Taking Knowledge Management To the Next Level: Directions for Future Knowledge Systems Research

Arthur J. Murray, D.SC.

Introduction

While approaches to data and information management are primarily computational, knowledge management requires both computational and non-computational approaches. This paper discusses areas of research needed to gain a deeper understanding of how to integrate the two approaches.

The Differences Between Computational and Non-Computational Approaches to KM

Data and information management benefited well from the evolution in computational capability and the corresponding increases in storage volume and processing speed. Without computers we would not have available such a vast amount of data and information at our fingertips. Knowledge management, however, requires more than just computationally-intensive solutions. Whether one takes Polyani's individualistic viewpoint of knowledge, [1] [2] or Popper's socio-cultural perspective, [3] [4] the knowledge systems of tomorrow will require approaches that effectively incorporate algorithmic, cognitive, biological, neuropsychological and socio-cultural paradigms.

Knowledge management has firmly taken the lead as the buzzword of the moment. Sadly, most of the KM landscape looks like a repackaging of information technology. We do not want to discard all that we have accomplished with IT. Brute force computation is the primary engine that brings the world together "at the speed of light." But even the most powerful supercomputer does not approach the human brain in terms of its ability to process and store sensory data.

Close your eyes for a moment, and think about your second grade classroom. Mentally compose a single phrase that describes that particular setting. Now think of another phrase that describes that same setting but in a way completely different from the first phrase. How long did that exercise take? For most individuals, many experiences have transpired between the second grade and the present. Yet within seconds we can reach back and retrieve mental images from that distant period.

If we are truly serious about managing knowledge, why aren't we looking at the world's most capable technology? Can we make use of both natural and artificial approaches? Can the human mind and physiology synergistically interact with high-speed processors, mass storage and wide bandwidth communications? Can carbon and silicon work together side by side?

The Digital Library and Human Memory

When reviewing the major events that helped shape the present age, the inventions of the printing press and digital computer rank near the top. Each resulted in a quantum leap in the availability and volume of knowledge created and disseminated for use by mankind. The challenge has been, and still is today, to rapidly find the exact knowledge artifact needed for a given situation. As the volume of available knowledge increases, so does the difficulty of locating the correct artifact (see Figure 1, solid line). Because of how we have structured our information- document-centric world, we are constantly engaged in this tug of war. The more voluminous our libraries get, the harder it becomes to recall anything with a high degree of precision.

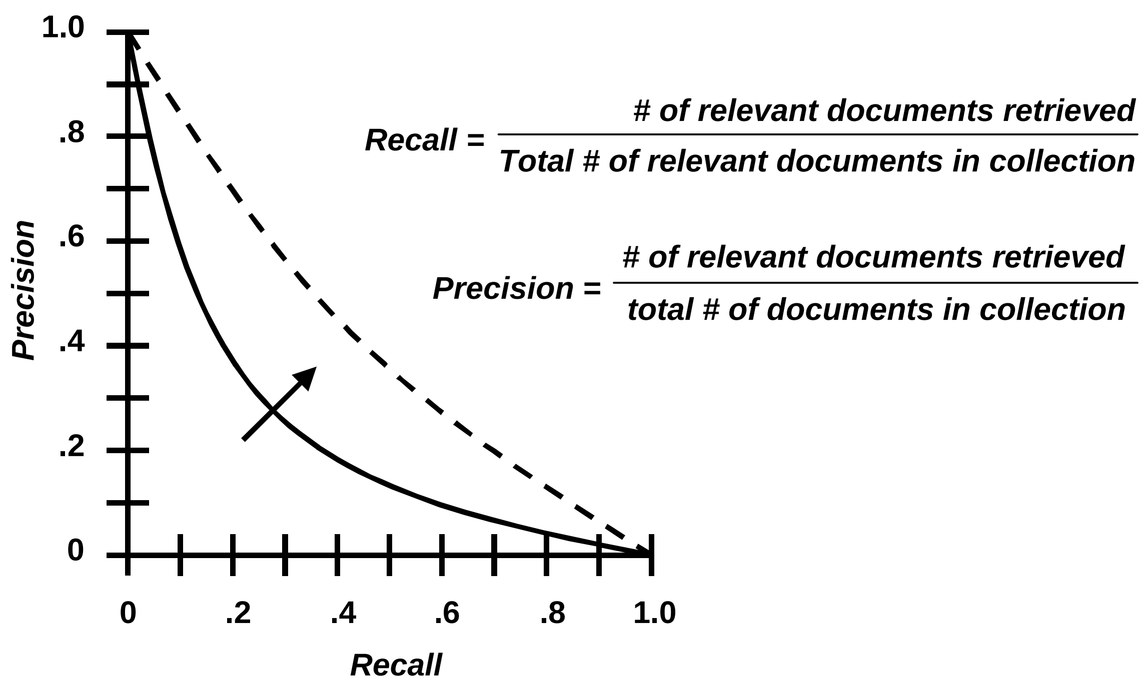

Figure One -- An Illustration of the Precision vs. Recall Problem

But what about recalling your 2nd grade classroom? Didn't you get a vivid picture of those wonderful days? Maybe you remembered the smell from a nearby bakery, or the stifling heat of a cramped room with no air conditioning, or squinting to see a dimly lit blackboard. It is likely that your recollection of the second grade didn't arrive cluttered with hundreds of other images --- fourth grade, fifth grade, college grades, the grade of the driveway in front of your second house. Remarkably, your natural retrieval system is able to achieve rapid recall with a much higher degree of precision than is possible computationally.

Two Reasons Why Computational Approaches Alone Are Insufficient

Recent advances in computational linguistics and the construction of large semantic networks (such as is found in Oracle's ConText Option) have begun to bend the precision recall graph outward (see dashed line in Figure 1). Temporal reasoning also attempts to improve precision and recall by incorporating the effects of the passage of time on the relevance to the current problem. However, most current information and knowledge processing systems can perform only within stable process compartments. [5] A process compartment is a domain-specific epoch in which a particular set of observable patterns emerges. Simple examples include the real estate boom of the late seventies and early eighties, or the stock market boom of the nineties. In each period, investors became conditioned to expect that the values of their investments only go in one direction—up. The error of getting caught in a squeeze at the top of the cycle is due in part to missing or ignoring time-invariant and cross-scale phenomena, whose relationship to the boom can be conjectured intuitively but not modeled algorithmically until the epoch finally resolves.

- *Reason #1: Stored Knowledge Becomes Invalid When the Associated Process Compartment Destabilizes*

The world is dynamic, cycling through periods of stability and transition (See Figure Two). Many different cycles with varying periodicities operate in superposition. Periods of stability may be short-lived or of epoch duration. Periods of transition may be smooth or extremely violent. When a compartment is stable, we make use of direct observation to enhance the world's knowledge corpora, especially given the ability of digital technology to rapidly collect and analyze massive amounts of data. For example, in the domain of unmanned space flight, supercomputers can perform the complex calculations needed to send probes to distant planets with remarkable accuracy, taking into account the rotation of the earth, the gravitational force of distant bodies, and the thrust of manmade propulsion systems. Even the adverse effects of the solar wind can be taken into account. This works because the solar system is a relatively stable, compartmentalized environment. The space and time scales are observable and predictable.

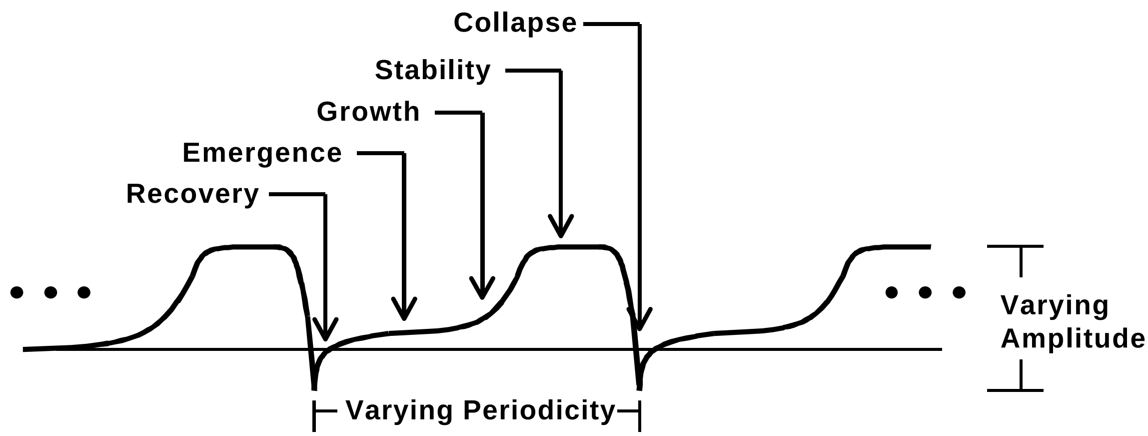

Figure Two -- Illustration of the Process Compartment Life Cycle

Such compartmentalization is not characteristic of social systems. The twentieth century was characterized by periods of stability interspersed among periods of violent transition. The great industrial age that began before the turn of the century was shaken apart by World War I and by the 1929 stock market crash and subsequent economic depression. World War II jolted heavy industry back into production. Then the cold war emerged, bringing with it a protracted period of stasis. Categorization was easy --- many nations were allied with either communist dictatorships or capitalist democracies. The remaining nations were neutral. Then the Berlin Wall and the Soviet Union collapsed, and the present era of globalization began.

It is easy to be seduced into thinking that knowledge management will help make the next century better by delivering the right knowledge to the right person at the right time. But social systems exhibit significantly different behaviors than physical systems. Knowledge management approaches that are heavily dependent on information technology and computational algorithms are simply not equipped to tackle problems involving the unpredictability of the human species.

On the other hand, human minds, working both individually and together, have made incredible discoveries, produced remarkable inventions and created many pockets of wealth and prosperity. This has been accompanied by an accelerated growth in technology, and in this new century we will soon have computers that are equal to the human brain in terms of nodes and connections. But many of the world's key events, good and bad, arose from actions that were non-computational. Because the computational approach depends upon limited observations of epiphenomena, the underlying cross-currents and interdependencies that give rise to the often violent transitional phases of society remain un-captured.

- *Reason #2: True Knowledge Flows Exhibit Quantum Behaviors Rarely Taken Into Account By Artificial Systems*

The data-information-knowledge-wisdom pyramid is a familiar illustration in knowledge management circles (see Figure Three). This creates the erroneous notion that each layer builds upon the other, when in fact each level can exist independently. Of greater interest to KM research are the points in the pyramid in which gap-crossing occurs, which may be between levels and within the same level (see Figure Four).

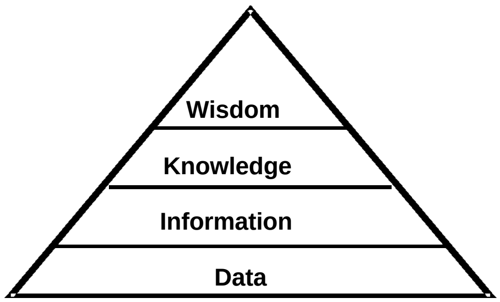

Figure Three -- Common Representation of the Transition from Data to Wisdom

Take for example, the generation and use of both physical and semiotic data in control systems theory. Physical data for a power plant may include voltages, frequencies and other physical measurements derived with physical (artificial) instruments. Other characteristics of the plant, such as the spatial arrangement of the equipment or the tactile response of the manual controls, are best derived from human observation and recorded in semiotic data formats such as natural language or engineering drawings. Clearly there is a boundary separating these two approaches to measurement. Humans cannot physically experience the precise voltage levels and waveforms with anything close to the accuracy of physical instruments. Computers cannot make intuitive observations regarding the tactile response or placement of a switch. Yet we need both forms of measurement. If we are to manage knowledge about a power plant in its totality, we need to integrate both sides of the gap into a unified system.

Other gap-crossings can be explored. Information presented in a display (information level, computational) is interpreted by an operator (knowledge level, natural). An operator's repetitive responses to certain circumstances (knowledge level, natural) are encoded, tested, and integrated into an automated system to improve productivity (knowledge level, computational).

Figure Four reflects the fact that the epistemic gaps encountered when crossing between natural and artificial systems are similar to stratification found in physics. Cross-gap phenomena have been observed in studies of holonomic brain processes, with impulses appearing, disappearing, and reappearing at different points in space without any visible propagation medium. [6] Figure Four shows, at least intuitively, that knowledge management necessitates moving back and forth between the two realms, Yet most of the effort in commercial knowledge management systems is focused primarily on the computational side.

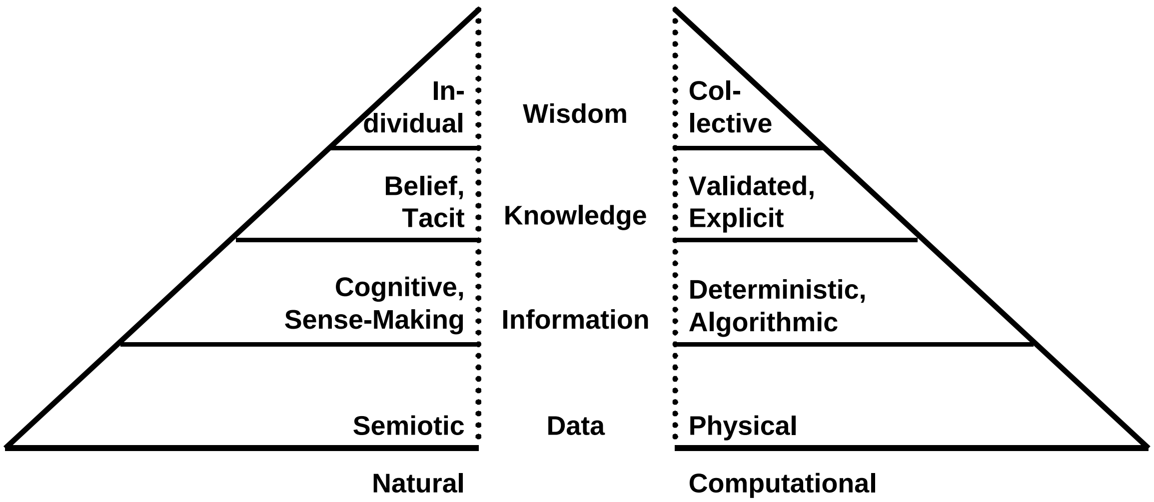

Figure Four -- The Natural Flow of Data, Information, Knowledge and Wisdom

Future Research Directions

Extensive research is needed to take knowledge management to the next level. This exploration must go beyond the current industry focus in which the majority of effort is spent getting the right information to the right place at the right time. If done correctly, knowledge management requires getting the right *knowledge* to the right *person* at the right time, and *using* that knowledge in the right manner.

The KM research landscape is extremely broad-based and multi-disciplinary. This paper will focus on two areas that particularly address the shortfalls discussed in the last section:

1) Improved knowledge representation and inference mechanisms
2) Advanced visualization techniques

Improved Knowledge Representation and Inference: Applied Semiotics

Many of today's KM toolsets spring from the document management realm. This is not surprising, given that most of the work performed in both the white and blue collar sectors is document-intensive. Whether in the form of policies and procedures, floor plans, source code or sacred scripture, documents are an integral part of our daily lives.

Until recently, most of the tools used to manage our vast document repositories have been built around identifying and indexing *what* is contained in the documents. Very little attention has been given to what the documents are *about*. This is an outgrowth of the computational paradigm - computers process symbols, so we turn our search engines loose at high speeds trying to match combinations of symbols. But our machines have no idea about what they are searching for. If we are to truly manage knowledge, our systems must have the ability to comprehend both the semantic and pragmatic meaning behind queries such as *"what do I need to know in order to accomplish X?"*

Most organizational knowledge is either tacit or implicit. If we want to capture deeply embedded knowledge, an understanding of mental models and cognitive processes is essential. In order to express this knowledge linguistically as an artifact, we need to construct rich domain ontologies that capture the appropriate context as well as the meaning of concepts. This has been the focus of much of the recent literature in the field of applied *semiotics.* [7] [8] Still, the smooth migration from mental models to knowledge artifacts is one of the major challenges of KM research.

Given recent advances in conceptual graphs, ontologies and neural networks, we can begin to build computational models that more closely correspond to the mental models that we began forming back in the second grade. In this way we can model the knowledge embedded within a document (or an entire system) by defining and modeling the concepts and relationships both within and beyond a local domain.

In applied semiotics, rich domain ontologies take the form of a concept space, as illustrated in Figure Five). A concept space represents the knowledge (concepts and relationships) contained (either explicitly or implicitly) in a document corpus, as it relates to a particular domain. The concepts and relationships within the space can be linked electronically to the portions of the documents that define or elaborate through explanation, story illustrations, or even raw data.

Care must be taken in the construction of a concept space. Combinatorics quickly degrade performance and utility, making modeling of the knowledge extremely difficult. For this reason, the concept space shown in the figure is layered. The lower layers contain fewer elements, implying that the more we understand about deep structure, the easier it will be for us to navigate the elements on the surface (i.e., the instances within our observable universe).

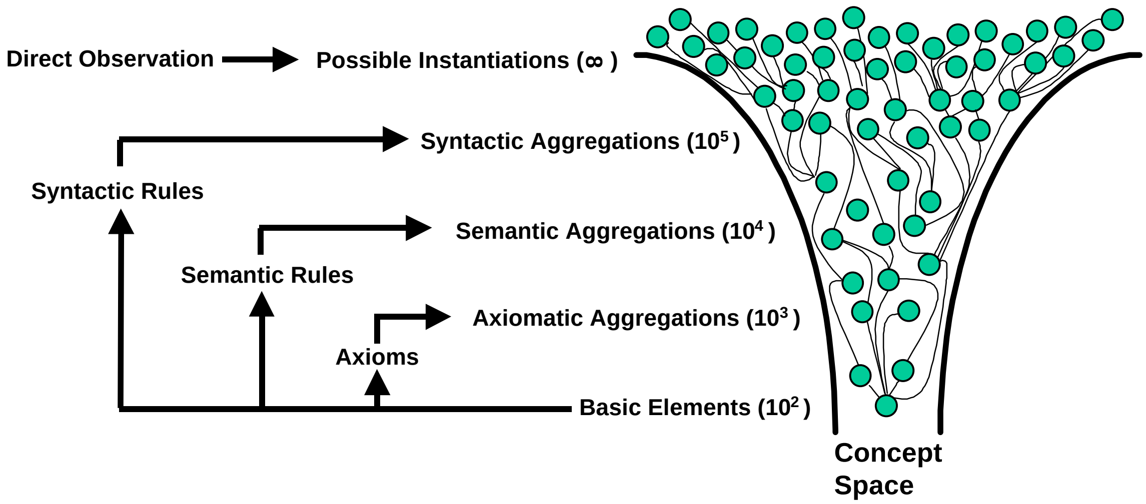

Figure Five -- Layers of the Concept Space

True knowledge management requires an understanding of the time-invariant structures beneath the surface level that give rise to the observed instances at the surface level. The substructural concepts and relationships can be encoded directly as conceptual graphs or as rules. We represent these substructures using a formal semiotic model developed by Pospelov. [9] The generic form of this model is:

M = < T, P, A, R >, where

M = a formal system
T = the set of basic elements that comprise M
P = the set of rules for constructing syntactically correct aggregations of T
A = all known axioms governing the attributes and behaviors of M
R = the set of rules for constructing semantically correct aggregations of T.

The subsurface layers in Figure Five are based on this model. The bottom layer consists of a set of basic elements or "atoms" from which more complex representations are built. In all, we expect that for any given domain there are no more than a few hundred of these basic elements. The building blocks used to construct T are derived from semantic primes and foundational conceptual relationships identified by Pospelov and linguists such as Wierzbicka. [10] Semantic primes are the most fundamental components that comprise the concepts we use to describe our world. These include such notions as time, space, causality, dependency and opposition.

The next layer represents A, the axioms that define all known valid aggregations of T. This is an open set, expandable up to and including all the elements of T. Just as semantic primes are used to construct the basic elements of T, ultra-structure theory [11] postulates a means for framing the fundamental time invariant laws that govern how those elements can be arranged or combined. Through the application of ultra-structure theory, a knowledge base consisting of many thousands of if-then relationships can be reduced to less than two-dozen classes. [12]

Ultra-structure theory uses an open logic approach to maintaining heuristics, in which the underlying rule structure is stable, and the outer levels are populated by an evolving rule set that gives rise to the diversity observed at the surface. "Open" logic in this sense means the use of a variety of interchangeable, semiotic logics in which the rules of derivation: 1) apply properties based on the subjective characteristics of human perception (such as time, space, causality); 2) can change along with changes in perception.

We expect that there are no more than a few thousand axioms inherent in any domain. Similarly, we believe that there are on the order of 10^4 semantically correct aggregations of basic elements, of which another order of magnitude more are syntactically correct.

The top layer represents the instances—specific cases or occurrences of precise combinations of variables at specific times under given conditions. The "best practice" approaches prevalent in KM and other consulting domains utilize this part of the concept space almost exclusively. The message of Figure Five is that it is hopeless to try to manage the infinitely large number of possible outcomes of events on the surface level.

From the earliest days of knowledge-based systems research, managing combinatorics has been a major obstacle. This is symptomatic of focusing most of the observation at surface-level phenomena. Further research into applied semiotics and ultra-structure theory may offer insights into how to make observations at deeper levels that help to unify principles across domains. By identifying foundational properties and relationships at sub-semantic levels, large portions of the search space can be quickly pared, thereby excluding possible yet implausible combinations of elements. Using this type of approach, the number of allowable aggregations becomes finite and manageable.

The application of Pospelov's model greatly enhances knowledge representation. Known universal truths and invariants are linked to direct observations as well as plausible outcomes not yet observed. This system of plausible reasoning uses a combination of John Stuart Mill's canons of induction and abductive inference to generate outcomes that might not have been considered otherwise [13].

For example, based on historical data, one could computationally estimate the cost of gasoline over the next five years. Through data mining one might enhance this projection by revealing correlations between the price of gasoline and: 1) the price of oil, 2) the number of negative news reports coming out of the middle east, and 3) the number of environmentalists serving in the US Congress.

Such knowledge seems adequate enough for a home delivery business to plan a five-year budget. But what if a terrorist attack disrupts the nation's fuel distribution system? The overall impact on price, availability, interest rates and commerce in general could very well put that firm out of business within a very short time. A true knowledge management system would use inductive reasoning to formulate possible deep structure relationships, and abduction to generate a set of plausible scenarios, which could be used to identify and correct a firm's most critical vulnerabilities.

Prueitt [14] provides a more extensive discussion of the application of Pospelov's and other Russian semiotic models to the problem of document understanding and the development of machine-readable ontologies. More exploration is needed, however, especially with regard to the formalization of semiotic logic systems.

Only a few initial prototypes of the conceptual space have been developed in Russian laboratories. A synthesis of Western ultra-structure theory and Russian applied semiotics would open new realms of possibilities for future generation knowledge architectures.

Advanced Visualization: Cognitive Computer Graphics

Knowledge generation typically occurs in two ways: through statistical analysis of data in which new insights are gained, or through Gestalt-type discoveries that emerge from human perception and cognition, stimulated by observation and interpretation based on situational awareness and past experience.

Problems arise when physical instruments are used to measure natural systems. Physical devices record physical attributes such as position, time, and rate of change. These are appropriate for physical models within a framework where Newtonian type interactions are plausible.

This is not necessarily the case for semiotic models, where measurement requires an intermediate organization of the data that can support visual cueing aimed at identifying invariants. This intermediate organization enables a "vetting" process that involves human perception and cognition about the correlational status between the events in the natural system and the events organized with the instrumentally derived data.

The specific outcome of an initial organization of instrumentally derived data can be considered to be somewhat arbitrary since the optimal organization of the elements comprising the target of investigation is not known. However, specific vetting processes can be designed and employed to produce a conversion of informational artifacts into knowledge, and eventually encode it in the form of semiotic model-based knowledge artifacts.

Placing the human-in-the-loop provides the opportunity to enhance conventional data mining with human perception for improved categorization based on non-computational similarity approaches. One such approach is the system of Cognitive Computer Graphics (CCG) developed by Professor Alexander A. Zenkin of the Russian Academy of Sciences.

CCG enables the creation of new knowledge, as evidenced by the numerous theorems discovered (and some old theorems disproved) through its use. The following simple example will serve to illustrate. One form of CCG visualization is the pythogram, in which a sequence of cells corresponds to the set of positive integers ($n \geq 1$), where each cell has a value of TRUE or FALSE. Graphs (a) and (b) on the left half of Figure Six show a pythogram of modulus 8 for the series of squares of natural numbers (1, 4, 9, 16, 25, 36, …). Graph (c) on the right half of Figure Six shows the visual patterns that emerged from expanding the modulus to 13 through 18, which resulted in the discovery of a repeating pattern under modulo 16. Graph (a) on the left hand side of Figure Seven shows the plot of n^2 vs. n, while graphs (b) and (c) of Figure Seven show the expanded view of the repeating pattern discovered at modulo 16.

The use of CCG pythograms by Zenkin in the analysis of Pall's Theorem led to the discovery of a new theorem. Pall's Theorem states:

> For $s \geq 6$, any natural number $n \geq 1$ is representable as a sum of exactly s squares of positive integers, except for the numbers 1, 2, 3, …, s-1, and the numbers of the form s + {1, 2, 4, 5, 7, 10, 13}.

Figure Eight shows Zenkin's visual proof of Pall's theorem, in which the 7-element invariant set (1, 2, 4, 5, 7, 10, 13) emerged at modulo 8, beginning with $s = 6$. Zenkin extended the visualization and discovered a new invariant set consisting of 75 elements: Z = (1, 2, 3, 4, 5, 6, 8, 9, 10, 11, 12, 13, 15, 16, 17, 18, 19, 20, 22, 23, 24, 25, 27, 29, 30, 31, 32, 34, 36, 37, 38, 39, 41, 43, etc.), which emerged at modulo 21 for $s \geq 14$ (see Figure Nine). The theorem, as revised by Zenkin, now states:

> For $s \geq 14$, any natural number $n \geq 1$ is representable as a sum of exactly s squares of positive integers, except for the numbers 1, 2, 3, …, s-1, and the numbers of the form s + Z, where the set Z is known explicitly.

CCG has demonstrated that at least for the domain of classical number theory, new knowledge can be created using a "human-in-the-loop" imaging process. By understanding this imaging process more fully, we expect to extend the creation of new knowledge to other domains.

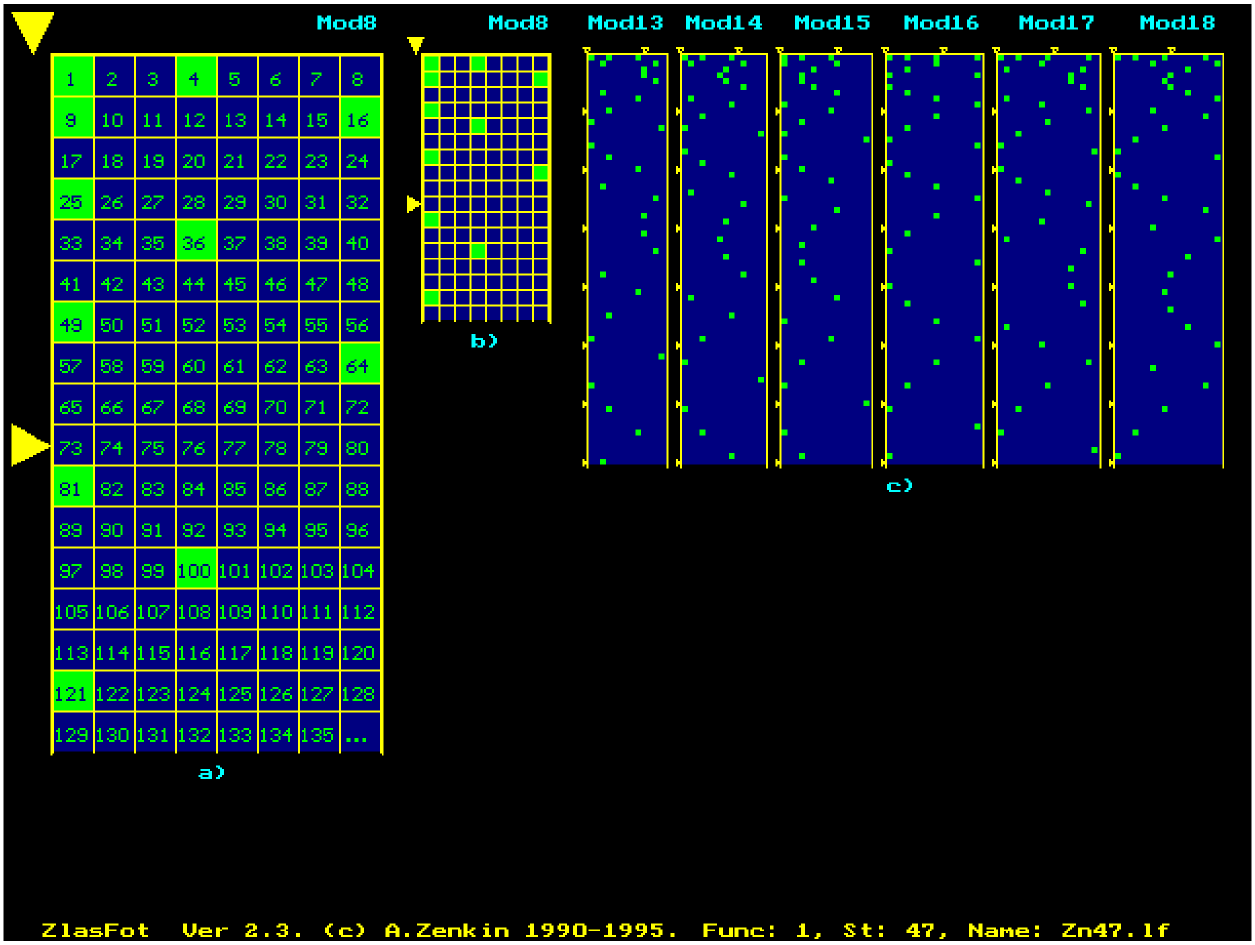

Figure Six -- Natural Pythograms of Moduli 8, 13-18) [15]

Summary

Taking knowledge management to the next level requires both computational and non-computational approaches. Applied semiotics bridges these two realms by applying a system of open logics (also known as second-order cybernetics) to an evolving set of axioms built on a foundation of time-invariant laws, yet allowing for a rich diversity of plausible aggregations of elements. This is how nature operates and this is how knowledge management should operate.

CCG provides a non-computational means of vetting these time-invariant laws. Together, the two technologies have the potential to break knowledge management free of the limitations of physical (stochastic) measurement systems, and into the more fluid semiotic measurement systems based on perception and cognition. Most knowledge management practitioners agree that the majority of human knowledge still resides in human minds. It makes sense, therefore, to move from purely computational methods to approaches that integrate computation, perception and cognition. It is here that our research efforts should be directed.

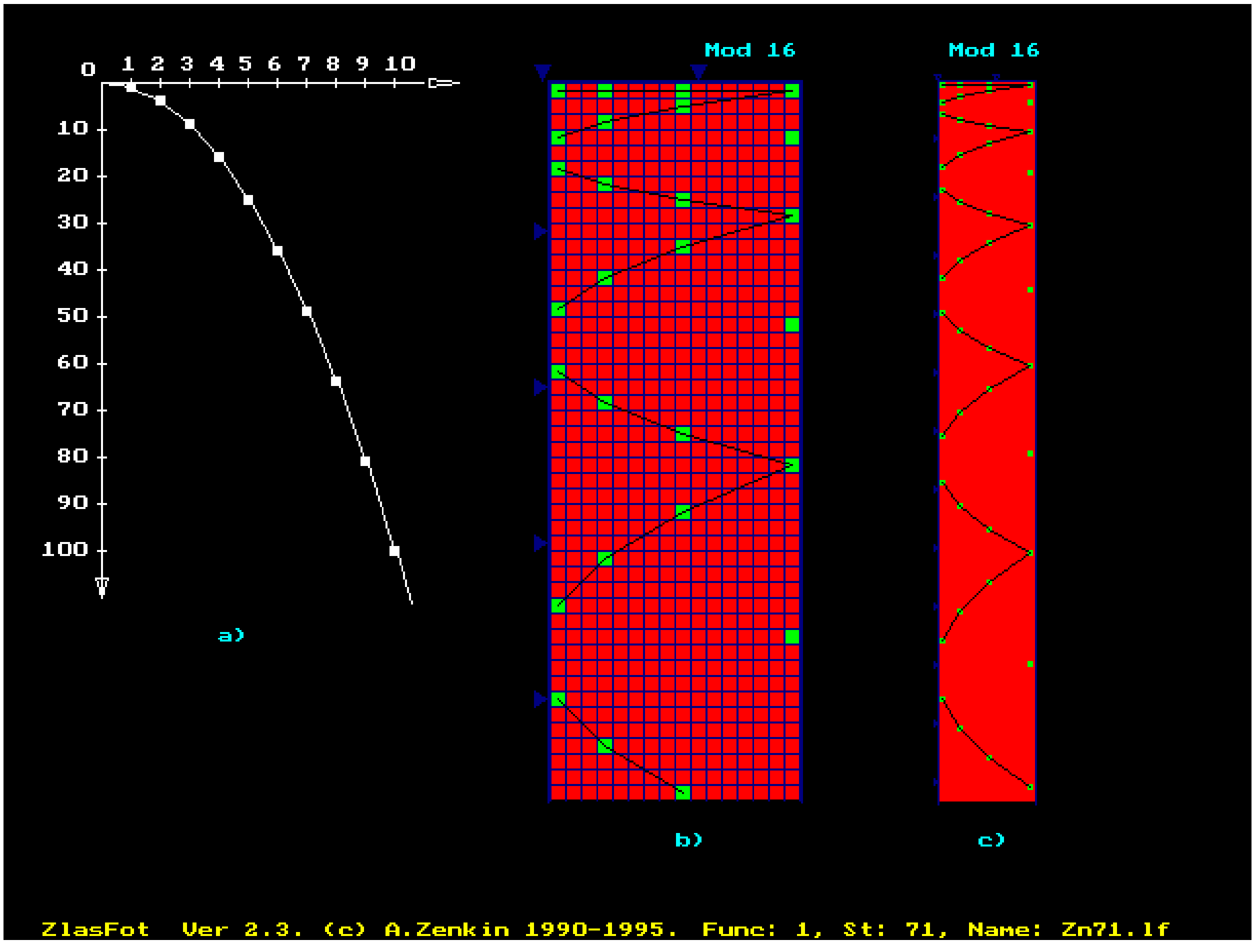

Figure Seven -- Natural Squares Pythogram, Modulo 16 [16]

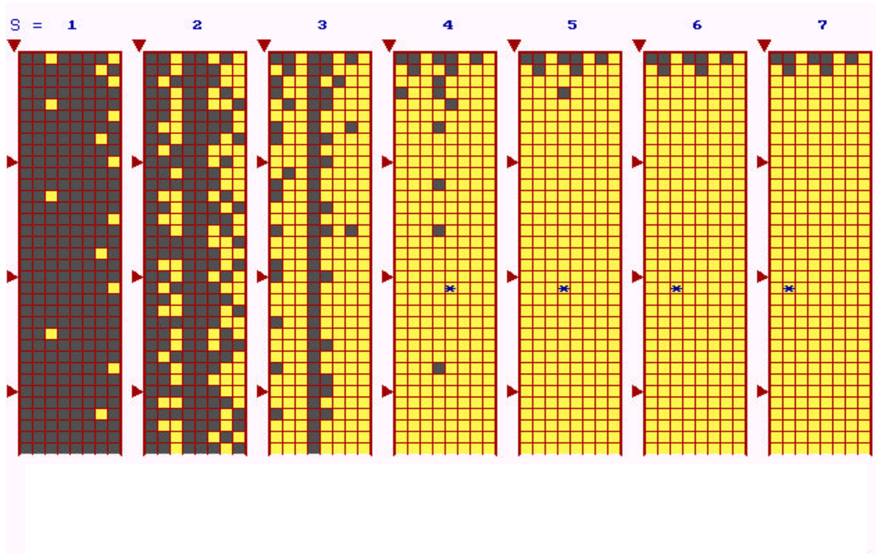

Figure Eight -- Visual Proof of Pall's Theorem, s ³ 6 (Modulo 8) [17]

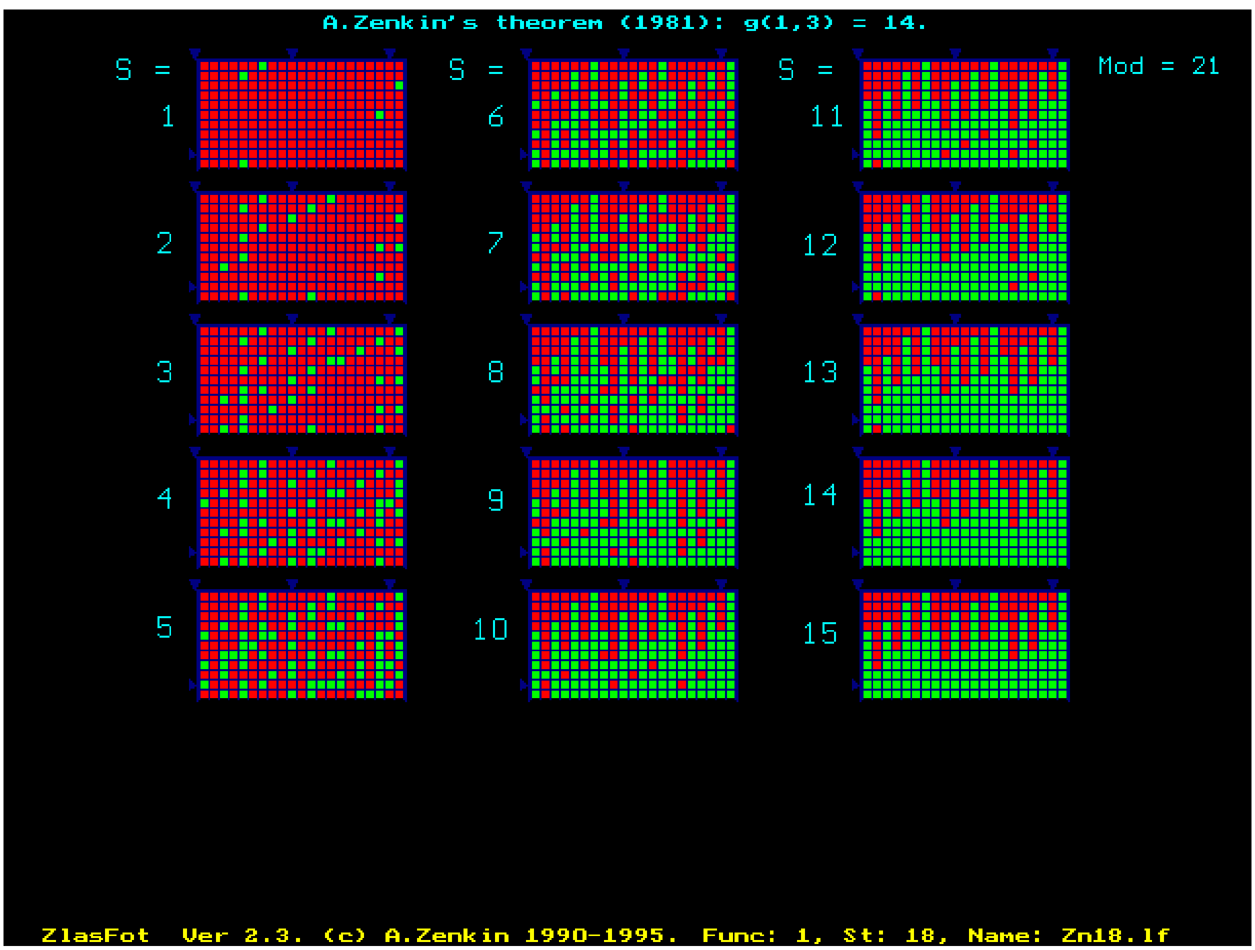

Figure Nine -- Visual Discovery of New Invariant Set at *s* ≥ 14 (Modulo 21) [17]

References

[1] Michael Polanyi. Personal Knowledge (Chicago, Ill: University of Chicago Press, 1958).

[2] Michael Polanyi. The Tacit Dimension (London, England: Routledge and Kegan Paul, 1966).

[3] Karl R. Popper, Objective Knowledge (London, England: Oxford University Press, 1972)

[4] Karl R. Popper and John Eccles, The Self and Its Brain (Berlin, Germany: Springer, 1977).

[5] Paul S. Prueitt, "A Theory of Process Compartments in Biological and Ecological Systems, in J. Albus, A. Meystel, D. Pospelov and T. Reader, eds., Architectures for Semiotic Modeling and Situational Analysis in Large Complex Systems, AdRem, Bala Cynwyd, PA, 1995.

[6] Karl H. Pribram, ed., Rethinking Neural Networks: Quantum Fields and Biological Data, ERA, Hillsdale, NJ, 1993.

[7] Dmitri A. Pospelov, "Large Semiotic Models in Control Systems," Proceedings of the 1996 International Workshop on Control Mechanisms for Complex Systems, Mike Coombs and Mark Sulcoski, eds., Las Cruces, NM, Dec 8-12, 1996, p. 287-292.

[8] Howard H. Pattee, "The Physics of Symbols and the Evolution of Semiotic Controls," Proceedings of the 1996 International Workshop on Control Mechanisms for Complex Systems, Mike Coombs and Mark Sulcoski, eds., Las Cruces, NM, Dec 8-12, 1996, p. 9-25.

[9] Dmitri A. Pospelov, Situational Control: Theory and Practice, Nauka Press, Russia, 1986.

[10] Anna Wierzbicka, Semantics: Primes and Universals, Oxford University Press, 1996.

[11] Jeffrey G. Long and Dorothy E. Denning, "Ultra-Structure: A Design Theory for Complex Systems and Processes," Communications of the ACM, January 1995, p. 103-120.

[12] Youngsuck Oh and Richard Scotti, "Analysis and Design of a Database Using Ultra-Structure Theory (UST)," 20th Annual ASEM Conference, Washington, DC, 21 October, 1999.

[13] Maria A. Mikheyenkova, "Application of JSM Reasoning to Problems of Sociology," in J. Albus, A. Meystel, D. Pospelov and T. Reader, eds., Architectures for Semiotic Modeling and Situational Analysis in Large Complex Systems, AdRem, Bala Cynwyd, PA, 1995.

[14] Paul S. Prueitt, "Semiotic Design for Document Understanding," Proceedings of the 1996 International Workshop on Control Mechanisms for Complex Systems, Mike Coombs and Mark Sulcoski, eds., Las Cruces, NM, Dec 8-12, 1996, p. 255-277.

[15] Alexander A. Zenkin, "Cognitive Explanation of the CCG Visualization Technology," Visiting Scholars' Seminar, The George Washington University Department of Computer Science, 25 April, 2000.

[16] Alexander A. Zenkin and Anton A. Zenkin, "Cognitive Computer Graphics: Beauty and Truth of Mathematical Abstractions," Proceedings of the International Conference on Mathematics and Art, Suzdal, Russia, 1996, p. 158-166.

[17] Alexander A. Zenkin, "Cognitive Graphics for Obtaining New Knowledge," Proceedings of the Workshop on Russian Situational Control and Cybernetic/Semiotic Modeling, Robert J. Strol, Ed., Columbus, OH, USA, 1996, p. 79-100.

Acknowledgments

The author gratefully acknowledges the sponsorship of the US Army Research Laboratory in the exploration of Cognitive Computer Graphics funded under ARL Contract #DAAD17-00-M-P021. Drs. Paul S. Prueitt and Alexander A. Zenkin, Co-Principal Investigators for the contract, contributed to the advanced visualization interfaces section of this paper.

Biography

Arthur J. Murray is president of Telart Technologies, Inc., a company he founded in 1993 for the purpose of helping organizations manage their corporate knowledge. He has been leading the successful Implementation of advanced information and knowledge systems within government and industry for over twenty-five years. He is an adjunct professor in the School of Engineering and Applied Science at The George Washington University, where he has taught at the graduate level for over fifteen years, including the school's first knowledge management course. He is Co-Founder of the Behavioral and Computational Neuropsychology (BCN) Group, an international research organization dedicated to achieving a deeper understanding of the natural structures and flows of knowledge.

He holds the D.Sc. and M.E.A. in Engineering Administration from The George Washington University and the B.S.E.E. from Lehigh University. He can be reached at artmurray@compuserve.com.