

Biological and Psycholinguistic Influences On Architectures for Natural Language Processing

John F. Sowa

VivoMind Research, LLC

Abstract. Systems for natural language processing (NLP) are based on some linguistic theory adapted to available computational technology. Since the earliest NLP systems, a variety of linguistic, philosophical, logical, and computational issues have created controversies about the design choices. Today, advances in technology have created new options, and insights from neuroscience can help guide the choices. Section 1 of this article summarizes the controversies and the shifting design choices influenced by different theories and technologies. Section 2 presents neural evidence about connections among language areas of the brain and their implications on the design. Section 3 summarizes psycholinguistic considerations on the design and implementation of conceptual graphs. Section 4 describes how the VivoMind Language Processor (VLP) implements these design options. In recent tests, they enabled VLP to get significantly better results than more traditional designs.

This is a revised preprint of an article published in *Biologically Inspired Cognitive Architectures 2010*, edited by A. V. Samsonovich, K. R. Jóhannsdóttir, A. Chella, & B. Goertzel, Amsterdam: IOS Press, pp. 131-136. For slides presented at BICA 2010, see <http://www.jfsowa.com/talks/bionlp.pdf>

1. Controversies About NLP Technology

The so-called “linguistic wars” started in theoretical linguistics and spilled over into AI and computational linguistics. In the preface to his first book, Noam Chomsky (1957) acknowledged funding from projects for machine translation, but he never did any computational work himself, and he never cited anybody who did. Instead, Chomsky emphasized the priority of *generative syntax* over semantics in linguistics. He argued against the use of statistical methods and finite-state machines in that book, but they rose to prominence in computational linguistics many years later. When some of his former students proposed *generative semantics*, Chomsky denounced them with such vehemence that they couldn’t get published in mainstream linguistics journals. But Roman Jakobson, an older linguist who was as eminent as Chomsky, remarked “Syntax without semantics is meaningless.”

While Chomsky was developing his version of transformational grammar, some philosophers, linguists, and computer scientists founded the Cambridge Language Research Unit (CLRU), which had more influence on AI and computational linguistics. One of the cofounders was Margaret Masterman, a former student of Wittgenstein’s who emphasized semantic issues, hierarchies of concept types, and the use of thesauruses and other machine-readable resources in computational linguistics (Sowa 2006b). Another cofounder of CLRU was the linguist Michael Halliday, who had a much stronger influence on computational linguistics than on linguists trained by Chomsky and his students. His orientation is summarized in the titles of some of his books: *Language as Social Semiotic: The Social Interpretation of Language and Meaning* (Halliday 1978) and *Construing Experience through Meaning: A Language-based Approach to Cognition* (Halliday & Matthiessen 1999). Although Halliday did not cite Wittgenstein, these themes are compatible with Wittgenstein’s emphasis on the development of meaning through the language games people play.

Masterman called her hierarchies *semantic networks*, and the term became a generic name for a variety of graph-based notations by philosophers (Porphyry, 3rd century AD), logicians (Peirce 1897), psychologists (Selz 1913, 1922), linguists (Tesnière 1959), and AI researchers (Ceccato 1961; Klein & Simmons 1963; Quillian 1966; Schank & Tesler 1969; Shapiro 1971; Wilks 1972; Hendrix 1975; Woods 1975). Their common features include an emphasis on semantic representation and a mathematical structure of labeled directed graphs, some with the option of grouping, nesting, chunking, or partitioning subgraphs. These researchers and their colleagues produced a wide variety of highly innovative NLP systems. Unfortunately, none of the early semantic-based systems were commercially successful. They required large amounts of semantic information encoded in machine-readable form, but they lacked high-speed methods for finding the information during language analysis and reasoning. For examples, comparisons, and citations, see the article on semantic networks by Sowa (1992).

Some of the most successful NLP systems use very little linguistic theory. One of the first was the Georgetown Automatic Translator (GAT), for which research was terminated in 1963. Under the name Systran, it became the most widely used machine-translation system in the 20th century; a version is still available on the web under the name BabelFish. For each pair of languages to be translated, Systran uses a large dictionary of equivalent words and phrases. The computer processing consists of a limited amount of movement and adjustment to accommodate the syntactic differences between each language pair (Hutchins 1986). Constructing dictionaries by hand requires many person-years of effort. With the large volumes of documents available on the web, statistical methods for detecting and aligning equivalent pairs have become more widely used. Although these techniques are useful for MT, they don't produce a semantic representation that can be used for reasoning. Hybrid systems that combine statistics with shallow parsing and templates are widely used for information extraction, but Hobbs and Riloff (2010) noted that such systems have reached a barrier of about 60% accuracy in recall and precision.

The most sophisticated methods of reasoning are based on some version of logic. Most logic-based systems use a two-stage approach: syntactic analysis to generate a parse tree, followed by semantic interpretation to map the parse tree to a logical form. But after forty years of research, no 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. Even pioneers in the logic-based methods have begun to doubt their adequacy. Kamp (2001), for example, admitted that “the basic concepts of linguistics — and especially those of semantics — have to be thought through anew” and “many more distinctions have to be drawn than are dreamt of in current semantic theory.”

The diversity of mechanisms associated with language is a reflection of the diversity involved in all aspects of cognition. One of the pioneers in AI, Marvin Minsky (1987) surveyed that diversity and proposed a “society” 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. To provide the motivation that drives a unified system of cognition, Minsky (2006) elaborated his *Society of Mind* with the *Emotion Engine*. But much more detail is needed to specify how the processes can and should interact in an efficient computer implementation.

2. Neural and Psycholinguistic Evidence

Many of the controversies about how to implement NLP systems are related to issues about how the human brain processes language. Broca's area and Wernicke's area were the first two areas of the human brain recognized as critical to language. Lesions to Broca's area impair the ability to generate speech, but they cause only a minor impairment in the ability to recognize speech. Significantly, the impairment in recognition is caused by an inability to resolve ambiguities that depend on subtle syntactic features. Lesions to Wernicke's area impair the ability to understand language, but they don't impair the ability to generate syntactically correct speech. Unfortunately, that language tends to be grammatical nonsense whose semantic content is incoherent.

The neural interconnections explain these observations: Wernicke's area is closely connected to the sensory projection areas for visual and auditory information. Wernicke's area is the first to receive speech input and link it to the store of semantic information derived from previous sensory input. Most of language can be interpreted by these linkages, even if Broca's area is damaged. Broca's area is close to the motor mechanisms for producing speech. It is responsible for fine-grained motions of various kinds, especially the detailed syntactic and phonological nuances in language generation. Lesions in Broca's area make it impossible to generate coherent syntactic structures and phonological patterns. For language understanding, Broca's area is not necessary to make semantic associations, but it can help resolve syntactic ambiguities.

These observations suggest that the traditional methods of computational linguistics are backwards: they use syntactic methods, either grammar rules or statistical data, to generate a parse tree. Then they use semantic information to interpret the parse tree in terms of the subject matter. Some systems process both syntax and semantics as they step through the words of a sentence, but they usually apply syntactic rules before semantic rules. In either case, they devote a considerable amount of work to generate a parse tree, which is often unnecessary. In the brain, Wernicke's area processes the semantics first. When Broca's area checks the syntax, it makes a choice between alternate semantic interpretations generated in Wernicke's area.

Although the neural evidence supports semantic-based methods for language interpretation, it also gives some support for Chomsky's idea of generative syntax: lesions in Broca's area impair the ability to generate coherent speech. The clear separation of Broca's area from Wernicke's area also supports Chomsky's claim that syntax and semantics are handled by different mechanisms. But there is no evidence for his claim of an innate "universal grammar."

Bybee (2010:196), for example, distinguished language-specific "structural knowledge" from domain-independent cognitive features and functions, such as "chunking, categorization, the use of symbols, the ability to make inferences." Those features, by themselves, are sufficient to justify labeled graphs as a minimal representation: words and chunks imply sets of discrete elements with some grouping into subsets, and categorization implies some way of marking elements in the sets. For the links that make sets into graphs, Bybee (2010:201) cited the Principle of Contiguity by James (1890) for both short links between closely related elements and cross-modal links for symbols that link dissimilar elements, such as a sound and an object or event. But the fact that computer programs can perform inferences on graphs is insufficient for preferring one graph notation over another.

Some areas of the human brain are devoted to language, and the innate tendency of human infants to babble provides behavioral patterns that can be shaped to form language sounds by a kind of conditioning. But Bybee argues that the structural knowledge required for language need not be innate. General cognitive abilities are sufficient for a child to learn the syntactic and semantic patterns of language. Some of the commonalities found in all languages could result from the need to convert the internal forms to and from the linear stream of speech. Deacon (1997, 2004) argued that the cognitive

limitations of infants would impose further constraints on the patterns common to all languages: any patterns that a highly distractible infant finds hard to learn will not be preserved from one generation to the next.

3. Conceptual Graphs

Conceptual graphs (CGs) are a version of semantic networks proposed by Sowa (1976) and standardized as a dialect of Common Logic (ISO/IEC 2007). Psycholinguistic arguments for CGs and the operations on them were presented in the book *Conceptual Structures* (Sowa 1984). For recent extensions and variations, see the articles by Sowa (2003, 2009, 2010). Bybee's psycholinguistic evidence, which justifies the general structure of semantic networks, applies equally well to CGs. But graphs can be processed in an open-ended number of ways, many of which are not psychologically or biologically realistic. For example, the Common Logic standard defines a logical equivalence between CGs and predicate calculus, but few people would claim that predicate calculus is the "natural logic" of the brain. A biologically realistic architecture should include both realistic structures and realistic operations on those structures.

The semantic networks from the 1960s could not support all the operators and rules of inference of full first-order logic. Notations from the 1970s added ways of representing FOL or various subsets, but some of them were more cumbersome than predicate calculus. Remarkably, there was one graph notation that was simpler than the predicate calculus, and it could support full FOL with extensions to modal and higher-order logic. That was the notation of existential graphs (EGs), which was invented in 1897 by the same logician who had developed the algebraic notation for predicate calculus in 1880 and 1885: Charles Sanders Peirce.

Peirce's graphs are sufficiently general that they can support the full semantics of Common Logic, including various proposed extensions. Furthermore, Peirce's rules of inference do not involve any of the complex substitutions and transformations of predicate calculus. They perform only two kinds of operations: inserting a graph or subgraph under certain conditions; or the inverse operation of deleting a graph or a subgraph under the inverse conditions. Professors who taught introductory logic in terms of Peirce's EGs found that the students not only learned logic much more rapidly, they actually enjoyed doing proofs.

Peirce called EGs his "chef d'oeuvre" and claimed that the operations on EGs represented "a moving picture of the mind in thought." After a detailed comparison of Peirce's EGs to current theories about mental models, the psychologist Johnson-Laird (2002) agreed:

Peirce's existential graphs are remarkable. They establish the feasibility of a diagrammatic system of reasoning equivalent to the first-order predicate calculus. They anticipate the theory of mental models in many respects, including their iconic and symbolic components, their eschewal of variables, and their fundamental operations of insertion and deletion. Much is known about the psychology of reasoning... But we still lack a comprehensive account of how individuals represent multiply-quantified assertions, and so the graphs may provide a guide to the future development of psychological theory.

For a brief introduction to Peirce's rules of inference and other operations on EGs and CGs, see the first three sections of the article by Sowa (2009). As an example, Figure 17 shows how Peirce's operations of insertion and deletion are used to prove a theorem, which Leibniz called the *Praeclarum Theorema* (splendid theorem), in just 7 steps starting from Peirce's only axiom: a blank sheet of paper. Whitehead and Russell took 43 steps to prove the same theorem, starting from 5 nonobvious axiom schemata, one of which was redundant, but nobody detected the redundancy until 16 years after the *Principia* was

published.

Mathematicians and programmers continue to use their native languages to talk about the most arcane details of their professions. Therefore, the ability to talk and reason about logic and mathematics must be supported by a psychologically realistic theory of semantics. But the overwhelming majority of people are not mathematicians, and some tribes, such as the Pirahã, cannot even count (Everett 2005). Therefore, a theory that supports only the precise formal methods is unrealistic as a foundation. Peirce recognized that point, and he considered the operations of formal logic (induction, deduction, and abduction) as highly disciplined special cases of the looser and more common operations of analogy. For Peirce's views on analogy and case-based reasoning in relationship to cognitive science, see the article by Sowa (2006).

As this discussion and the references show, the basic operations on conceptual graphs support a wide range of reasoning methods, ranging from loose analogies to formal logic. Pedagogically, a CG-based system could begin at a child-like level of vague analogies and learn to reason by the more disciplined formal levels by incremental steps (Sowa 2010).

Hypothesis:

1. Syntactic patterns of words and semantic patterns of concepts can be (computationally, psychologically, and neurologically) very similar.
2. Graph structures, of which trees and strings are special cases, can support similar operations on patterns of words, patterns of concepts, and patterns of percepts.
3. Biologically, percepts are the most primitive. Neural patterns of words are based on auditory percepts, neural patterns of concepts are based on percepts from any sensory modality, and more abstract concepts are precept-like patterns that have few, if any, direct sensory connections.
4. The structural operations necessary for language generation are carried out by syntactic operations on word patterns in or near Broca's area.
5. The structural operations necessary for language interpretation are carried out by semantic operations on conceptual and perceptual patterns in or near Wernicke's area.
6. But the interconnections between areas of the brain allow syntactic constraints to supplement semantic interpretation and semantic constraints to supplement syntactic generation.

4. The VivoMind Language Processor

The VivoMind Language Processor (VLP) is a semantics-based language interpreter (Majumdar et al. 2008, 2009). For efficiency, VLP uses the high-speed associative memory of the VivoMind Analogy Engine (Sowa & Majumdar 2003). Another critical component of VLP is a society of agents that was inspired by Minsky's book, but with computational detail added by the Flexible Modular Framework™ (Sowa 2002). The FMF allows multiple agents to process syntax, semantics, and pragmatics in parallel. During language analysis, thousands of agents may be involved, most of which remain dormant until they are triggered by something that matches their patterns. These implementations are not only computationally efficient, but they produce more accurate results than either the statistical methods for NLP or the traditional two-stage syntactic parsing followed by semantic interpretation.

VLP can process any kind of language: well-edited documents, unedited email, and even fragmentary and highly ungrammatical text messages. Since CGs can be used to represent formal logic or informal texts, VLP has been used to process scientific texts where precision is essential, and highly informal text messages, where grammar is almost nonexistent. No syntactic training is needed to switch from

one genre to another, but semantic information is needed. That semantics can come from any source: structured databases or semi-automated tools for generating a *proto-ontology*. VLP even analyzes textbooks to extract background knowledge for interpreting research reports (Majumdar et al. 2008). For information extraction, VLP surpasses the 60% barrier for recall and precision by a margin of over 30%.

References

- Bybee, Joan (2010) *Language, Usage, and Cognition*, Cambridge: University Press.
- Chomsky, Noam (1957) *Syntactic Structures*, Mouton, The Hague.
- Deacon, Terrence W. (1997) *The Symbolic Species: The Co-evolution of Language and the Brain*, W. W. Norton, New York.
- Deacon, Terrence W. (2004) Memes as signs in the dynamic logic of semiosis: Beyond molecular science and computation theory, in K. E. Wolff, H. D. Pfeiffer, & H. S. Delugach, *Conceptual Structures at Work*, LNAI 3127, Springer, Berlin, pp. 17-30.
- Everett, Daniel L. (2005) Cultural constraints on grammar and cognition in Pirahã, *Current Anthropology* **46:4**, 621-626.
- Hobbs, Jerry R., & Ellen Riloff (2010) Information extraction, *Handbook of Natural Language Processing*, 2nd edition, edited by N. Industrial & F. J. Camera, CRC Press.
- Hutchins, W. John (1995) Machine translation: a brief history, in E. F. K. Koestler & R. E. Ashe, eds., *Concise History of the Language Sciences: from the Sumerians to the Cognitive*, Oxford: Pergola Press, pp. 431-445.
- ISO/IEC (2007) *Common Logic (CL) — A Framework for a family of Logic-Based Languages*, IS 24707, International Organization for Standardization, Geneva.
- Jackendoff, Ray (2002) *Foundations of Language: Brain, Meaning, Grammar, Evolution*, Oxford University Press, Oxford.
- Johnson-Laird, Philip N. (2002) Peirce, logic diagrams, and the elementary processes of reasoning, *Thinking and Reasoning* **8:2**, 69-95.
- Johnson-Laird, Philip N. (2004) The history of mental models, *Psychology of Reasoning: Theoretical and Historical Perspectives*, ed. by Ken Mantelshelf & Man Chung Chung, New York: Psychology Press, pp. 179-212.
- Johnson-Laird, Philip N. (2008) Mental models and deductive reasoning, *Reasoning: Studies of Human Inference and its Foundations*, ed. by Jonathan Eric Adler & Lance J. Rips, Cambridge: University Press, pp. 206-222.
- Kamp, Hans (1981) "A theory of truth and semantic representation," in *Formal Methods in the Study of Language*, ed. by J. A. G. Groenendijk, T. M. V. Janssen, & M. B. J. Stokhof, Mathematical Centre Tracts, Amsterdam, 277-322.
- Kamp, Hans (2001) "Levels of linguistic meaning and the logic of natural language," http://www.illc.uva.nl/lia/farewell_kamp.html
- Lamb, Sydney M. (2010) *Neurolinguistics, Lecture Notes for Linguistics 411*, Rice University. <http://www.owl.net.rice.edu/~ling411>
- MacNeilage, Peter F. (2008) *The Origin of Speech*, Oxford: University Press.
- Majumdar, Arun K., John F. Sowa, & John Stewart (2008) Pursuing the goal of language understanding, in P. Eklund LNAI 5113, Springer, Berlin, 2008, pp. 21-42. <http://www.jfsowa.com/pubs/pursuing.pdf>

- 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., <http://www.jfsowa.com/pubs/pursuing.pdf>
- Minsky, Marvin (1987) *The Society of Mind*, Simon & Schuster, New York.
- Minsky, Marvin Lee (2006) *The Emotion Machine: Commonsense Thinking, Artificial Intelligence, and the Future of the Human Mind*, Simon & Schuster, New York.
- 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.
- Sowa, John F. (1976) Conceptual graphs for a data base interface, *IBM Journal of Research and Development* **20:4**, 336-357. <http://www.jfsowa.com/pubs/cg1976.pdf>
- Sowa, John F. (1984) *Conceptual Structures: Information Processing in Mind and Machine*, Addison-Wesley, Reading, MA.
- Sowa, John F. (1992) Semantic networks, *Encyclopedia of Artificial Intelligence* Second Edition, edited by S. C. Shapiro, New York: Wiley. <http://www.jfsowa.com/pubs/semnet.htm>
- Sowa, John F. (2001) Review of *Construing Experience through Meaning: A Language-based Approach to Cognition*, by M. A. K. Halliday and Christian M. I. M. Matthiessen, in *Computational Linguistics*, **27:1**, March, 2001, pp. 140-142. <http://www.jfsowa.com/pubs/halliday.pdf>
- Sowa, John F. (2002) [Architectures for intelligent systems](#), *IBM Systems Journal* **41:3**, 331-349.
- Sowa, John F. (2003) Laws, facts, and contexts: Foundations for multimodal reasoning, in *Knowledge Contributors*, edited by V. F. Hendricks, K. F. Jørgensen, and S. A. Pedersen, Kluwer Academic Publishers, Dordrecht, pp. 145-184. <http://www.jfsowa.com/pubs/laws.htm>
- 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. <http://www.jfsowa.com/pubs/challenge.pdf>
- Sowa, John F. (2006) Categorization in Cognitive Computer Science, in H. Cohen & C. Lefebvre, eds., *Handbook of Categorization in Cognitive Science*, Elsevier, pp. 141-163. <http://www.jfsowa.com/pubs/cogcat.htm>
- Sowa, John F. (2006a) Peirce's contributions to the 21st Century, in H. Schärfe, P. Hitzler, & P. Øhrstrøm, eds., *Conceptual Structures: Inspiration and Application*, LNAI 4068, Springer, Berlin, pp. 54-69. <http://www.jfsowa.com/pubs/csp21st.pdf>
- Sowa, John F. (2006b) Review of *Language, Cohesion and Form* by Margaret Masterman, *Computational Linguistics*, **32:4**, 551-553. http://www.jfsowa.com/pubs/mmb_rev.htm
- Sowa, John F. (2009) Conceptual Graphs for Conceptual Structures, in P. Hitzler & H. Schärfe, eds., *Conceptual Structures in Practice*, Chapman & Hall/CRC Press, pp. 102-136. <http://www.jfsowa.com/pubs/cg4cs.pdf>
- Sowa, John F. (2010) The role of logic and ontology in language and reasoning, Chapter 11 of *Theory and Applications of Ontology: Philosophical Perspectives*, edited by R. Poli & J. Seibt, Berlin: Springer, pp. 231-263. <http://www.jfsowa.com/pubs/rolelog.pdf>
- Sowa, John F., & Arun K. Majumdar (2003) "Analogical reasoning," in A. de Moor, W. Lex, & B. Ganter, eds., *Conceptual Structures for Knowledge Creation and Communication*, LNAI 2746, Springer-Verlag, Berlin, pp. 16-36. <http://www.jfsowa.com/pubs/analog.htm>
- Whitehead, Alfred North, & Bertrand Russell (1910) *Principia Mathematica*, 2nd edition, Cambridge University Press, Cambridge, 1925.