic-based system and in *rhetorical structure theory* (Mann & Thompson 1988), which is widely used for discourse analysis and generation.

As this summary indicates, many aspects of syntax and semantics can be formalized, but no current formalism is as dynamic and flexible as language. The far more complex subject of pragmatics, however, is key to understanding the nature, origin, and function of language. Infants use language to satisfy their needs with their very first words, the protolanguage of early hominins probably had little or no syntax, and foreigners with a rudimentary knowledge of the local language can communicate effectively with gestures and isolated words. The idea that syntax is the foundation began with the logical positivists, who focused on written symbols as concrete “observables.” To avoid the complex relationship of language to the world, they replaced the world with abstract sets, which serve as surrogates for the messy objects, events, and people. They left pragmatics as an afterthought, which was almost totally ignored in the writings by Carnap, Tarski, Montague, and Quine. Yet any formal definition that cannot be adequately explained in the semantics and pragmatics of ordinary language is destined to be misunderstood by the people who need it the most.

4. A Neo-Wittgensteinian Approach

In the *Tractatus*, Wittgenstein claimed that “the totality of facts” about the world can be stated clearly in language or logic, and “Whereof one cannot speak, thereof one must be silent.” That book set the agenda for formal semantics in the 20th century. Yet those formal systems were also brittle, inflexible, and incapable of representing the kinds of language that people normally speak and write. In his later philosophy, Wittgenstein replaced the monolithic logic and ontology of his first book with an open-ended family of language games. As an alternative to definitions by necessary and sufficient conditions, he used the term *family resemblance* for the “complicated network of overlapping and criss-crossing similarities, in the large and the small” (1953, §66). Unlike his mentors, Frege and Russell, he did not consider vagueness a defect in language:

One might say that the concept ‘game’ is a concept with blurred edges. — “But is a blurred concept a concept at all?” — Is an indistinct photograph a picture of a person at all? Is it even always an advantage to replace an indistinct picture with a sharp one? Isn’t the indistinct one often exactly what we need?

Frege compares a concept to an area and says that an area with vague boundaries cannot be called an area at all. This presumably means that we cannot do anything with it. — But is it senseless to say: “Stand roughly (ungefähr) there”? (§71).

Frege’s view is incompatible with natural languages and with every branch of empirical science and engineering. Vagueness is not the result of a careless use of language, and it cannot be eliminated by replacing natural languages with artificial languages. Its ultimate source is the attempt to describe the continuous physical world with a finite vocabulary of discrete symbols. For representing vagueness, fuzzy logic (Zadeh 1975) and logics of ambiguity (van Deemter & Peters 1996) have been useful for some applications, but they don’t address the core issues. Vagueness caused by semantic discrepancies cannot be remedied by context-independent syntactic rules.

Natural languages can be as precise as any formal language or as vague as necessary in planning, negotiation, debate, and empirical investigation. In a formal language, the meaning of a sentence is completely determined by its *form* or syntax together with the meaning of its components. In this sense, natural languages are informal because the meanings of nearly all sentences depend on the situation in which they’re spoken, the background knowledge of the speaker, and the speaker’s assumptions about the background knowledge of the listeners. Since nobody ever has perfect knowledge of anyone else’s background, communication in natural language is an error-prone process that requires frequent questions, explanations, objections, concessions, distinctions, and stipulations. Precision and clarity are the goal of analysis, not the starting point. Whitehead (1937) aptly summarized that point:

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.

During his career as an experimental physicist and a practicing engineer, Peirce learned the difficulty of stating any general principle with absolute precision:

It is easy to speak with precision upon a general theme. Only, one must commonly surrender all ambition to be certain. It is equally easy to be certain. One has only to be sufficiently vague. It is not so difficult to be pretty precise and fairly certain at once about a very narrow subject. (CP 4.237)

This quotation summarizes the futility of any attempt to develop a precisely defined ontology of everything, but it offers two useful alternatives: an informal classification, such as a thesaurus or terminology designed for human readers; and an open-ended collection of formal theories about narrowly delimited subjects. It also raises the questions of how and whether the informal resources might be used as a bridge between informal natural language and formally defined logics and programming languages.

A novel theory of semantics, influenced by Wittgenstein's language games and related developments in cognitive science, is the *dynamic construal of meaning* (DCM) proposed by Cruse (2002). The fundamental assumption of DCM is that the most stable aspect of a word is its spoken or written sign; its meaning is unstable and dynamically evolving as it is used in different contexts or language games. Cruse coined the term *microsense* for each subtle variation in meaning. That is an independent rediscovery of Peirce's view: sign types are stable, but each interpretation of a sign token depends on its context in a pattern of other signs, the physical environment, and the background knowledge of the interpreter. Croft and Cruse (2004) suggested an integration of DCM semantics with a version of construction grammar, but their definitions are not sufficiently detailed for a computer implementation.

A computable method directly inspired by Wittgenstein's language games was developed by Margaret Masterman, one of six students in his course of 1933-34 whose notes were compiled as *The Blue Book*. In the late 1950s, Masterman founded the Cambridge Language Research Unit (CLRU) as a discussion group, which became one of the pioneering centers of research in computational linguistics. Her collected papers (Masterman 2005) present an approach with many similarities to DCM:

- A focus on semantics, not syntax, as the foundation for language: "I want to pick up the relevant basic-situation-referring habits of a language in preference to its grammar" (p. 200).
- A context-dependent classification scheme with three kinds of structures: a thesaurus with groups of words organized by areas of use, a *word fan* radiating from each word type to each area of the thesaurus in which it occurs, and dynamically generated combinations of fans for the word tokens of a text.
- Emphasis on images as a language-independent foundation for meaning with a small number (about 50 to 100) of combining elements represented by ideographs or monosyllables, such as IN, UP, MUCH, THING, STUFF, MAN, BEAST, PLANT, DO.
- Recognition that analogy and metaphor are fundamental to the creation of novel uses of language, especially in the most advanced areas of science. Maxwell's elegant mathematics, for example, is the final stage of a lengthy process that began with Faraday's analogies, diagrams, and vague talk about lines of force in electric and magnetic fields.

Figure 4 shows a word fan for *bank* with links to each area in Roget's *Thesaurus* in which the word occurs (p. 288). The numbers and labels identify areas in the thesaurus, which, Masterman claimed, correspond to "Neo-Wittgensteinian families."

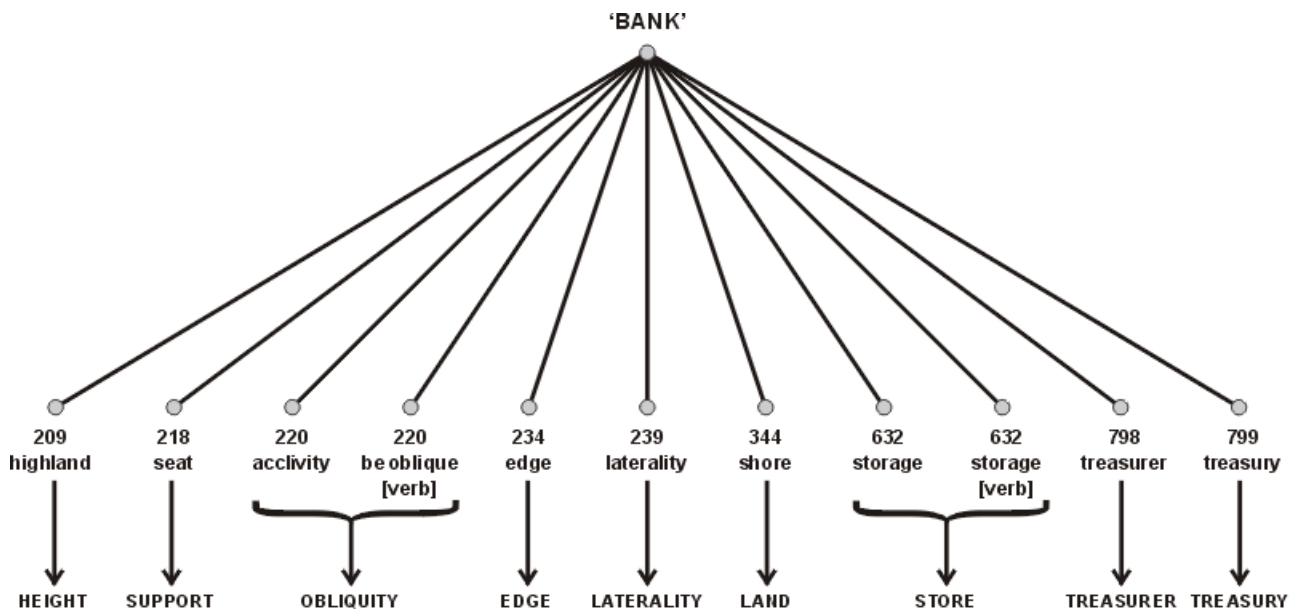


Figure 4: A word fan for *bank*

To illustrate the use of word fans, Masterman analyzed the phrases *up the steep bank* and *in the savings bank*. All the words except *the* would have similar fans, and her algorithm would “pare down” the ambiguities “by retaining only the spokes that retain ideas which occur in each.” For this example, it would retain “OBLIQUITY 220 in *steep* and *bank*; whereas it retains as common between *savings* and *bank* both of the two areas STORE 632 and TREASURY 799.”

Although Masterman’s work is over forty years old, it has some claim to be a more plausible cognitive theory than systems of “mentalese” for a private language of thought:

1. Emphasis on images as the primary semantic representation.
2. Actual English words as the units of meaning rather than abstract or artificial markers, features, or categories.
3. Language games as the basis for organizing and using patterns of words.
4. A context-dependent organization of word senses by usage (thesaurus style) instead of the more common dictionary style of grouping all senses of each word type.
5. Word fans as a secondary (context-independent) method for finding all senses of a word type when the area of usage is not known.

These principles are compatible with Wittgenstein’s later philosophy, but more is needed to capture the dynamics of the games, the way they mutate and evolve, and their relationships to one another. Halliday was a cofounder of CLRU who explored other aspects of language games with his emphasis on the use of language in social interactions. Yet neither Masterman nor Halliday addressed the language games of mathematics and logic, which were Wittgenstein’s starting point and a topic he addressed repeatedly in his later teaching, writing, and notebooks. To support both formal and informal languages, Masterman’s word fans must be extended with links to logic, but in a flexible way that permits an open-ended variety of options.

5. Steps Toward Formalization

In his *Philosophical Remarks* from the transitional period of 1929-30, Wittgenstein analyzed some “minor” inconsistencies in the *Tractatus*. His analysis led to innovations that form a bridge between his early system and the far more flexible language games. Shanker (1987) noted two new terms that are key to Wittgenstein’s transition:

1. *Satzsystem*: a system of sentences or propositions stated in a given syntax and vocabulary.
2. *Beweissystem*: a proof system that defines a logic for a Satzsystem.

Formally, the combination of a Satzsystem with a Beweissystem corresponds to what logicians call a *theory* — the deductive closure of a set of axioms. Informally, Wittgenstein’s remarks about Satzsysteme are compatible with his later discussions of language games. In conversations reported by Waismann (1979:48), Wittgenstein said that outside a Satzsystem, a word is like “a wheel turning idly.” Instead of a separate mapping of each proposition to reality, as in the *Tractatus*, the Satzsystem is mapped as a complete structure: “The Satzsystem is like a ruler (Maßstab) laid against reality. An entire system of propositions is now compared to reality, not a single proposition.” (Wittgenstein 1964, §82).

For a given logic (Beweissystem), each Satzsystem can be formalized as a theory that defines the ontology of a narrow subject. The multiplicity of Satzsysteme implies that any word that is used in more than one system will have a different sense in each. For natural languages, that principle is far more realistic than the monolithic logic and ontology of the *Tractatus*. Yet Wittgenstein illustrated his *Philosophical Remarks* primarily with mathematical examples. That *turning point*, as Shanker called it, implies that the goal of a unified foundation for all of mathematics, as stated in the *Principia Mathematica*, is impossible. The implication alarmed Russell, who observed “The theories contained in this new work of Wittgenstein’s are novel, very original, and indubitably important. Whether they are true, I do not know. As a logician who likes simplicity, I should wish to think that they are not.”

From the mid 1930s to the end of his life, Wittgenstein focused on language games as a more general basis for a theory of meaning. But he continued to teach and write on mathematical topics, and he compared language games to the multiple ways of using words such as *number* in mathematics: “We can get a rough picture of [the variety of language games] from the changes in mathematics.” These remarks imply that Satzsysteme can be considered specialized language games. The crucial addition for natural language is the intimate integration of language games with social activity and even the “form of life.” As Wittgenstein said in his notebooks, language is an “extension of primitive behavior. (For our language game is behavior.)” (Zettel, §545) The meaning of a word, a chess piece, or a mathematical symbol is its use in a game — a Sprachspiel or a Beweissystem.

A formal definition of language game is probably impossible, primarily because the games are integrated with every aspect of life. Even an informal characterization is difficult because Wittgenstein traversed many different academic boundaries in his examples: syntax, semantics, pragmatics, logic, ontology, speech acts, scenarios, sublanguage, and genre. As he admitted, “the very nature of the investigation... compels us to travel over a wide field of thought criss-cross in every direction... The same or almost the same points were always being approached afresh from different directions, and new sketches made” (1953, Preface). In effect, each Satzsystem uses a formal logic to define a formal ontology that may be used in one or more language games. But even a simple language game, which might use the ontology of a single Satzsystem, has a pragmatics that integrates the moves of the game with human behavior in a social setting. Chess, for example, is a game that can be formalized and played by a computer, but no computer experiences the struggle of combat, the joy of winning, or the disappointment of losing. Those experiences determine human intentionality, which might be represented in an ontology or *metalevel* ontology of a language game for talking about language games and the people who play them.

As formal theories, Satzsysteme require a metalevel theory to support reasoning about theories and their relationships with language games. The basis for such reasoning is generalization and specialization: a theory can be specialized by adding detail (more axioms) or generalized by deleting axioms. For example, a theory that describes properties of all animals is more general than a theory that describes mammals, because any true statement about all animals must also be true about mammals. Similarly, a theory about dogs or cats is more specialized than a theory about mammals. The generalization-specialization operator defines a partial ordering of theories, which happens to form a *lattice*. General theories near the top of the lattice are, in Peirce's terms, "sufficiently vague" to characterize a wide range of subjects. Specialized theories at lower levels are sufficiently "narrow" to be "pretty precise and fairly certain" No formal logic can be truly vague, but the axioms of a theory may be underspecified to accommodate multiple options. When precision is necessary, a theory may be specialized by tightening the constraints and adding detail.

Theories in the lattice may be large or small, and they may be stated in any version of logic. For most practical applications, the ISO standard for Common Logic is sufficient (ISO/IEC 2007). If \mathbf{L} is the set of all possible theories expressed in a given logic, then the lattice over \mathbf{L} is specified by a partial ordering \leq and two dyadic operators \cap and \cup . Let x , y , and z be any theories in the lattice \mathbf{L} :

- If the theory x is true of a subset of the cases or models in which y is true, written $x \leq y$, then x is said to be a *specialization* of y , and y is said to be a *generalization* of x . Every theory is a generalization and a specialization of itself.
- The *supremum* of x and y , $x \cup y$, is their most specialized common generalization: $x \leq x \cup y$; $y \leq x \cup y$; and if $x \leq z$ and $y \leq z$, then $x \cup y \leq z$.
- The *infimum* of x and y , $x \cap y$, is their most general common specialization: $x \cap y \leq x$; $x \cap y \leq y$; and if $z \leq x$ and $z \leq y$, then $z \leq x \cap y$.
- The top of the lattice \top , called the *universal theory*, is a generalization of every theory. It contains all tautologies, and it is true of everything.
- The bottom of the lattice \perp , called the *absurd theory*, is a specialization of every theory. It contains every statement expressible in the given logic (including all contradictions), and it is true of nothing.

For every imaginable subject, every true or false theory x lies somewhere between the top and the bottom: $\perp \leq x \leq \top$. Even inconsistent theories are in the lattice, because they collapse into the absurd theory at the bottom. The complete lattice of all possible theories is infinite, but only a finite subset can ever be implemented in an actual system.

To relate language to logic, Masterman's word fans can link each word type to multiple word senses, each represented by a *monadic predicate* or *concept type*. Figure 5 illustrates a word fan that maps word types to concept types to *canonical graphs* and finally to a lattice of theories. In this article, the canonical graphs are represented as *conceptual graphs* (CGs), one of the three standard dialects of Common Logic. Equivalent operations may be performed with any notation for logic, but graphs have important formal advantages (Sowa 2008).

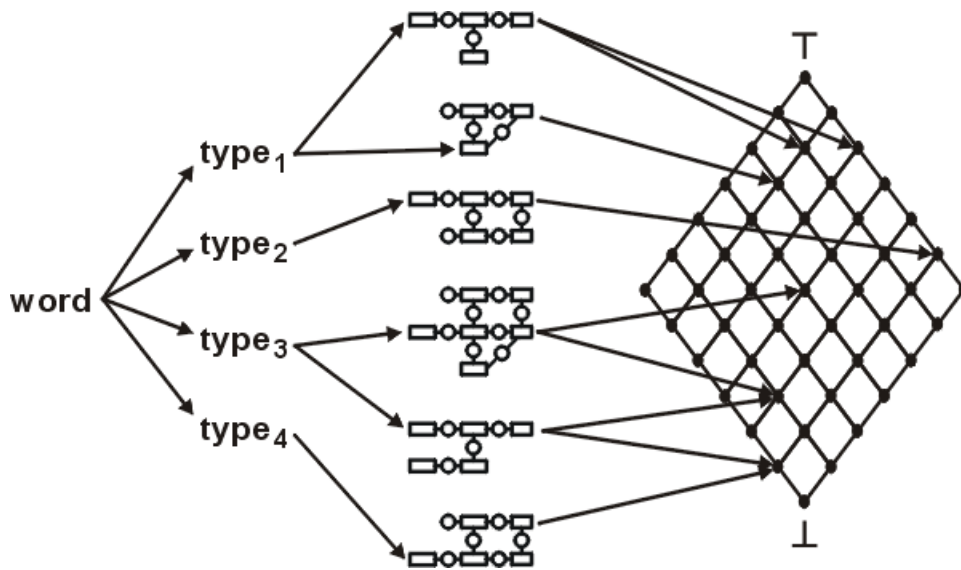


Figure 5: words → concept types → canonical graphs → lattice of theories

The fan on the left of Figure 5 links each word to an open-ended list of *concept types*, each of which corresponds to some area of a thesaurus in Masterman's system. The word *bank*, for example, could be linked to types with labels such as **Bank799** or **Bank_Treasury**. In various applications or language games, those types could be further subdivided into fine-grained subtypes, which would correspond to Cruse's microsenses. The selection of subtypes is determined by canonical graphs, which specify the characteristic patterns of concepts and relations associated with each type or subtype. Figure 6 illustrates three canonical graphs for the types **Give**, **Easy**, and **Eager**.




Figure 6: Canonical graphs for the types Give, Easy, and Eager

A canonical graph for a type is a conceptual graph that specifies one of the patterns characteristic of that type. On the left, the canonical graph for **Give** represents the same constraints as a typical *case frame* for a verb. It states that the agent (**Agnt**) must be **Animate**, the recipient (**Rcpt**) must be **Animate**, and the object (**Obj**) may be any **Entity**. The canonical graphs for **Easy** and **Eager**, however, illustrate the advantage of graphs over frames: a graph permits cycles, and the arcs can distinguish the directionality of the relations. Consider the following two sentences:

Bob is easy to please. Bob is eager to please.

For both sentences, the concept [**Person: Bob**] would be linked via the attribute relation (**Attr**) to the concept [**Easy**] or [**Eager**], and the act [**Please**] would be linked via the manner relation (**Manr**) to the same concept. But the canonical graph for **Easy** would make Bob the object of the act **Please**, and the graph for **Eager** would make Bob the agent. The first sentence below is acceptable because the object may be any entity, but the constraint that the agent of an act must be animate would make the second sentence unacceptable:

The book is easy to read. * *The book is eager to read.*

Chomsky (1965) used the easy/eager example to argue for different syntactic transformations associated with the two adjectives. But the canonical graphs state semantic constraints that cover a wider range of linguistic phenomena with simpler syntactic rules. A child learning a first language or an adult reading a foreign language can use semantic constraints to interpret sentences with unknown or even ungrammatical syntax. Under Chomsky's hypothesis that syntax is a prerequisite for semantics, such learning is inexplicable.

Canonical graphs with a few concept nodes are adequate to discriminate the general senses of most words, but the canonical graphs for detailed microsenses can become much more complex. For the adjective *easy*, the microsenses occur in very different patterns for a book that's easy to read, a person that's easy to please, or a car that's easy to drive. For the verb *give*, a large dictionary lists dozens of senses, and the number of microsenses is enormous. The prototypical act of giving is to hand something to someone, but a large object can be given just by pointing to it and saying "It's yours." When the gift is an action, as in giving a kiss, a kick, or a bath, the canonical graph used to parse the sentence has a few more nodes. But the graphs required to understand the implications of each type of action are far more complex, and they're related to the graphs for taking a bath or stealing a kiss.

The canonical graph for *buy* typically has two acts of giving: money from the buyer to the seller, and some goods from the seller to the buyer. But the canonical graphs needed to understand various microsenses may require far more detail about the buyers, the sellers, the goods sold, and other people, places, and things involved. Buying a computer, for example, can be done by clicking some boxes on a screen and typing the billing and shipping information. That process may trigger a series of international transactions, which can be viewed by going to the UPS web site to check when the computer was airmailed from Hong Kong and delivered to New York. All that detail is involved in one microsense of the verb *buy*. In a successful transaction, the buyer can ignore most of it, but somebody must be able to trace the steps if something goes wrong.

6. Formal and Informal Reasoning

Historically, Aristotle developed formal logic and ontology as an abstraction from the arguments and patterns of reasoning in ordinary thought and language. In his pioneering work on symbolic logic, Boole (1854) continued the tradition by calling his version *The Laws of Thought*. Frege (1879) developed a complete system of *first-order logic* with a tree notation, but Peirce (1880, 1885) extended Boolean algebra to a version that Peano adopted for *predicate calculus*. Although Peirce gave equal attention to applications in language and mathematics, most 20th-century logicians emphasized mathematics to the almost complete exclusion of other uses for logic. Montague tried to force natural language semantics into the same deductive forms, and other logicians extended his approach. Finally, Kamp (2001), a former student of Montague's, admitted

Natural language semantics increasingly takes on the complexion of a branch of a general theory of information representation and transformation. The role of logical inference in the processes of linguistic interpretation indicates an interleaving of inferential and other representation-manipulating operations. This suggests that the inferential relations and operations that have often been considered the essence of logic are better seen as an integral part of a wider repertoire. Thus logic comes to look much more like a general theory of information, than as a discipline concerned more or less exclusively with deduction.

Although Peirce knew that deduction is important, he realized that it can only derive the implications of already available premises. Without some method for deriving the premises, deduction is useless. He

observed that two other methods, induction and abduction, are required for deriving the starting assumptions:

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.*
Tweety, Polly, and Hooty fly.
 Assume: *Every bird flies.*

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

Given: *Vampy flies. Vampy is a bat. Vampy and Tweety have wings.*
 Guess: *Every animal with wings flies.*

Deduction, which is the most precise and disciplined method of reasoning, is the only method that is certain. But as Peirce observed, discipline is “purely inhibitory. It originates nothing” (CP 5.194). The patterns of induction and abduction, which can derive new premises, are, at best, methods of plausible reasoning. Aristotle listed those methods among the fallacies, but he admitted that they were necessary for deriving theories from empirical data. In his skeptical writings, Sextus Empiricus noted that deduction in mathematics could be certain, but any generalization about the physical world, such as *Every bird flies*, would have to be derived by induction from observations. Any conclusion derived by deduction would be just as uncertain as the premises derived by induction and abduction.




Figure 7: The world, a model, and a theory

Figure 7 shows the source of vagueness and uncertainty in the discrepancies between the physical world and Tarski-style models about the world. On the left is an illustration of the world. On the right is a theory stated in some version of logic. In the center is the kind of model Tarski (1933) assumed: a set of entities, represented by dots, and relationships between entities represented by lines. For each sentence in the theory, Tarski’s evaluation function would determine a unique truth value in terms of the model. Scientific methodology, including theories of probability or fuzziness, would evaluate the degree of approximation of the model to the world.

The evolution of Wittgenstein's thought can be summarized by the way the model is related to the world. In the *Tractatus*, he considered the model to be an exact picture of some aspect of the world: the dots and lines are isomorphic to objects in the world and atomic facts that relate objects. Each true sentence is a Boolean combination of atomic facts that makes a true statement about the world: "Everything that can be said can be said clearly" (§4.116), and "Whereof one cannot speak, thereof one must be silent" (§7.0). In his transitional period, Wittgenstein weakened the assumption by claiming that the theory as a whole (the Satzsystem) can be mapped to the world, but not each sentence by itself. That implies that the model is a good picture of the world, but not all parts of the model have a one-to-one correspondence with objects and relations in the world. A vague imperative, such as "Stand roughly there," could be meaningful in a language game, even though the phrase *roughly there* does not determine a unique object. In his late philosophy, he dropped the requirement that words and sentences need any referents outside the pattern of behavior associated with the language game. Language games involving prayers, singing, or even nonsense syllables would acquire their meaning solely from the associated behavior, even though many words might have no observable referents. The new language games could accommodate all the sentences of the early philosophy, but as special cases, not as the mainstream.

Of the three methods of reasoning, induction draws a generalization from instances, but only abduction can introduce a truly novel idea. In Peirce's semiotics, abduction is the basis for creativity, and there is no need for Descartes's innate ideas or Kant's synthetic *a priori* judgments. In a computer implementation, abduction can be implemented by a process for selecting appropriate chunks of information from memory and reconfiguring them in a novel combination. It can be performed at various levels of complexity:

- *Reuse*. Search for a previous fact, rule, pattern, theory, or chunk that approximately matches the current context, problem, or goal.
- *Revise*. Modify any promising chunk from an approximate match to an exact match.
- *Combine*. Revise and combine as many theories and chunks as needed to solve the problem or reach the goal. If a new chunk is consistent with the current theory, conjunction is appropriate. If it is inconsistent, revise it.

These steps can be iterated indefinitely. After a hypothesis is formed by abduction, its implications must be tested against reality. If the implications are not confirmed, the hypothesis must be revised or replaced by another abduction. In Peirce's logic of pragmatism, the novel ideas 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)

Abduction reassembles previously observed elements in novel combinations. Each combination is a new concept, whose practical meaning is determined by the totality of purposive actions it implies — in Wittgenstein's terms, by the language games associated with those actions. As Peirce said, meanings grow as new information is received, new implications are derived, and new actions become possible. Unlike other logicians, Peirce put learning at the center of his system and assumed a continuity from an infant's early experience to the most sophisticated theories of science. There are three kinds of learning, each of which modifies an old theory to create a new theory:

1. *Rote*. Rote learning accumulates information by adding low-level facts about instances. With only rote learning, each fact is a new axiom, no deduction is possible, and the number of propositions in the theory is equal to the number of axioms.

2. *Induction*. The number of axioms is reduced when a general principle derived by induction subsumes multiple instances that can be derived as needed by deduction. In effect, a generalization, such as “Every bird flies,” compresses the data and reduces the memory load by enabling a smaller set of axioms to imply all the old statements.
3. *Abduction*. Like induction, abduction can also reduce the number of axioms. Unlike induction, which does not change the vocabulary, abduction can introduce new terminology for concepts that are not mentioned in the old axioms. Induction, for example, might lead to the generalization “Every bird and bat flies.” Abduction, however, forms new kinds of connections, which may require new terminology: “Every animal with wings flies.”

The lattice of theories provides a framework for analyzing and relating all the methods of learning and reasoning. Classical deduction is the only method that stays within the bounds of a single theory. All other methods can be viewed as plans or strategies for walking through the lattice of theories. Each step along the path adds, deletes, or changes some axiom to make a new theory that is more suitable for some purpose. Every method of learning specializes the starting theory by adding new axioms. Pure rote learning accumulates facts without relating them. In effect, each new fact adds a new axiom to the old theory, but it doesn't form connections to the previous axioms. Induction and abduction add general principles, which may make some old axioms redundant, but they increase the number of implications.

For reasoning about defaults and exceptions, nonmonotonic logics introduce new kinds of operators and rules of inference. Methods of belief revision achieve the same results by changing the theories instead of changing the logic (Makinson 2005; Peppas 2008). Like learning, belief revision can be treated as a walk through the lattice to select an appropriate revision of the current theory. The parallels between learning and nonmonotonic reasoning are appropriate because defaults and exceptions add new information that had not been incorporated in the old theory. Unlike learning, which always specializes a theory by adding axioms, belief revision may generalize a theory by deleting axioms or move sideways in the lattice by changing axioms. Figure 8 shows the four basic operators for navigating the lattice: *contraction*, *expansion*, *revision*, and *analogy*.




Figure 8: Four operators for navigating the lattice of theories

The operators of contraction and expansion follow the arcs of the lattice, revision makes short hops sideways, and analogy makes long-distance jumps. The first three operators obey the AGM axioms for belief revision (Alchourrón et al. 1985). The analogy operator (Sowa 2000) relabels one or more types or relations. If the original theory is consistent, any relabeled version is guaranteed to be consistent, but it may be located in a remote branch of the lattice. To illustrate the moves through the lattice, suppose that A is Newton's theory of gravitation applied to the earth revolving around the sun and F is Niels Bohr's theory about an electron revolving around the nucleus of a hydrogen atom. The path from A to F is a step-by-step transformation of the old theory to the new one. The revision step from A to C replaces the gravitational attraction between the earth and the sun with the electrical attraction between the electron and the proton. That step can be carried out in two intermediate steps:

- *Contraction.* Any theory can be contracted to a smaller, more general theory by deleting one or more axioms (and their implications). In the move from A to B, axioms for the gravitational force would be deleted. Contraction has the effect of blocking proofs that depend on the deleted axioms.
- *Expansion.* Any theory can be expanded to a larger, more specialized theory by adding one or more axioms. In the move from B to C, axioms for the electrical force would be added. The result of both moves is a substitution of electrical axioms for gravitational axioms.

Unlike contraction and expansion, which move to nearby theories in the lattice, analogy jumps to a remote theory, such as C to E, by systematically renaming the types, relations, and individuals that appear in the axioms: the earth is renamed the electron; the sun is renamed the nucleus; and the solar system is renamed the atom. Finally, the revision step from E to F uses a contraction step to discard details about the earth and sun that have become irrelevant, followed by an expansion step to add new axioms for quantum mechanics.

Formal and informal reasoning should not be considered incompatible or conflicting. Instead, formal reasoning is a more disciplined application of the techniques used for informal reasoning. Analogy, the process of finding common patterns in different structures, is the foundation for both. The logical methods of reasoning by induction, deduction, and abduction are distinguished by the constraints they impose on analogy:

- *Deduction.* A basic rule of deduction is *modus ponens*: given an assertion p and an axiom of the form p implies q , derive the conclusion q . In most applications, the assertion p is not identical to the p in the axiom, and a version of analogy is necessary to *unify* the two ps before the rule can be applied. The most time-consuming task is not the application of a single rule, but the repeated use of analogies for finding patterns that may lead to successful rule applications.
- *Induction.* When every instance of p is followed by an instance of q , induction may derive the general principle that p implies q . Since the ps and qs are rarely identical in every occurrence, the generalization requires a version of analogy that subsumes all the instances.
- *Abduction.* The operation of guessing or forming an initial hypothesis by abduction uses the least constrained version of analogy, in which some parts of the matching structures may be more generalized while other parts are more specialized.

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.” The basis for analogy is pattern matching, which is also used in more disciplined versions of reasoning. But analogy is also used in *case-based reasoning*, which can have the combined effect of deduction, induction, and abduction.

7. Foundations for Dynamic Ontology

In every science, the vocabulary and theories are developed in close connection with the data on which they are based. Although they may be distinguished during the analysis, the data and the vocabulary influence one another inextricably. No ontology, formal or informal, is independent of the vocabulary and the methodologies (i.e., language games) used to analyze the data. Natural language terms have been the starting point for every ontology from Aristotle to the present. Even the most abstract ontologies of mathematics and science are analyzed, debated, explained, and taught in natural languages. For computer applications, the users who enter data and choose options on menus, think in the words of the NL vocabulary. Any options that cannot be explained in words the users understand

are open invitations to mistakes, confusions, and system vulnerabilities. Therefore, every ontology that has any practical application must have a mapping, direct or indirect, to and from natural languages.

These observations imply that any foundation for applied ontology must support a systematic mapping to and from natural languages. Yet both logicians and linguists recognize that Montague's claim that there is no difference between the semantics of formal and natural languages is false. In a series of studies with a strong Wittgensteinian orientation, Wilks (2006, 2008a,b) has shown that the widely used linguistic resources such as WordNet (Miller 1995) and Roget's *Thesaurus* are fundamentally different from and irreconcilable with the axiomatized ontologies expressed in formal logics. The linguistic resources blur distinctions that are critical for the ontologies. Even worse, translations from language to logic based on them introduce contradictions that would make a theorem prover unusable. Furthermore, there is no way to "correct" them in order to align them with formal ontologies because any such correction would make them unusable for their primary purpose of interpreting a broad range of natural language texts. Attempted alignments merely blur critical distinctions and introduce contradictions in the formal systems. Using WordNet to align independently developed ontologies is as futile as using it to align Linux with Microsoft Windows and Apple's OS X.

The comparison with operating systems is significant because every program or collection of programs is a formal system that could be described by some theory in the infinite lattice. Such a theory would be a detailed formal specification of the program, with preconditions and postconditions for every change to the data. Every revision or update to the programs would be described by a different theory somewhere in the lattice. Even a fix to a bug creates another program described by another theory. For minor revisions, all the programs would be described by theories located close to one another in the lattice. The lattice operators could be used to organize and relate different versions of the systems. For example, suppose that some program X has multiple versions distinguished by numbers such as 4.1, 4.2, 4.21:

- In usual practice, version 4.x would be more similar to version 4.y than either would be to any version 3.z.
- For any version X.Y, let ThX.Y be the theory that describes it.
- For any two versions X.Y and Z.W, the supremum $\text{ThX.Y} \cap \text{ThZ.W}$ would be a generalization that correctly describes all the features common to both versions and was silent about their differences.
- For a major version number, such as 4.x, the least common generalization of Th4.0, Th4.1, ..., Th4.9 would be a theory that describes all and only those features common to every version. Any other program that uses only those features would be guaranteed to run correctly with every version from 4.0 to 4.9.

This summary, which characterizes common programming practice, is just as applicable to any application of formal ontology. The lattice of theories is key to implementing the approach.

The words used in programming illustrate the variations in terminology that typically occur in all engineering applications. The word *file*, for example, is commonly used to describe a collection of data managed by an operating system. But every version of every operating system has a different definition. In Linux and other UNIX-like systems, a file contains an ordered list of character strings separated by newline characters. In Apple's operating systems, the strings of a file are separated by carriage-return characters. In all versions of Windows, the strings are separated by two characters — a carriage-return followed by a newline character. In IBM's mainframe operating systems, the strings are called *records*, and a file with fixed-length records might have no character strings between records. This example is just one way in which the meaning of the word *file* differs among systems, and each

version of every operating system for the past half century has added new microsenses to the meaning of all such terms, including the term *operating system* itself. Yet a typical English dictionary lumps them all under a single word sense, such as definition 2c of the *Miriam-Webster Collegiate*: “a collection of related data records (as for a computer).” In computational linguistics, assigning word sense 2c to an occurrence of *file* is considered “disambiguation,” but it is only the first step toward mapping it to a formal ontology.

The problem of matching language to logic is unsolvable if the two are considered totally different, irreconcilable systems. Montague simplified the problem by adopting Wittgenstein’s early assumption: both language and logic can have a monolithic model-theoretic semantics, along the lines developed for formal logics. Forty years of research in logic, linguistics, and AI has not produced a successful implementation: no computer system based on that approach can read one page of a high-school textbook and use the results to answer the questions and solve the problems as well as a B student. Wittgenstein’s later philosophy makes the semantics of formal logic a special case of the much richer semantics of natural languages. Instead of a monolithic semantics, each language game has its own semantics that is intimately connected with the more general methods of perception and behavior. The diversity of mechanisms associated with language is a reflection of the diversity involved in all aspects of cognition. In his book *The Society of Mind*, Minsky (1987) surveyed that diversity and proposed an organization of active processes as a computational model that could simulate the complexity:

What magical trick makes us intelligent? The trick is that there is no trick. The power of intelligence stems from our vast diversity, not from any single, perfect principle. Our species has evolved many effective although imperfect methods, and each of us individually develops more on our own. Eventually, very few of our actions and decisions come to depend on any single mechanism. Instead, they emerge from conflicts and negotiations among societies of processes that constantly challenge one another. (§30.8)

This view is radically different from the assumption of a unified formal logic that cannot tolerate a single inconsistency. Minsky’s goal is to build a flexible, fault-tolerant system out of imperfect, possibly fallible components. Different processes supported by different components could implement different language games or even different aspects of the same language game. Versions of analogy would support the reasoning methods used by most or all of the processes, some of which might use the more disciplined methods of analogy called formal logic.

A foundation for ontology inspired by Peirce, Wittgenstein, and Minsky sounds intriguing, but any implementation would require more detailed specifications. The infinite lattice of theories is formally defined, but no actual implementation can be infinite. A less ambitious term is *open-ended hierarchy*, which implies that only a finite subset is ever implemented, the implemented theories do not form a complete lattice, but there is no limit to the possible revisions and extensions. Following is a proposed foundation for a dynamic ontology that would relate natural language terminologies to a hierarchy of formal theories related by generalization and specialization:

1. Separate, but related lexical and formal resources: *lexical resources* for natural language terminologies, syntactic information, informal semantic information, and links to the formal types and relations; and a *hierarchy of formal theories*, represented in any dialect of Common Logic, including subsets such as RDF(S) and OWL.
2. A *registry* for recording all the resources of point #1 and a *repository* or a distributed collection of repositories for storing them.
3. *Metadata* about the resources, including who developed them and used them for what purposes and evaluations of the results.

4. The lexical resources could contain a variety of terminologies, dictionaries, lexicons, grammars, and resources such as WordNet. The links to logic could be represented as word fans and canonical graphs, as in Figures 5 and 6, but they could use any dialect that could be mapped to Common Logic, including frame-like notations.
5. The hierarchy of formal theories would be organized in a partial ordering by generalization and specialization, but the full set of lattice operations need not be fully implemented. As time goes by, more gaps would be filled in and more of the implications of the partial ordering would be discovered.
6. The theories in the hierarchy could be of any size. Some would represent the *microtheories* of Cyc, which are devoted to specialized subjects, and others would resemble typical upper-level ontologies. Many of them would represent useful mathematical structures or standard relationships for times and dates, weights and measures, or geographical coordinates. Others might contain large mergers of a variety of other theories that can be used together. In general, any useful ontology of any size that is represented in any dialect of Common Logic could be contributed to the hierarchy.
7. The generalization hierarchy would be especially valuable for stating interoperability constraints and version controls. If the interfaces to two programs are compatible with theory X in the ontology, then they would be compatible with anything stated in any ontology that is a generalization of X. If one program is compatible with theory X and another with a different theory Y, then they could interoperate on anything stated in an ontology that is a common generalization of X and Y.

Nothing in this proposal requires any further research, since parts of every point above have been implemented in one or more working systems. This proposal does not require a unified upper-level ontology, but it allows any ontology with such an upper level to be registered in the hierarchy as one among many. It also allows specialized low-level theories, such as an ontology of times and dates, to be placed in the hierarchy as a specialization of many different upper levels. A program that uses times and dates, but no other specialized information, would then be able to interoperate with other systems that used different upper-level ontologies.

For linking lexical resources to formal ontologies, the hierarchy provides the mechanism for resolving the multiple microsenses. As an example, the dictionary definition 2c for the word *file* cannot distinguish the multiple microsenses for every operating system. That word sense could be linked to a simple, general theory about file systems as collections of records or strings without any detail about how the records are represented. Theories about the multiple versions of file systems for Unix, Microsoft, Apple, and IBM mainframes would be specializations of that theory further down in the hierarchy. The word *file* in the context of a Microsoft manual would narrow the meaning to a subtype in a theory several levels deeper in the hierarchy. The single definition 2c in the English dictionary would correspond to hundreds of different microsenses for every version of every operating system that ever used the word *file*.

The sharp distinction between lexical resources and a hierarchy of theories enables multiple agents to use them concurrently without requiring a prior merger into a single consistent theory. Majumdar and Sowa (2009) take advantage of that option in the VivoMind Language Processor (VLP), which implements a society of heterogeneous agents as suggested by Minsky (1987, 2006). During the process of language analysis and reasoning, different agents can use different dictionaries and ontologies to analyze different parts of the same text or even the same sentence. To avoid inconsistency, deductive reasoning must be restricted to a single theory with a single ontology. For abduction, however, different agents, using different ontologies, may propose new hypotheses derived from any

resource of any kind. Other agents can then test those hypotheses on competing interpretations of the current text. The VLP system does not yet implement the full range of ideas discussed in this article, but the framework has proved to be highly flexible and efficient for language analysis and reasoning.

References

- Ajdukiewicz, Kazimierz (1935) Die syntaktische Konnexität, *Studia Philosophica*, **1**, 1-27; translated as Syntactic Connexion, in S. McCall, ed., *Polish Logic 1920-1939*, Clarendon Press, Oxford, 1967.
- Alchourrón, Carlos, Peter Gärdenfors, & David Makinson (1985) On the logic of theory change: partial meet contraction and revision functions, *Journal of Symbolic Logic* **50:2**, 510-530.
- Arieti, Silvano (1978) The psychobiology of sadness, in S. Arieti & J. Bemporad, *Severe and Mild Depression*, Basic Books, New York, 109-128.
- Aristotle, *Works*, Loeb Library, Harvard University Press, Cambridge, MA.
- Austin, John L. (1962), *How to do Things with Words*, second edition edited by J. O. Urmson & Marina Sbisa, Harvard University Press, Cambridge, MA, 1975.
- Barwise, Jon, & John Perry (1983) *Situations and Attitudes*, MIT Press, Cambridge, MA.
- Bartlett, Frederic C. (1932) *Remembering*, Cambridge University Press, Cambridge.
- Bharati, Akshar, Vineet Chaitanya, & Rajeev Sangal (1995) *Natural Language Processing: A Paninian Perspective*, Prentice-Hall of India, New Delhi.
- Bohnert, H., & P. Backer (1967) Automatic English-to-logic translation in a simplified model, technical report RC-1744, IBM, Yorktown Heights, NY.
- Brentano, Franz (1874) *Psychologie vom empirischen Standpunkte*, translated as *Psychology from an Empirical Standpoint* by A. C. Rancurello, D. B. Terrell, & L. L. McAlister, Routledge, London.
- Bundy, Alan (2007) Cooperating Reasoning Processes: More than just the Sum of their Parts, *Proceedings of IJCAI-07*, pp. 2-11.
- Bundy, Alan, & Fiona McNeill (2006) Representation as a fluent: An AI challenge for the next half century, *IEEE Intelligent Systems* **21:3**, pp. 85-87.
- Ceccato, Silvio (1961) *Linguistic Analysis and Programming for Mechanical Translation*, Gordon and Breach, New York.
- Chomsky, Noam (1957) *Syntactic Structures*, Mouton, The Hague.
- Chomsky, Noam (1965) *Aspects of the Theory of Syntax*, MIT Press, Cambridge, MA.
- Chomsky, Noam (1995) *The Minimalist Program*, MIT Press, Cambridge, MA.
- Croft, William, & D. Alan Cruse (2004) *Cognitive Linguistics*, Cambridge University Press, Cambridge.
- Cruse, D. Alan (1986) *Lexical Semantics*, Cambridge University Press, New York.
- Cruse, D. Alan (2002) Microsenses, default specificity and the semantics-pragmatics boundary, *Axiomathes* **1**, 1-20.
- Fillmore, Charles (1988) The mechanisms of construction grammar, *Berkeley Linguistics Society* **35**, 35-55.
- Förster, Eckart (2000) *Kant's Final Synthesis: An Essay on the Opus postumum*, Harvard University Press, Cambridge, MA.
- Frege, Gottlob (1879) *Begriffsschrift*, English translation in J. van Heijenoort, ed. (1967) *From Frege to Gödel*, Harvard University Press, Cambridge, MA, 1-82.
- Grice, H. Paul (1975) Logic and conversation, in P. Cole & J. Morgan, eds., *Syntax and Semantics 3: Speech Acts*, Academic Press, New York, 41-58.
- Halliday, M.A.K. (1978) *Language as Social Semiotic: The Social Interpretation of Language and Meaning*, University Park Press, Baltimore.
- Halliday, M.A.K. & Christian M.I.M. Matthiessen (1999) *Construing Experience Through Meaning: A Language-Based Approach to Cognition*, Cassell, London.
- Harris, Zellig S. (1951) *Methods in Structural Linguistics*, Chicago.

- ISO/IEC (2007) *Common Logic (CL) — A Framework for a family of Logic-Based Languages*, IS 24707, International Organisation for Standardisation.
- Jackendoff, Ray (1983) *Semantics and Cognition*, MIT Press, Cambridge, MA.
- Kamp, Hans (2001) Levels of linguistic meaning and the logic of natural language, http://www.ilic.uva.nl/lia/farewell_kamp.html
- Kamp, Hans, & Uwe Reyle (1993) *From Discourse to Logic*, Kluwer, Dordrecht.
- Lakoff, George (1987) *Women, Fire, and Dangerous Things*, University of Chicago Press, Chicago.
- Langacker, Ronald W. (1999) *Grammar and Conceptualization*, de Gruyter, Berlin.
- Lenat, Douglas B. (1995) Cyc: A large-scale investment in knowledge infrastructure, *Communications of the ACM* **38:11**, 33-38.
- Lord, Albert B. (1960) *The Singer of Tales*, Harvard University Press, Cambridge, MA.
- Majumdar, Arun K., & John F. Sowa (2009) Two paradigms are better than one and multiple paradigms are even better, in S. Rudolph, F. Dau, and S.O. Kuznetsov, eds., *Proceedings of ICCS'09*, LNAI 5662, Springer, pp. 32-47.
- Makinson, David (2005) *Bridges from Classical to Nonmonotonic Logic*, King's College Publications, London.
- Mann, William C., & Sandra A. Thompson (1988) Rhetorical structure theory: Towards a functional theory of text organization, *Text* **8:3**, 243-281.
- Masterman, Margaret (2005) *Language, Cohesion and Form*, edited by Yorick Wilks, Cambridge University Press, Cambridge.
- Miller, George A. (1995) WordNet: A lexical database for English, *Communications of the ACM* **38:11**, 39-41.
- Minsky, Marvin (1975) A framework for representing knowledge, in P. Winston, ed., *The Psychology of Computer Vision*, McGraw-Hill, New York, 211-280.
- Minsky, Marvin (1987) *The Society of Mind*, Simon & Schuster, New York.
- Minsky, Marvin (2006) *The Emotion Machine: Commonsense Thinking, Artificial Intelligence, and the Future of the Human Mind*, Simon & Schuster, New York.
- Mohanty, J. N. (1982) *Husserl and Frege*, Indiana University Press, Bloomington.
- Montague, Richard (1970) English as a formal language, reprinted in R. Montague, *Formal Philosophy*, Yale University Press, New Haven, pp. 188-221.
- Newell, Allen, & Herbert A. Simon (1972) *Human Problem Solving*, Prentice-Hall, Englewood Cliffs, NJ.
- Ockham, William of (1323) *Summa Logicae*, Johannes Higman, Paris, 1488 (the edition owned by C. S. Peirce).
- Ogden, C. K., & I. A. Richards (1923) *The Meaning of Meaning*, Harcourt, Brace, and World, New York, 8th edition 1946.
- 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 (1911) Assurance through reasoning, MS 670.
- Peirce, Charles Sanders (CP) *Collected Papers of C. S. Peirce*, ed. by C. Hartshorne, P. Weiss, & A. Burks, 8 vols., Harvard University Press, Cambridge, MA, 1931-1958.
- 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.
- Peppas, Pavlos (2008) Belief revision, in F. van Harmelen, V. Lifschitz, & B. Porter, *Handbook of Knowledge Representation*, Elsevier, Amsterdam, pp. 317-359.
- Post, Emil L. (1943) Formal reductions of the general combinatorial decision problem, *American J. of Mathematics* **65**, 197-268.
- Propp, Vladimir (1928) *Morfologia Skazki*, translated as *Morphology of the Folktale*, University of Texas Press, Austin, 1958.
- Reid, Thomas (1785) *Essays on the Intellectual Power of Man*, Edinburgh University Press, Edinburgh.
- Saussure, Ferdinand de (1916) *Cours de Linguistique Générale*, translated by W. Baskin as *Course in General Linguistics*, Philosophical Library, New York, 1959.

- Schank, Roger C., ed. (1975) *Conceptual Information Processing*, North-Holland Publishing Co., Amsterdam.
- Schank, Roger C., & Robert P. Abelson (1977) *Scripts, Plans, Goals and Understanding*,
- Selz, Otto (1913) *Über die Gesetze des geordneten Denkverlaufs*, Spemann, Stuttgart.
- Shanker, Stuart G. (1987) *Wittgenstein and the Turning Point in the Philosophy of Mathematics*, SUNY Press, Albany.
- Sowa, John F. (1976) Conceptual graphs for a data base interface, *IBM Journal of Research and Development* **20:4**, 336-357.
- Sowa, John F. (1984) *Conceptual Structures: Information Processing in Mind and Machine*, Addison-Wesley, Reading, MA.
- Sowa, John F. (2000) *Knowledge Representation: Logical, Philosophical, and Computational Foundations*, Brooks/Cole Publishing Co., Pacific Grove, CA.
- Sowa, John F. (2005) The Challenge of Knowledge Soup, in J. Ramadas & S. Chunawala, *Research Trends in Science, Technology, and Mathematics Education*, Homi Bhabha Centre, Mumbai, pp. 55-90.
- Sowa, John F. (2008) Conceptual graphs, in F. van Harmelen, V. Lifschitz, and B. Porter, eds., *Handbook of Knowledge Representation*, Elsevier, Amsterdam, pp. 213-237.
- Sperber, Dan, & Deirdre Wilson (1986) *Relevance: Communication and Cognition*, Blackwell, Oxford.
- Talmy, Leonard (2000) *Toward a Cognitive Semantics, Volume I: Concept Structuring Systems, Volume II: Typology and Process in Concept Structure*, MIT Press, Cambridge, MA.
- Tarski, Alfred (1933) Der Wahrheitsbegriff in den formalisierten Sprachen, English trans. as The concept of truth in formalized languages, in A. Tarski, *Logic, Semantics, Metamathematics*, Second edition, Hackett Publishing Co., Indianapolis, pp. 152-278.
- Tesnière, Lucien (1959) *Éléments de Syntaxe structurale*, 2nd edition, Librairie C. Klincksieck, Paris, 1965.
- Tolman, Edward C. (1948) Cognitive maps in rats and men, *Psychological Review* **55:4**, 189-208.
- van Deemter, Kees, & Stanley Peters (1996) *Semantic Ambiguity and Underspecification*, CSLI, Stanford, CA.
- Vygotsky, Lev Semenovich (1934) *Thought and Language*, MIT Press, Cambridge, MA, 1962.
- Waismann, Friedrich (1979) *Ludwig Wittgenstein and the Vienna Circle. Conversations Recorded by Friedrich Waismann*, Blackwell, Oxford.
- Wertheimer, Max (1925) *Über Gestalttheorie* translated as Gestalt Theory, by W. D. Ellis, *Source Book of Gestalt Psychology*, Harcourt, Brace and Co, New York 1938.
- Whitehead, Alfred North (1937) Analysis of Meaning, *Philosophical Review*, reprinted in A. N. Whitehead, *Essays in Science and Philosophy*, Philosophical Library, New York, 122-131.
- Wierzbicka, Anna (1996) *Semantics: Primes and Universals*, Oxford University Press, Oxford.
- Wilks, Yorick (2006) Thesauri and ontologies, in Rapp, Sedlmeier and Zunker-Rapp, eds., *Perspectives on Cognition: a Festschrift for Manfred Wettler*, Pabst.
- Wilks, Yorick (2008a) The Semantic Web as the apotheosis of annotation, but what are its semantics? *IEEE Transactions on Intelligent Systems*, May/June 2008.
- Wilks, Yorick (2008b) What would a Wittgensteinian computational linguistics be like? *Proceedings of AISB'08*, Workshop on Computers and Philosophy, Aberdeen.
- Winograd, Terry (1972) *Understanding Natural Language*, Academic Press, New York.
- Wittgenstein, Ludwig (1921) *Tractatus Logico-Philosophicus*, Routledge & Kegan Paul, London.
- Wittgenstein, Ludwig (1953) *Philosophical Investigations*, Basil Blackwell, Oxford.
- Wittgenstein, Ludwig (1964) *Philosophische Bemerkungen*, ed. by Rush Rhees, translated as *Philosophical Remarks* by Raymond Hargreaves and Roger White, University of Chicago Press, Chicago, 1980.
- Woods, William A. (1968) Procedural semantics for a question-answering machine, *AFIPS Conference Proc.*, 1968 FJCC, 457-471.
- Zadeh, Lotfi A. (1975) Fuzzy logic and approximate reasoning, *Synthese* **30**, 407-428.