Quantum Cognition

Arun K. Majumdar and John F. Sowa
VivoMind Research, LLC

Article. In *The Emperor's New Mind*, Roger Penrose claimed that quantum-mechanical effects are critical to human intelligence. But those effects need not be represented at the atomic level. A method of encoding conceptual graphs in a continuous representation called a *cognitive signature* supports a kind of *quantum knowledge representation* (QKR). It exhibits the key properties of superposition, entanglement, and uncertainty. The operations of searching and graph matching, when performed on a QKR, are analogous to the measurements in quantum-mechanical systems. A quantum computer would be ideal for processing them. But even with today's digital computers, searching and graph matching on a QKR can be performed with floating-point computations that scale in logarithmic time. For processing Big Data, they enable an ordinary laptop to outperform a supercomputer. This article shows how a QKR is used to analyze large volumes of documents and answer questions about them.

1. Holistic Perception and Path-Based Analysis

Mammalian perception evolved over 250 million years ago, but humans diverged from the apes about 6 million years ago. In that short time, human languages and the human brain evolved a new way thinking and communicating. Evidence from neuroscience shows the difference: fMRI scans indicate that both sides of the cerebral cortex are involved in language processing, but with different specializations. The left hemisphere (LH) is critical for the detailed syntactic processing, but the right hemisphere (RH) is more active in classification, recognition of metaphors, and distinguishing word senses. In computational terms, RH uses the more general perceptual mechanisms for recognizing the overall structure or Gestalt, while LH does the path-based analysis in parsing sentences, recognizing fine details, and mapping word patterns to and from perceptual patterns.

Penrose (1989, 2011) suggested that human intelligence results from quantum mechanical operations in the neural cells. But he admitted "In my view the conscious brain does not act according to classical physics. It doesn't even act according to conventional quantum mechanics. It acts according to a theory we don't yet have" (2009). For their theory of quantum cognition, Busemeyer and Bruza (2012) observed that human beliefs don't jump from one state to another. Instead, people feel an ambiguous superposition of all the options or "eigenstates." Although Peirce (1877) didn't use the terminology of quantum mechanics, he aptly characterized the uneasy feeling:

> Doubt is an uneasy and dissatisfied state from which we struggle to free ourselves and pass into the state of belief; while the latter is a calm and satisfactory state which we do not wish to avoid, or to change to a belief in anything else. On the contrary, we cling tenaciously, not merely to believing, but to believing just what we do believe.

As Penrose observed, a theory that predicts the probability of one option or another is useful. But a computational system for language understanding and reasoning requires a theory that can represent the details of those options. Some method of quantum knowledge representation (QKR) is necessary. Some authors propose a cognitive version of quantum mechanical wave functions: an infinite dimensional Hilbert space of complex functions whose eigenstates represent alternative beliefs (Bruza et al. 2009; Aerts & Sozzo 2011; Wittek et al. 2013). The "fixation of belief," as Peirce called it, would correspond to a measurement that resolves the superposition by forcing the QKR for a confused state of doubt into an eigenstate that has a definite answer. The challenge for those theories is to specify the QKR functions, map them to and from language and logic, and show how they can support perception, action, and reasoning.

2. Mapping Graphs to Continuous Functions

For over a century, logicians, philosophers, linguists, and computer scientists have been mapping language to various linear and graphic notations and specifying systematic rules for reasoning with them. Of all the versions, Peirce's existential graphs (EGs) are the simplest forms that have the expressive power of first-order logic. Some notations — including RDF, discourse representation structure (DRS), and conceptual graphs (CGs) — can be translated directly to EGs (Sowa 2013). Other notations, such as predicate calculus and description logics, can also be translated to EGs, but the mappings require more transformations.

A good notation is important, but the methods for reasoning with and about the notation are critical. Peirce's rules for reasoning with EGs are based on just two operations: insert a graph or erase a graph. The conditions for recognizing when to insert or erase a graph are visual: they can be seen or imagined by the same mechanisms as perception. The result is a simplification and generalization of Gentzen's version of natural deduction (Sowa 2011). Because of the simplicity of the operations and their relationship to perception, the psychologist Johnson-Laird (2003) recommended EGs as a candidate for a logic that could be supported by neural mechanisms.

René Thom was a mathematician who proposed mappings from language to graphs to continuous functions. The basis for the mapping is *catastrophe theory*, which he invented and applied to a variety of physical phenomena. As he broadened the range of applications, Thom (1969) discovered biological phenomena that displayed related patterns. For language, he used the dependency graphs by Lucien Tesnière (1959) to represent linguistic patterns and relate them to the patterns of catastrophe theory [14]. Petitot (2011) and Wildgen (1994, 2010) developed Thom's ideas further. Their mappings demonstrate methods for translating graphs to continuous functions, but their real-valued functions don't exhibit the superposition and entanglement of a QKR.

The graphs for representing molecules in organic chemistry are similar to EGs and CGs. The similarity is not a coincidence because Peirce had a BS degree in chemistry, and his goal for EGs was to represent "the atoms and molecules" of logic. He adopted the chemical term *valence* for the number of arguments of a relation. For graph computations, chemists have been in the forefront of developing efficient algorithms for finding and comparing large numbers of large graphs. For searching a database with millions of chemical graphs, Rhodes et al. (2007) mapped each graph to a numeric vector. To find all chemical graphs that are similar to a particular graph *g*,

- Represent each graph by its International Chemical Identifier (InChI).
- Map the InChI codes to numeric vectors that encode both the graph structure and the labels of the atoms and bonds.
- Estimate the semantic distance between graphs by a measure based on both the graph structure and the labels on the nodes and arcs (atoms and bonds).
- Index the vectors by a locality-sensitive hashing algorithm (LSH) that stores similar graphs in nearby locations.
- For any query graph *q* and semantic distance *d*, find all graphs that are within a distance *d* of the query graph *q*.

The time to encode a single graph as a numeric vector takes time that is polynomial in the size of the graph. But the time to find all graphs that are similar to a given graph increases logarithmically with the number of graphs in the database (DB).

The time to encode a single graph as a numeric vector varies as a polynomial in the size of the graph. But once the graphs are encoded, the time to find all graphs similar to a query graph increases logarithmically with the number of graphs in the DB.

The InChI-based algorithm is highly efficient for finding all graphs in the DB that are similar to the query graph *q*. But it cannot compare a query graph to the subgraphs of any graph. For the search algorithm to access all subgraphs of every graph in the DB, each one would have to be encoded and stored as a separate entry. That would cause an exponential increase in the number of entries. But a QKR encoding would not require the subgraphs to be stored separately. Each of them would be a "virtual graph," accessible as an eigenstate of a superposition.

3. Cognitive Memory

The real-valued mappings discussed in Section 2 can be generalized to complex numbers or to geometric algebras, which include complex numbers as a special case. The VivoMind software uses such methods to encode conceptual graphs in complex-valued vectors or the multivectors of geometric algebras. For any conceptual graph *g*, such an encoding is called the *cognitive signature*™ of *g*. Figure 1 shows *Cognitive Memory*™ as a system for storing the signatures of all CGs in a knowledge base. The lower level of Figure 1 shows how a query CG *q* is mapped to a cognitive signature and compared to a database of stored signatures. That DB would also store links back to the original CGs or to other resources from which the CGs were derived.

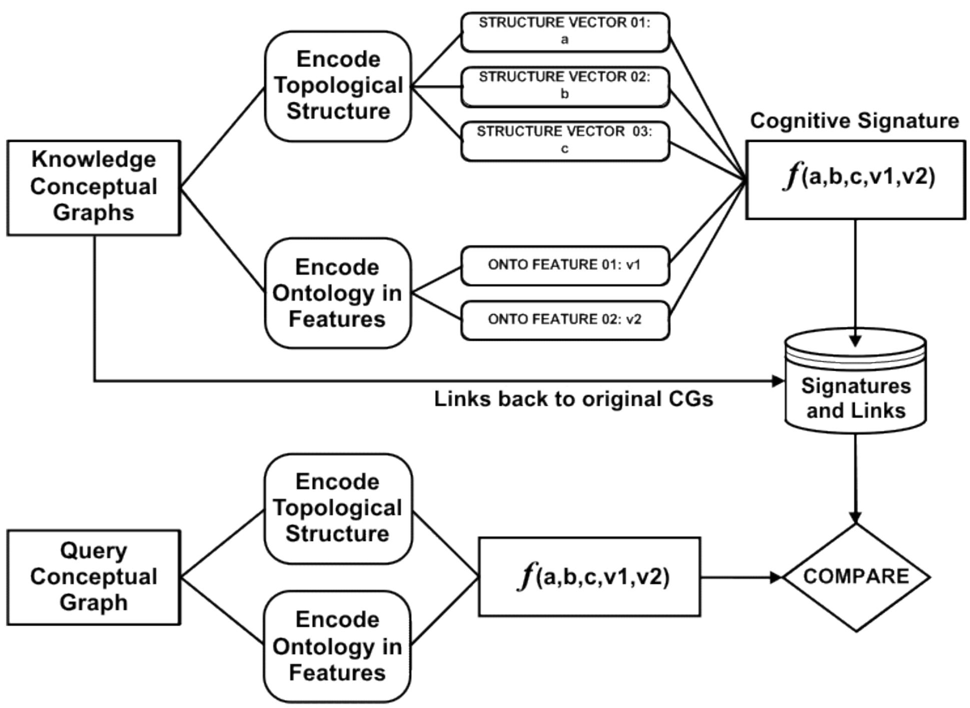

Figure 1. Cognitive Memory™ for conceptual graphs

The mapping shown in Figure 1 involves two kinds of vectors or multivectors: *topological vectors* that encode the structure of the graphs, and *ontological vectors* that encode the labels on the nodes and arcs. The topological vectors represent the connectivity, branching, and cycles of the graphs. The ontological vectors represent the types of concepts and relations. The type labels are organized in a hierarchy (partial ordering) that may be specified by the axioms and definitions of some ontology. The details of how the hierarchy was specified are not encoded in the cognitive signatures. An informal taxonomy such as WordNet or a partial ordering derived from a set of axioms would be encoded in the same way.

The function labeled *f* in Figure 1 combines the two vectors to form a cognitive signature. But *f* may be any member of a family of encodings that could emphasize different aspects of the structure or the ontology. A function that ignores the ontology would encode a structure of unlabeled graphs. A function that ignores the structure would encode a bag of concepts and relations. For most applications, a combination of the two is best, but *f* may give greater weight to certain aspects of the structure or the ontology. During the development of VivoMind technology, many algorithms for representing *f* were implemented and tested. In the best versions, the encoding of a graph *g* is a superposition of the encodings for each of its subgraphs. As in quantum mechanics, the subgraphs are eigenstates that are virtually present in the encoding. They can be detected by a measurement (i.e., a search with a query graph), but they do not occupy space in any physical medium. As a result, a QKR encoding has the effect of making all subgraphs accessible to a search, but without increasing the space requirements.

The time to encode a graph scales as a polynomial in the size of the graph. The time to find matching graphs scales logarithmically with the number of graphs in the database. For encoding a single graph, the exponent of the polynomial depends on the complexity of the interconnections. In encoding a document, for example, the connections within a sentence may be complex, but the references that link sentences within a document or collection of documents are much sparser. The observed encoding time for the CGs derived from a document or a corpus of documents is proportional to (N log N), where N is the total number of sentences.

In searching cognitive signatures, a distance measure is necessary to find all graphs within a given semantic distance. But distances in the ontology depend on the way levels are counted. In a biological classification, for example, the number of levels that separate molluscs from vertebrates is less than the number that separate cats from dogs. In general, distances at the lower levels should be considered much shorter than distances at the upper levels. But context is also important: a guppy, for example, is not a typical fish nor a typical pet, but it is a typical pet fish. When considered as food, salmon would be closer to a clam than to a guppy. To compute a semantic distance that takes the levels and the context into account, Majumdar (2013) specified a method of *conceptual relativity* that fixes the distance from the top of any ontology to the bottom at 1.0. Distances decrease at the lower levels, and they are affected by qualifiers such as pet or food. But the distances are not affected by added levels when scientists and engineers refine their classifications.

By contrast, path-based methods for finding similar graphs take N3 time for a database of N graphs [19]. That performance is acceptable for processing a small number of graphs retrieved by faster methods, but not for an associative memory that finds all similar graphs in a large database. Retrieval in logarithmic time is necessary to support the almost instantaneous processes in human perception and reasoning.

To find all similar graphs, some measure is necessary to evaluate the semantic distance between any two graphs. But that measure depends on the distance assigned to each level of the ontology. In a biological hierarchy, for example, more levels separate cats from dogs than vertebrates from molluscs. To account for the greater number of instances, a larger distance should be assigned to each level near the top of an ontology than to the levels near the bottom. But context is also important. A guppy, for example, is not a typical fish nor a typical pet, but it is a typical pet fish. When considered as food, a salmon would be closer to a clam than to a guppy.

To compute a semantic distance that takes the levels and the context into account, Majumdar [20] specified a method of conceptual relativity that fixes the distance from the top of any ontology to the bottom at 1.0. Distances decrease at the lower levels, and they are affected by qualifiers such as pet or food. But the distances are not affected by new levels that scientists add when they refine a classification.

4. Applications

A high-speed system that can retrieve all patterns similar to a given query is essential for perception, analogy, and all methods of reasoning, formal or informal. Figure 2 compares the use of analogy for *case based reasoning* with a formal system based on induction followed by deduction (Sowa 2006). The new case in Figure 2 corresponds to a query graph. It presents a novel pattern, which may be a threat, an opportunity, or some irrelevant or neutral occurrence. A computer system might have a stored theory that had been derived by induction from cases 1 to 4. By deduction, it could apply the theory to derive a conclusion that characterizes the new case and suggests some way of dealing with it. But given the same four cases, a system that uses analogy for case-based reasoning could derive the same conclusion in just one step. A theory derived from thousands of cases, however, would be more convenient and efficient than reasoning by analogy for every new application.

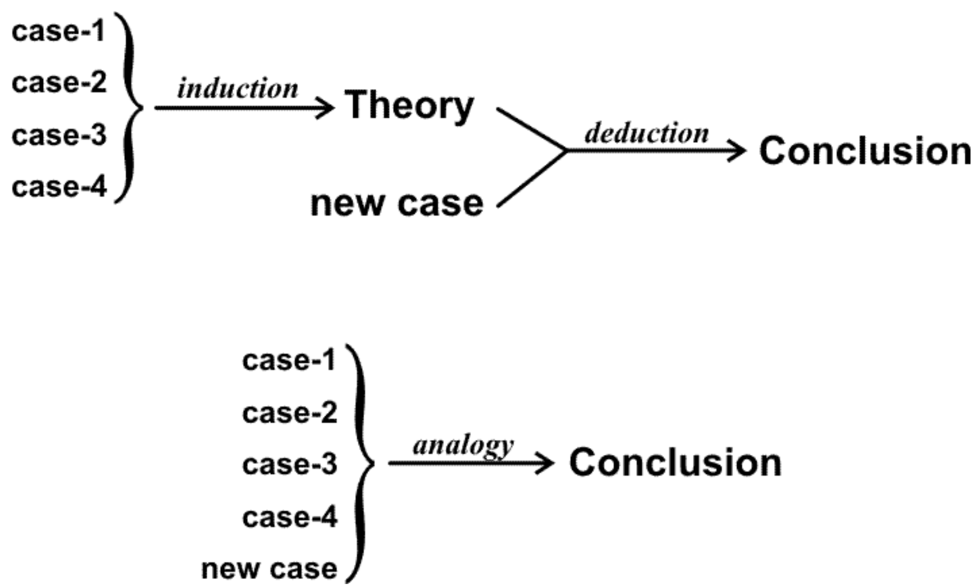

Figure 2. A comparison of logical and analogical reasoning

By analogical reasoning, a novel pattern triggers a search for similar patterns in long-term memory. That search brings back memories of previous experiences and the favorable or unfavorable results of actions taken or not taken. Peirce observed that the three methods of formal reasoning — induction, deduction, and abduction — are refinements of the informal methods of analogy and case-based reasoning. But the formal methods can also benefit from Cognitive Memory or related QKR systems for high-speed retrieval.

For examples of practical applications of Cognitive Memory, see the slide presentations on The Goal of Language Understanding [22, 23]. Slides 19 to 25 of [22] show an example of case-based reasoning for checking student answers to word problems in algebra. Slides 7 to 15 of [23] discuss an application to legacy re-engineering. It involved mapping two kinds of languages to conceptual graphs and encoding them in Cognitive Memory: software written in COBOL, SQL, and JCL; and documents about the software written in English. The task was to develop a cross-reference index that related passages in the

English documentation to the software modules they described. The second part of the task was to find inconsistencies between what the software did and what the documents described. Slides 16 to 24 of [23] discuss an application that derived CGs from a textbook on geology, encoded them in Cognitive Memory, and used them to interpret research reports about oil and gas fields. Then a geologist could describe a different geological site and ask the system to find related information from the reports.

As these examples show, CGs and the cognitive signatures that encode them can be derived from informal texts in natural languages or from formally specified notations. For the interface to other systems, Cognitive Memory uses the Conceptual Graph Interchange Format (CGIF). It is general enough to be translated to and from graph notations, frame notations, and any logic specified by or translated to the ISO standard for Common Logic [24]. For the legacy re-engineering task, CGs derived from programming languages provided the ontology and background knowledge used to interpret the English documentation. That is a much simpler task than translating English to any formal language.

References

Aerts, Diederik, & Sandro Sozzo (2011) *Quantum Structure in Cognition: Why and How Concepts Are Entangled*, in D. Song et al., *Lecture Notes in Computer Science* **7052**, Berlin: Springer, pp. 116-127.

Aaronson, Scott (2013) *Quantum Computing since Democritus*, Cambridge: Cambridge University Press.

Busemeyer, Jerome R., & Peter D. Bruza (2012) *Quantum Models of Cognition and Decision*, Cambridge: Cambridge University Press.

Bruza, Peter D., Dominic Widdows, & John Woods (2009) A quantum logic of down below, in K. Engesser, D. Gabbay, & D. Lehmann, eds., *Handbook of Quantum Logic and Quantum Structures*, Elsevier, pp. 625-660.

Gärdenfors, Peter (2000) *Conceptual Spaces: The Geometry of Thought*, Cambridge, MA: MIT Press.

Johnson-Laird, Philip N. (2002) Peirce, logic diagrams, and the elementary processes of reasoning, *Thinking and Reasoning* **8:2**, 69-95. http://mentalmodels.princeton.edu/papers/2002peirce.pdf

Majumdar, Arun K. (2013) Relativistic concept measuring system for data clustering, United States Patent, 8566321.

Majumdar, Arun K., John F. Sowa, & John Stewart (2008) Pursuing the goal of language understanding, in P. Eklund & O. Haemmerlé, eds, *Proceedings of ICCS'08*, LNAI 5113, Berlin: Springer, pp. 21-42.

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. http://www.jfsowa.com/pubs/paradigm.pdf

Majumdar, Arun K., & John F. Sowa (2013) Big data: structured and unstructured, in J. Liebowitz, ed., *Big Data and Business Analytics*, Boca Raton: CRC Press, pp. 249-268.

Peirce, Charles Sanders (1877) The fixation of belief, *Popular Science Monthly* **12**, Nov 1877, 1-15. http://www.peirce.org/writings/p107.html

Petitot, Jean (2011) *Cognitive Morphodynamics: Dynamical Morphological Models of Constituency in Perception and Syntax*, Bern: Peter Lang.

Penrose, Roger (1989) *The Emperor's New Mind*, Oxford University Press, Oxford.

Penrose, Roger (2009) Interview with S. Kruglinski and O. Chanarin, *Discover Magazine*, September 2009. http://discovermagazine.com/2009/sep/06-discover-interview-roger-penrose-says-physics-is-wrong-string-theory-quantum-mechanics

Penrose, Roger, & Stuart Hameroff (2011) Consciousness in the universe: neuroscience, quantum space-time geometry and Orch OR theory, *Journal of Cosmology* **14**, http://journalofcosmology.com/Consciousness160.html

Rhodes, James, Stephen Boyer, Jeffrey Kreulen, Ying Chen, & Patricia Ordonez (2007) Mining patents using molecular similarity search, *Pacific Symposium on Biocomputing* **12**, 304-315. http://psb.stanford.edu/psb-online/proceedings/psb07/rhodes.pdf

Sowa, John F. (2008) Conceptual graphs, in F. van Harmelen, V. Lifschitz, and B. Porter, eds., *Handbook of Knowledge Representation*, Amsterdam: Elsevier, pp. 213-237. http://www.jfsowa.com/cg/cg_hbook.pdf

Sowa, John F. (2011) Peirce's tutorial on existential graphs, *Semiotica* **186:1-4**, 345-394. http://www.jfsowa.com/pubs/egtut.pdf

Sowa, John F. (2013) From existential graphs to conceptual graphs, *International Journal of Conceptual Structures* **1:1**, 39-72. http://www.jfsowa.com/pubs/eg2cg.pdf

Tesnière, Lucien (1959) *Éléments de Syntaxe structurale,* 2nd edition, Paris: Librairie C. Klincksieck, 1965.

Thom, René (1969) Topological models in biology, *Topology* **8**, 313–335.

Thom, René (1988) *Esquisse d'une Sémiophysique*, Paris: InterEditions.

Way, Eileen C. (1991) *Knowledge Representation and Metaphor*, Dordrecht: Kluwer.

Wildgen, Wolfgang (1994) *Process, Image, and Meaning: A Realistic Model of the Meaning of Sentences and Narrative Texts*, John Benjamins Publishing Co., Amsterdam.

Wildgen, Wolfgang, & Per Aage Brandt (2010) *Semiosis and Catastrophes: René Thom's Semiotic Heritage*, Bern: International Academic Publishers.

Wittek, Peter, Bevan Koopman, Guido Zuccon, & Sándor Darányi (2013) Combining word semantics within complex Hilbert space for information retrieval, *Lecture Notes in Computer Science: Quantum Interaction* **8369**, Berlin: Springer, pp. 160-171.

[5] Peirce, Charles Sanders (1877) The fixation of belief, Popular Science Monthly 12, Nov 1877, 1-15.

[7] Aerts, Diederik, & Sandro Sozzo (2011) Quantum structure in cognition: Why and how concepts are entangled, in D. Song et al., Lecture Notes in Computer Science 7052, Berlin: Springer, pp. 116-127.

[8] K Wittek, Peter, Bevan Koopman, Guido Zuccon, & Sándor Darányi (2013) Combining word semantics within complex Hilbert space for information retrieval, Lecture Notes in Computer Science: Quantum Interaction 8369, Berlin: Springer, pp. 160-171.

[9] Sowa, John F. (2013) From existential graphs to conceptual graphs, International Journal of Conceptual Structures 1:1, 39-72. http://www.jfsowa.com/pubs/eg2cg.pdf .

[10] Sowa, John F. (2011) Peirce's tutorial on existential graphs, Semiotica 186:1-4, 345-394. http://www.jfsowa.com/pubs/egtut.pdf

[11] Johnson-Laird, Philip N. (2002) Peirce, logic diagrams, and the elementary processes of reasoning, Thinking and Reasoning 8:2, 69-95. http://mentalmodels.princeton.edu/papers/2002peirce.pdf

[12] Thom, René (1969) Topological models in biology, Topology 8, 313–335.

[13] Tesnière, Lucien (1959) Éléments de Syntaxe structurale, 2nd edition, Paris: Librairie C. Klincksieck, 1965.

[14] Thom, René (1988) Esquisse d'une Sémiophysique, Paris: InterEditions.

[15] Petitot, Jean (2011) Cognitive Morphodynamics: Dynamical Morphological Models of Constituency in Perception and Syntax, Bern: Peter Lang.

[16] Wildgen, Wolfgang (1994) Process, Image, and Meaning: A Realistic Model of the Meaning of Sentences and Narrative Texts, Amsterdam: John Benjamins Publishing Co.

[17] Wildgen, Wolfgang, & Per Aage Brandt (2010) Semiosis and Catastrophes: René Thom's Semiotic Heritage, Bern: International Academic Publishers.

[18] Rhodes, James, Stephen Boyer, Jeffrey Kreulen, Ying Chen, & Patricia Ordonez (2007) Mining patents using molecular similarity search, Pacific Symposium on Biocomputing 12, 304. http://psb.stanford.edu/psb-online/proceedings/psb07/rhodes.pdf

[19] Falkenhainer, Brian, Kenneth D. Forbus, Dedre Gentner (1989) The structure mapping engine: algorithm and examples, Artificial Intelligence 41, 1-63.

[20] Majumdar, Arun K. (2013) Relativistic concept measuring system for data clustering, United States Patent, 8566321.

[21] Sowa, John F. (2006) Categorization in cognitive computer science, in H. Cohen & C. Lefebvre, eds., Handbook of Categorization in Cognitive Science, Elsevier, 2006, p. 141-163.

[22] Sowa, John F. (2013) Slides on the goal of language understanding, Chapter 6, Analogy and case-based reasoning, http://www.jfsowa.com/talks/goal6.pdf

[23] Sowa, John F. (2014) Slides on the goal of language understanding, Chapter 7, Learning by Reading, http://www.jfsowa.com/talks/goal7.pdf