Cognitive Principles for Information Management: The Principles of Mnemonic Associative Knowledge (P-MAK)

Michael Huggett
Holger Hoos
Ron Rensink

August 31, 2006

Department of Computer Science
University of British Columbia
2366 Main Mall, Vancouver, Canada V6T 1Z4

Abstract. Information management systems improve the retention of information in large collections. As such they act as memory prostheses, implying an ideal basis in human memory models. Since humans process information by association, and situate it in the context of space and time, systems should maximize their effectiveness by mimicking these functions. Since human attentional capacity is limited, systems should scaffold cognitive efforts in a comprehensible manner.

We propose the Principles of Mnemonic Associative Knowledge (P-MAK), which describes a framework for semantically identifying, organizing, and retrieving information, and for encoding episodic events by time and stimuli. Inspired by prominent human memory models, we propose associative networks as a preferred representation. Networks are ideal for their parsimony, flexibility, and ease of inspection. Networks also possess topological properties—such as clusters, hubs, and the small world—that aid analysis and navigation in an information space.

Our cognitive perspective addresses fundamental problems faced by information management systems, in particular the retrieval of related items and the representation of context. We present evidence from neuroscience and memory research in support of this approach, and discuss the implications of systems design within the constraints of P-MAK's principles, using text documents as an illustrative semantic domain.

Key words: information management, memory prosthesis, associationism, semantic similarity, co-occurrence, spatio-temporal context, episodic events, associative networks, spreading activation

Contents

1 Introduction

Memory provides the raw materials for intelligence. Without memory an intelligent agent, whether mechanical or biological, would be unable to store and compare ideas, or respond appropriately to changing circumstances. Memory's associative moves fluidly between related facts, and to narrow in on distinct information. It provides us with the sense of a continuous personal identity, and a of our world. Improved memory forms the basis for greater intelligence by providing an expanded store of knowledge and experience. One theory concerning our evolutionary cousins speculates (based on physiology and anthropological artifacts) that *Homo Neanderthalensis* had, despite a smaller working memory, an excellent long-term memory that enabled them to compete with and sometimes even surpass us (Wynn and Coolidge, 2004). [1] Clearly then, enhanced memory expands our cognitive abilities, much as a crane extends our ability to lift, or a vehicle our ability to travel.

How can we design more sophisticated information-management systems still comprehensible to our fallible "stone-age" brains? Words on paper have always been subject to physical constraints—such as printing, shipping, filing, and shelving—that make their management burdensome. In modern times, *information management* has evolved from tactile filing of ink-based books, papers, and folders, to channeling an increasing flow of digital media—a task beyond human capacity. Instead, automated processes should summarize what users have available to them, facilitate their searches, and direct them to the information that best suits their needs.

Humans are predisposed to a particularly *human* structure of ideas: we can communicate across language and culture barriers and anticipate how others think (Ekman, 1971; Osgood et al., 1975). Thus to make information management systems more useful to a wider range of people, it seems reasonable to apply functional cognitive principles to data storage and retrieval. Daily experience shows us that our memories are highly associative: given a particular cue (e.g. *red*) our thoughts and perceptions stimulate a wealth of contextually relevant memories (*fire*, *apples*, *cherries*, *stop*, etc.). The general principle of *associationism* holds that higher mental processes—and memory in particular—result from connections between sensory and mental elements based on experience (Anderson and Bower, 1973), and associative network models are prominent in cognitive science as descriptions of memory, knowledge, and reasoning (e.g. Anderson, 1983; Collins and Loftus, 1975; Quillian, 1969; Raaijmakers and Shiffrin, 1981).

The physical medium underlying such mental functions is less intuitive: the brain is massively parallel processor with billions of neurons, each of which typically gains synaptic inputs from thousands of others. All neurons fire continually at varying rates, and researchers have only the vaguest notion of how this roar of activity resolves to our subjective experience of reality. Although connectionist models that simulate brain activity at a neural level (e.g. Rumelhart et al., 1986) have provided important insights into pattern learning, it would be an overwhelming task to model higher-level memory processes in such detail. Instead we embrace the motto "to the desirable via the possible" (Marr, 1982), and take a pragmatic interest in the overt *functions* of memory, which are more easily represented and implemented. We simultaneously avoid an emphasis on human *behaviour*, since behaviour is influenced by cultural contexts. Our goal then is to suggest culturally-neutral systems that are automated, simple, and easily understood: systems are ultimately created to fulfill human needs, and in doing so should use metaphors that are easily human-interpretable.

We propose the *Principles of Mnemonic Associative Knowledge* (P-MAK), a framework for information management systems that is both inspired and guided by several related fields: information retrieval, cognitive modeling, knowledge representation, evolutionary psychology, and neurophysiology. The resulting set of principles combines the functional strengths of machine computation and human cognition. In the short term P-MAK can suggest improvements to current information systems; in the long term it suggests the convergence of semantics, contextual interaction, and *human-centred information management* (HCIM). P-MAK exploits the brain's useful information-processing paradigms and presents information in a way that is familiar to human users—and thus easier to use. Although we make no claims of whether particular cognitive models are complete and accurate descriptions of mental function, they are useful starting points for information systems in the principles that they embody. Human memory has strengths to be mimicked, but also weaknesses to be scaffolded (as discussed in Section 2).

The P-MAK framework describes four sets of theoretical principles that structure and guide the development of information systems: they are divided in pairs into the *fundamental* principles of the properties of brains and the constraints of machines, and the *organizational* principles of the accumulation of semantic knowledge, and its contextual links to the real world. The *fundamental* principles of the framework are comprised of two sets of principles describing the necessary basis of processing in mind and machine (Section 4):

> *Mechanistic principles*. These concern the computational constraints on the design and operation of computing machines and information management systems, which should be familiar to

every computer scientist: minimal use of resources, scalability, heterogeneity, extensibility, and reliability. [2]

Anthropic principles. These concern the properties and constraints of human memory and knowledge that machines should incorporate to provide more effective user interactions: human information processing is fundamentally associative. It also requires simplicity, which leads the mind to quantize the world into discrete objects, and form meaningful abstractions as summary 'maps'.

The *organizational* principles are based on these, and describe how the structure of information can facilitate information-management operations, especially when supplemented by the context of time and place and using cognitive paradigms (Section 5):

Epistemic principles. These describe the basics of inducting, organizing, and retrieving information objects with an associative knowledge structure. Information objects are described pragmatically in terms of discrete *attributes*, which are assigned by a *perceptual* classification process. The similarity of objects is then judged based on the attributes that they have in common. Objects are retrieved both by queries that match specified attributes, and also *ostensively* via similar objects.

Situational principles. These describe the main contextual characteristics of human memory: encoding regular and co-occurring events, the intervals at which they occur, and the physical conditions that accompany them. The re-occurrence of events can be predicted from this contextual structure.

Together, P-MAK's principles imply a particular representation: associative networks that optimally and explicitly encode relations between information objects. When such networks are sparse and scale-free, they show advantages for organizing, exploring, and analyzing the information space (Section 6).

Some of the terms in this paper have been used elsewhere in different contexts. Since we take a functional approach, which focuses on information processing common to all brains, P-MAK applies first to information management for individuals. *Human-centred information management* (HCIM) is broadly defined by both human information-seeking behaviours, and to the cultural (usually business) context in which information-seeking takes place (Schlögl, 2005). We believe that our functional approach is complementary, and encourage the combination of functional and behavioural factors in user-friendly systems. The challenges of information management lie between the perspectives of function and behaviour, as determined by human characteristics and shaped by cultural context. Ultimately, we view HCIM as a cognitive prosthetic that extends memory to recover information that we have already seen, and augments perception to find new, useful information.

2 The Challenge of Human Memory

The *sapiens* brain was well-matched to our ancestral environment (Wynn and Coolidge, 2004), but a fast-paced computer-mediated society requires increased powers of recognition and recall in both *semantic* memory for general factual knowledge, and *episodic* memory for events and experiences (Tulving, 1972). [3] As a starting point for improved human interaction with information systems, we first examine the properties of human memory by way of its deficits, and how it has been explained by computational models. We then look at how artificial memory aids have been used to support memory-based tasks, and at systems specifically inspired by models of human memory.

2.1 Failures of Memory

Brains can simultaneously execute diverse memory-based tasks, such as planning, talking, risk assessment and the identification of stimuli. But as impressive as this may be, human memory also experiences functional lapses every day. Some key examples are:

Forgetting — The timely recall of important information is important to survival. Human memory can be fast and efficient, but its failures are often frustrating, even dangerous. The proposition that forgetting is nature's way of keeping our minds focused on important aspects of a changing environment (Anderson and Schooler, 1991) is small consolation when facts need to be accurately recalled on demand. Memory performance can be affected by poor initial encoding, such as through wandering attention, or may be the result of poor cues, such as being forced to recall context-dependent details in a different context. If memories are permanent, as some researchers believe (e.g. Bahrick, 1984), then there is much to be gained by making their re-activation more reliable.

Fallacious Integration — Given several descriptions of similar situations, humans tend to *integrate* the details erroneously into new hybrid memories. This phenomenon increases with the complexity of the descriptions: facets of a situation are combined into a single representation of the general idea, and this *gist*—not the details—is well-remembered. Memories created in this way are taken as fact unless some detail is so corrupted as to contradict the gist of the original descriptions (Bransford and Franks, 1971).

Presupposition & Inference — The gist itself is also subject to corruption. For example, experiments on eyewitness testimony have found that choice of words during questioning can alter memories of actual experiences, and subjects may even claim to have seen details that did not occur. Memory is therefore a blend of true facts and erroneous information implied later (Loftus and Palmer, 1974). Distortions also occur when people find themselves in a particular situation, and based on expectations choose an inappropriate behavioural *script* for that situation, biasing the gist and leading to misinterpretations that are mistaken for true memories (Brewer and Treyens, 1981).

Biased Memories — Changes in context can also have a significant effect on the interpretation of a past experience during recall. Although people tend to believe that, when queried, they draw from a stable semantic memory that has been faithfully abstracted from hundreds of episodes, experiments on episodic influence have shown that even a single experience can have a biasing effect and produce low-probability responses (e.g. Jacoby and Witherspoon, 1982).

Machines have properties that are complementary to memory's failures. Machine-based data storage is less prone to decay, and can be backed up for ensured permanence; where data indexing fails, brute-force searching can recover data reliably in a way that brains cannot. Machines can retain raw data with high accuracy, and different interpretations of the same information objects can then be evaluated as goals and perspectives change. The relative permanence of machine-based data storage makes it impervious to suggestibility, and although bias in brains is hidden and difficult to detect, the biases of machines are inspectable in the algorithms they use, and in how their code is written. Brains have powerful abilities, but machines are more reliable, and seem promising as supports and scaffolds for the miss-steps of human memory.

2.2 Computational Models of Memory

A great deal of work has been done to understand how human long-term memory operates, by developing cognitive memory models that describe the functions of human memory in implementable mathematical detail. These models follow the *information-processing* paradigm that treats the brain as a computing device, and the computer as a "thinking machine". As we prefer a simple, discrete approach, we focus on functional symbol-manipulating models, rather than more detailed and resource-dependent neural models; however these are not incompatible with our approach, as we shall see later. Each of these models accurately, though not completely, models various aspects of empirically-observed properties of memory.

The pioneering ***Teachable Language Comprehender*** **(TLC)** was written as a computer program with the goal of recreating human inferential ability. It organizes knowledge into a hierarchical conceptual organization much like a taxonomy tree, with a general root node (e.g. animal) connected to subordinate nodes (e.g. fish, bird, mammal, etc.) that are each connected to subordinates (e.g. salmon, tuna, pike, etc.). At each node, defining properties are included by listing physical characteristics and abilities. All of the characteristics of a superordinate node (e.g. animal → has-skin) are inherited by its subordinate nodes, so that a shared attribute only has to be defined once, giving a *cognitive economy* of attributes. The structure of TLC models the human *category size effect*: questions in larger domains take more time to search, and the time required is directly related to the number of links between nodes in the hierarchy (Collins and Quillian, 1969; Quillian, 1969).

Building on TLC, the ***Spread of Activation Model*** (Collins and Loftus, 1975) uses links of different lengths to indicate relative strength of association. The length of the links reflects the time required to activate related concepts, thereby encoding their *semantic distance* and *typicality*. Spread of activation also accumulates activation in the nodes related to the activated node; even when this may not be enough to make them fire, there is still a *priming effect* that allows nodes to reach full activation faster. In ***Associative Strength Theory*** (Anderson, 1983), memories are recovered according to how strongly related they are to a presented cue. Activation spreads to related memory traces, which rise into consciousness if their activation exceeds a given threshold. The theory explains observed phenomena such as slower response times when faced with more choices; this is modeled by the *fan effect*, which treats activation is treated as a finite quantity to be shared among all connected nodes.

Each of these uses spreading activation for relatively constrained computation in long-term memory, which appeals to our goals of simplicity and pragmatism. By contrast, ***global matching models*** perform retrieval based on the combination of cues in short-term memory, and operate in highly parallel fashion. One example is the ***Search of Associative Memory*** **(SAM)** model (Raaijmakers and Shiffrin, 1981; Gillund and Shiffrin, 1984), which describes long-term memory as a set of “images” that are “closely interconnected, relatively unitized, permanent sets of features” describing context as well as semantic content. For retrieval to occur, a set of cues is assembled in short-term memory; these cues activate the images to which they are connected. Parallelism reaches its zenith in models such as ***MINERVA 2*** (Hintzman, 1984), which assumes that a query is matched against all memories in parallel, and that all memories then respond in parallel, “the retrieved information reflecting their summed output”.

Although these computational models successfully mimic specific characteristics of human memory, they do not generalize well and fail outside the boundaries of their assumed conditions. A more general approach has been the development of ***cognitive architectures*** comprised of cognitively-justified tools and theoretical constraints, with the goal of performing a full range of human cognitive tasks; they are used to develop and test new cognitive models. The two most well-known architectures are ***Soar*** (Nason and Laird, 2005) and ***ACT-R*** (Anderson et al., 2004). ACT-R describes neural-like computation, and assumes that human cognition is optimally evolved to reflect statistical trends in the environment. By contrast, Soar is based on the premise that humans use knowledge in a rational way in order to achieve goals, and assumes that human cognition is a symbol system built upon a connectionist neural physiology (Johnson, 1997). Although more complete in their description of human cognition, cognitive architectures are too broad and powerful to serve as the basis of a cognitive information retrieval system. Together with global matching models, cognitive architectures are beyond the scope of information management; thus we look to more computationally tractable models for the storage and retrieval of information objects.

2.3 Memory Prosthesis

The challenge of *human-centered information management* (HCIM) is to compensate for the weaknesses of human memory while still taking advantage of its strengths. Its development has taken two paths. The first has been to develop systems that *support* human memory by whatever means—not that model, mimic, or explain it. The scope of these systems has been relatively narrow, acting as reminders or trackers, or sometimes as special-purpose task assistants. Their operation is seldom based on functional cognitive processes, and as such could be further improved. The second path has been to develop systems with a deep structure designed to work with the known characteristics of human memory. There have been relatively few of these, and as research testbeds they have seldom provided full functionality or been widely used.

Artificial Associative Memory Aids

The use of *mnemonics*—devices intended to assist the memory—began with cave paintings and progressed through clay tablets to the invention of paper. Computerized systems have since given us powerful indexing and search functions to organize large sets of items with relative speed and ease. Today portable personal devices (such as phones and digital notepads) are probably the most ubiquitous and familiar memory aids. They can access online information, and serve a useful prospective (or forward-looking) task if used to plan schedules and chime notices, but they do not yet fulfill the promise of human-like *associative* information management, as presciently described by Vanevar Bush (Bush, 1945). The invention of the computer led Bush to imagine a hypothetical multimedia information system called *memex* that would contain all of a person's books, records and communications in "an enlarged intimate supplement" to memory, "mechanized so that it may be consulted with exceeding speed and flexibility". It would store enduring associative "trails" of items collected on a subject, to provide a form of clustering, navigation, and memory cuing. [4]

Although Bush's vision exceeded the technology of his time, its inspiration has since driven researchers to create personal information management systems that record the salients of a person's activities (Want et al., 1992), automatically index everything saved on their computer (Dumais et al., 2003), and display personally-relevant items on a timeline (Fertig et al., 1996). However, these systems do not allow users to create "trails" between objects as would Bush's memex. Although other systems allow users to manually create their own link trails (e.g. Gemmell et al., 2002), they do not automatically induct and link objects. None leverage associative similarity between objects to build memory-like knowledge structures. Rather, modern memory aids typically focus on indexing and mining items from a database, primarily by leveraging user annotations and searching on existing property fields.

A few commercially-available systems employ an associative model, with templates describing various ways of organizing and displaying archived information. MindManager (nd) and TheBrain (nd) provide little in the way of automated classification of information, but are rather enterprise-wide "knowledge platforms" that consolidate data into a single repository and provide users a common graphical interface; they require users to input and link items together explicitly to create a communal knowledge structure "map". Where the goal is the "automation of unstructured information", associative enterprise systems can be complex and expensive. Autonomy (nd) is a state-of-the-art Bayesian-inference system that provides real-time information retrieval: it analyzes words as they are typed and opens new windows showing related news releases, archived reports, and diagrams, and also displays the contact information of related experts.

<u>Cognitively-Inspired Systems</u>

Although the fields of cognition and computation are sometimes directly compared (e.g. Anderson, 1989; Foltz, 1991), the field of memory has not yet enjoyed the same migration of concepts between human sciences and computation as seen for vision in (Marr, 1982). Computer simulations are commonly used in cognitive science to support particular cognitive theories, but are rarely used to form the basis of an actual information system. One notable exception is the ***Memory Extender*** (Jones, 1986), a "personal filing system" that seeks to combine the benefits of human memory and electronic storage. Its network representation reflects the principle of *associationism*, while its weighted node-and-link architecture connects computer files with the terms that describe them best, in an analog of symbolic memory. A query using a particular term activates the term's node, which then passes activation to the nodes of files that contain the term. Queries with multiple terms increase activation in the nodes that share those terms, which then rise to the top of a ranked list of retrieved items. Terms can also be combined with reference to a *context* node, which is associated with new files as they are created under particular circumstances. The system also decays the activation levels of unused nodes so that they are eventually "forgotten". The system is intended as an enhancement (or replacement) of the standard file-and-folder (FaF) desktop idiom. As such, it is concerned only with file retrieval—it does not extend its model, for example by connecting files directly based on shared characteristics, or by offering an overview of themes in the corpus.

A very different but significant development is ***Latent Semantic Analysis* (LSA)** (Deerwester et al., 1990), an indexing system that scores candidate documents against a representative corpus, basing its similarity scores on the co-occurrence patterns of keywords within them. It has been successfully used for various tasks such as cross-language information retrieval, information filtering, text analysis and essay grading. While LSA has shown expert human-like classification abilities (Landauer et al., 1998), it has some drawbacks. As a matrix-based method in a high-dimensional space, it requires significant amounts of storage space and computation, and as a black-box method its parameters must be hand-tuned for each collection. It is a batch process that requires the induction of a large training set before use, the semantic dimensions of which are determined algorithmically at an abstract mathematical level and difficult to describe in humanly-meaningful terms. Once the knowledge structure has been built, it cannot be easily edited or updated with new concepts

(Lemaire and Denhière, 2004; Zha and Simon, 1999). While LSA has been advanced convincingly as a cognitive knowledge model of synonymy (Landauer and Dumais, 1997), its vector-based approach assumes symmetrical similarity between terms, which contradicts psycholinguistic findings [5].

3 Human-Centred Information Management

For computing machines, "the performance of the device is characterized as a mapping from one kind of information to another" (Marr, 1982, p.24), a transformative process that takes a stream of raw input and interprets it to produce a structured output. Here we apply this perspective to information-management systems with respect to the functions of human memory. The goals of information retrieval are two-fold: the prospective finding of new relevant information, which maps new information into a growing knowledge structure, and the retrospective recovery of information that has already been seen, which maps user queries and behaviour onto relevant retrieved objects. These mappings are equally pertinent to human memory and computational information systems.

3.1 The Basic Operations of Information Management

Human memory varies according to changing goals, which affect how items are perceived, encoded, organized, and retrieved (Barsalou, 1983; Mandler, 1984). As humans begin to recall information and use it to decide their next actions, they refine their search iteratively by repeating an interaction loop that brings them closer to their goal. The data stream then involves more than a direct mapping from input to output: learning and retrieval are inextricably entangled, as retrieval primes memory to prioritize items similar to those that have previously proven themselves of interest (Figure 1). This sort of retrieval interaction between attentional processes and long-term memory is typical of prominent human memory models (e.g. Anderson, 1983; Raaijmakers and Shiffrin, 1981).

Information management systems should be designed to support and exploit this human characteristic, and for this require an appropriate, simple, pragmatic, easily-understood ontology: discrete *objects* that may be organized into sets, and a set of *basic operations* that create and manipulate these objects. In the broad sense, an information object is comprised of two parts: a physically-instantiated *entity*—typically a discrete stored datum such as a document or image—and a descriptive *reference* that serves as a pointer to that

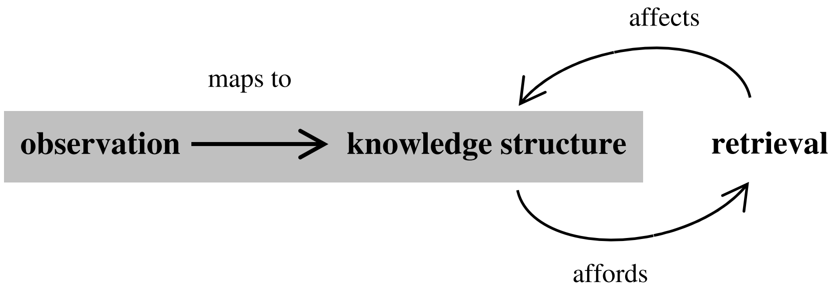

Figure 1: Information-mapping within minds and machines. The shaded area reflects a mapping from raw data to representation. Memory retrieval adds a feedback loop in which user actions adjust the knowledge structure to reflect the context of a user's interests. The mapping process itself, if adaptive, can be dynamically biased toward higher-value objects.

entity. The object reference sits within the information management system, which uses it to retrieve the entity, to compare it to other data, and to determine clusters of topics. By contrast, the entity is stored in a database that may be at some remote location. These basic operations and objects necessarily apply to any agent—human or otherwise—that interacts with an information source:

- Specification — the definition of an object's ontology: its structure, composition, properties and elements, and how objects may be related to each other.
- Induction — the incorporation of objects into a knowledge structure.
- Modification — the mutation of objects and their relations intended to increase their utility, either by inferring trends or through direct user intervention.
- Retrieval — the recovery of desired objects from within the knowledge structure.

3.2 Mind-Machine Symbiosis

Machines need not *be* like humans, but human-machine *symbiosis* is important for systems to function as an extension of human memory. Physical human-machine comparisons are misleading. How machines achieve this symbiosis cannot depend on mimicking the details of brain architecture: the brain contains billions of massively-connected neurons, versus an upper bound for machines of only a few thousand linked processors. This difference in scale is likely to hold for some time. Wang and Liu (2003) estimate the relative capacity of human memory at a staggering $10^{8,432}$ bits. Comparative estimates of processing power show a similar imbalance: Moravec (1998) estimates human processing power at an impressive 100 million MIPS (Million Instructions Per Second), or 10^{14} individual 'commands'. Yet despite the massive parallelism of the brain, the human capacity for attending to multiple information sources is very limited (e.g. Baars, 1993; Kahneman and Treisman, 1984). For this reason human interaction with information sources typically involves a *sequence* of simple actions that narrow in on a target: the serial presentation of a few well-chosen items is less overwhelming than presenting an entire corpus all at once.

The feedback loop in Figure 1 applies also the interactive nature of information management, which requires a sequence of interaction between user and information system. Interaction is improved if the system's function is easily understood by the user, and if the user's actions are properly interpreted by the system. The design of effective systems requires that users should be presented with a set of rational and predictable behaviours, while machines should relieve users of more repetitive or complex operations, and adapt to user preferences Hoffman et al. (2002).

Current information systems are prone to a number of frustrating problems. Even technically-fluent users experience operational failures on a regular basis, such as hardware failures and accidental deletions, and so continue to use paper documents (Whittaker and Hirschberg, 2001). Automated sorting in a file-and-folder (FaF) system is rare, which forces users to judge the relevance of their files repeatedly, to create (often naïve) *ad hoc* organization schemes, to remember where things are stored, to remember what search terms are applicable, to search exhaustively for lost files, and so forth. Other problems touch on the nature of knowledge, such as how machines should recognize and manipulate semantic subtleties, or how they should adapt to unanticipated changes of context. Other problems are purely algorithmic, and affect the efficient management of data in large, ever-growing archival storage systems. The goal of the P-MAK framework is to define constraints for any system that wishes to avoid these problems.

3.3 Introduction to Principles: P-MAK

If minds are to interact well with machines, care must be taken to design systems that support the best aspects of both. The principles that apply are divided by two general aspects. The first concerns the scope of the constraints: *universal* principles are applicable to any intelligent agent, while *human-centred* principles are specific to the nature of human cognition. Universal principles originate in the information-management task itself, independent of the agent that performs it, and define general computational considerations for the

		APPLICATION	
		OPERATIONAL	ORGANIZATIONAL
GENERALITY	UNIVERSAL	mechanistic Parsimony Scalability Scope Plasticity Robustness	epistemic Identification Perception Similarity Navigation
	HUMAN-CENTRED	anthropic Associationism Simplicity Quantization Abstraction	situational Persistence Retrospective Indexing Prospective Indexing Environmental Indexing

Figure 2: The P-MAK framework. In terms of *scope*, the universal principles apply to any optimal general-purpose intelligent agent that interacts with information, while the human-centred principles describe the nature of human information processing: its key strengths and constraints. In terms of *function*, fundamental principles define the (internal) constraints and qualities that affect knowledge structures, while organizational principles concern an intelligent system's interpretation of and interaction with the external world, of which the system's user is a part.

design of optimal systems. For example, that an operation should remain tractable as the size of a dataset increases is clearly desirable, regardless of how an agent processes information. By contrast, human-centred principles include cognitive properties that would improve interaction between users and information management systems; as such, they are descriptive rather than prescriptive. For example, that humans require meaningful summaries is clearly human-centred, since it is uncertain that any non-human agent (natural or artificial) would necessarily operate this way.

The second general aspect concerns a principle's functional application: *fundamental* principles define the (internal) constraints and qualities that directly affect an agent's knowledge structures, while *organizational* principles concern a system's interpretation of and interaction with the external world including the user. That data structures be scalable is a fundamental principle, since it is independent of both the semantics of the data and the goals of the user. Meanwhile, comparing objects based on shared attributes pertains to an agent's organization of information to make sense of its environment. Although fundamental and organizational principles can be separated out in theory, in practice they often interact. Indeed, the fundamental principles set the constraints by which the organizational principles function.

The aspects of scope and function divide the principles as follows (Figure 2):

1. Mechanistic principles (universal,fundamental):

   The necessary properties for efficient computation and data retrieval.

2. Anthropic principles (human-centred,fundamental):

   The inherent properties of human memory that would usefully be incorporated into information-management systems.

3. Epistemic principles (universal, organizational):

   The processes necessary for an intelligent agent to induct, classify, and retrieve information.

4. Situational principles (human-centred, organizational):

   The environmental aspects of how humans organize knowledge: by combinations of co-occurrence, time, and physical context.

Next we lay the cognitive and computational foundations, after which we apply them to the creation of knowledge and its contexts.

4 The Fundamental Principles

The purpose of computers is to help people think: computers then become as much extensions of human cognitive efforts as reading, writing or drawing. But computers often do not appear collaborative, since they are primarily designed to maximize various hardware properties such as bit depth and MIPS. The need for improved usability has led to the founding of *human-computer interaction* (HCI), a field that seeks to facilitate the flow of information between humans and computers (see e.g. Baecker et al., 1995). The goal of HCI is not to improve the fundamental operation of machines *per se*, but to interpose a layer of *translation* between human and machine. Thus HCI seeks to accommodate and support cognitive properties at the interface level, but seldom applies them directly to machine *function*.

P-MAK takes the opposite approach: it proposes that information within machines should be stored and retrieved in a manner that is inherently biomimetic (i.e. based on forms in nature) so that it is *inherently* comprehensible. Human memory represents a clearly successful approach to information retrieval and processing. One goal then is to transfer this efficient memory structure to machines for better organization and retrieval of information. Conversely, the implementation of evolved human information-management solutions, such as associative retrieval, makes machines more comprehensible by providing a familiar (in fact *ingrained*) information-management paradigm, instead of an *ad hoc* system-specific one.

4.1 *Mechanistic Principles*: Making Machines Effective

Compared to brains, the relative simplicity of machines leads to a set of fundamental computational concerns that must be assumed for effective information-management implementation. These are encompassed by the *mechanistic* principles:

Parsimony — *Systems should minimize the use of storage and computational resources.* Computing systems can clearly maximize simplicity and efficiency by reducing the usage of computational resources (such as cycles or storage space) to a practical minimum. With good design of algorithms and data types, more can be computed in less time. Even as computers become more powerful, the questions that users ask of their systems will become more demanding, requiring continued use of appropriately minimal designs that maximize efficiency. Parsimony has a direct influence on choice of data representation and tractability of computation (Smith, 1996; Moravec, 1998).

Scalability — *Systems should use scalable structures for efficient retrieval in growing datasets.* As information repositories grow, systems should continue to retrieve their items in a reasonable amount of time. But information systems do not automatically scale well as more data is added, and while most people are willing to wait a few seconds for a search to complete, beyond that they become impatient (Wickens and Hollands, 1999). *Indexing* algorithms build reference structures that group the items of a corpus semantically so that items can be more easily found. Indexing should partition the semantics of an archive in a balanced and organized way, allowing rapid navigation from the general to the specific. For example, indexing is straightforward where items can be sorted alphabetically in binary search trees, guaranteeing fast log-time access, but complex semantic spaces can resolve to many dimensions—hundreds or even thousands (Burgess and Lund, 2000; Deerwester et al., 1990)—which would be intractable to index in real time. Access times can be kept low by pre-calculating and storing relationships, but then the incorporation of new data may require an expensive recalculation of corpus-wide properties.

Portability — *Systems should apply readily to different and diverse domains.* Systems should be adaptable to new uses and new applications with a minimum of effort. In the domain of information retrieval, the classification of data is dependent on its type; a classifier appropriate to sorting textual documents is useless for images or sound (Witten et al., 1999). A heterogeneous dataset will therefore require a classifier appropriate to each of its expected data types. For machines to fulfill their function as "thinking tools", ideally they should be able to accept data from any domain, since mining through such an enlarged and diverse set could find new and interesting relations.

Plasticity *— Systems should use structures that are easy to reconfigure.* A simple representation that can be quickly elaborated and updated to reflect changes in data relations would be ideal, and contributes to *parsimony* by improving efficiency and reducing usage of resources during reconfiguration. Plasticity is useful where inferring categories depends upon the needs and perspective of users; the data representation should promote assembly of ad hoc categories that match user interests. The antithesis of plasticity is found in the traditional database model, where every field in a data record is of a particular type (e.g. an integer), is allocated a set amount of memory (e.g. 32 bits), and represents a fixed property (e.g. name, address, age, income, etc.). Few such presuppositions can be made about the content of data in a heterogeneous, ever-changing, real-world situation.

Robustness *— System operation should degrade minimally as the quality of information deteriorates.* The brain is extremely good at reasoning under uncertainty, since human survival has depended upon making mostly correct choices under uncertain conditions. Information systems, on the other hand, follow fixed instructions; traditional database models do not easily support robust retrieval with imprecise queries, which confounds users who don't know the words best describing an item, or who erroneously consider different combinations of words to be equivalent. To counter the requirement of exact query terms, information retrieval systems have used statistical term weighting, truncation, and synonymy (Witten et al., 1999); other methods such as vector-space, singular-value decomposition, or Bayesian comparison of terms between documents have also been successfully employed to increase robustness (Foltz, 1991).

4.2 *Anthropic Principles*: Making Knowledge Comprehensible

Human memory exhibits various characteristics that should be applied to information management systems; these are described by the *anthropic* principles. The goal of these principles is to meet ingrained user expectations of how information should be managed, and to suggest possible new algorithms.

Associationism *— Human memory is—functionally speaking—associative, and the most important associations are semantic.* Cognitive science has provided strong evidence for the semantic associativity of human memory, which affects most aspects of acquiring, structuring, and exploring knowledge. *Semantic priming* has been observed where subjects judge word pairs to be real words more quickly if they are semantically related (Meyer and Schvaneveldt, 1971), and words seem to be encoded by their semantic relatedness (Blank and Foss, 1978). Associationism lends itself to straightforward implementation on a machine, essentially as a stimulus-response model in the functional tradition of Skinner (1977). As such it includes well-specified learning behaviours such as the association of co-occurring pairs and the principle of reinforcement (Hebb, 1949).

Simplicity — *Humans cannot easily use information if it is too plentiful or complex.* To manage information effectively, humans need clear and simple representations. Humans can hold only about 4 to 7 items in working memory (Cowan, 2000; Miller, 1956), and have inherent bounds on their rationality which require shortcuts in reasoning (Todd and Gigerenzer, 2000). Information systems should therefore emphasize simple knowledge structures where possible, to avoid overwhelming human comprehension—simplicity refers to the limits on quantity of information that can be absorbed at one time.

Quantization — *Humans perceive the world as a collection of discrete objects and concepts.* Quantization describes how humans divide up the world into types and tokens. This "chunking" seems to derive from evolutionary pressures: symbolic abstraction is a fast and efficient way to identify and reference known entities and their descriptive properties. Since language is commonly modeled as a symbol system with a generative grammar (Chomsky, 1965), and since human thought seems to occur in the context of discrete words (Carroll and Whorf, 1956), biomimetic information management can be justifiably based on a symbolic cognitive paradigm. The ***Language of Thought*** hypothesis (Fodor, 1975) further supports the notion that thought is explainable by the manipulation of symbol tokens, that complex ideas are compositions of simpler (atomic) symbols, and that symbols are combined in the same *structure-sensitive* compositions as language. Indeed, quantization appears to be a *necessary* principle: the alternative continuous ontology, the representation of one's world by a real-valued density function, would require an infinite amount of information and calculation to produce reactions and decisions (Smith, 1996).

Abstraction — *Humans require meaningful summaries in order to make sense of their world.* Although instance theories of human memory claim that people store images of all their experiences (Medin and Schaffer, 1978; Hintzman, 1984; Nosofsky, 1984), there is no doubt that human intelligence is primarily defined by the ability to abstract—essentially, to generalize and make theorems for rapid, informed decision-making.

Complex information can be made comprehensible through good organization, which can be described in two ways. First in informatic terms, organization can be modeled by similarity functions and by clusters of concepts. Second, organization is central to comprehension, as humans construct interpretations around a central idea (e.g. Bransford and Johnson, 1973), classify identical items differently depending on current goals (e.g. Barsalou and Sewell, 1984), and attribute different meanings to items depending on context (e.g. Labov, 1973; Lakoff, 1987; Medin and Schaffer, 1978). Abstraction produces better quality of information: inadequately accurate or descriptive summaries tax our limited human attention.

5 The Organizational Principles

The organizational *epistemic* and *situational* principles describe how knowledge is created and used. The epistemic principles describe universal processes for encoding and retrieving *semantic* information, while the situational principles describe *non-semantic* organization according to regularities of time and environment.

5.1 *Epistemic Principles*: Building Knowledge

The *epistemic* principles describe the transformation of raw information into useful knowledge; as such they encompass the *basic operations of information management* described earlier. The principles of *identification*, *perception*, and *similarity* describe how information is encoded and organized into knowledge, while the principle of *navigation* describes how stored information is retrieved. Raw information is inducted into discrete objects (the principle of *identification*) that are identified by their salient properties (*perception*), organized based on alikeness (*similarity*), and retrieved through an interactive process of one or more steps (*navigation*).

Identification *— Objects are discrete, and described in terms of semantically discrete attributes.*

Identification is a necessary information management, following on the principle of *quantization*: the alternative to discreteness, a continuous density function, is intractably complex. Identification specifies how a semantic object describes the entity that it represents. It is based on the idea that humans tend to notice the salient *attributes* of an object, and use those attributes to compare objects and generalize abstractions about their world, as described in a number of *feature-based* knowledge models (e.g. Medin and Schaffer, 1978; Posner and Keele, 1970; Rosch and Mervis, 1975). *Feature Set Theory* (Smith et al., 1974) is typical of these, and describes objects in terms of *defining* features (those shared by all objects of a class) and *characteristic* features (those common but not essential). In information management, words that consistently identify a type of document would be descriptive, while words giving minor differentiation between documents of that type would be characteristic.

Attributes in the P-MAK framework need not be limited to keywords, but could also encode continuous values in discrete dimensions. In an information-space of related concepts, an *attribute* can be defined as any semantic description that is meaningful to humans. Osgood (1952) uses a set of Likert scales of opposing qualities to describe a concept; choosing a real-valued number on a scale denotes the relative strength of its qualities. For example, the qualities wet–dry and active–passive, rain would score toward wet and active, whereas sand toward dry and passive. Other methods define words in terms of semantic *microfeatures* such as humanness, softness, and form (McClelland and Kawamoto, 1986), or derive the semantic components of words by collating the subjective valuations of a large number of people (McRae et al., 1997). These are good ways to describe entities in terms of *meaningful* continuous attributes (cf. LSA: Landauer and Dumais, 1997; Lemaire and Denhière, 2004), but they also require labour-intensive surveys of human subjects to determine where a particular object lies on each scale. Automated extraction of attributes would clearly be preferable in real-world situations with large corpora.

Although *identification* occurs at a level above the operation of individual neurons, some fundamental concepts may be represented in the brain in the form of neuronal clusters (such as the impression of *redness* being contained in a cluster in the visual cortex), while compound concepts are synchronously distributed in clusters throughout the brain (such as the colour, texture, and behaviour components of a *dog* node). Thus an associative knowledge structure can be mapped onto the physical structure of the brain by assuming that an object or attribute is an abstract representation of a neuronal *cell assembly* comprised of some 10K to 100K neurons (Goertzel, 1997; Huyck, 2001; Pulvermüller, 1999). Indeed, the firing patterns of actual neural signalling systems have been be interpreted as "the state symbol-alternatives of the higher-level symbol-processing description" (Cariani, 2001). Such comparisons suggest a link between low-level, high-granularity connectionist neural networks and more abstract symbolic structures.

A symbolic approach to *identification* has several clear advantages. Whereas file-and-folder indexing and the desktop metaphor of typical information management systems emphasize an object's physical location, an object identified by its attributes can be retrieved through *content-based addressing*, so that remembering where it is is less important than specifying *what* it is. Attributes are human-readable, conforming to the principle of *simplicity*—objects are then comprehensible given the set of component attributes by which they are identified. Thus *identification* is fundamental to any information management system that uses a language or symbol system to reference distinct objects.

Perception *— Objects should be distinguished and assigned descriptive attributes by perceptual classification.*

Perception isolates and identifies individual objects by extracting and registering their attributes. Automated attribute extraction is essential for information management systems, particularly where large numbers of new, uncategorized items are inducted. This requires the design of individual classifiers sensitive to the key attribute types of a particular domain—for instance, the features of images are different from those of sounds, and are processed differently. The constraints of human perception and comprehension usually determine what qualities are important in each case, but since what is notable about an object often varies according to the perspective of the observer, the best strategy is to extract as many of an object's salient attributes as possible so that the object may be interpretable in many different contexts.

Information systems use the equivalent of perception to perform adequately in their "environment" of data streams. For example, data-mining systems may be considered *perceptual* since they detect patterns in data sets. However, the patterns that they find may be difficult to define in human terms—this is a larger challenge of automated intelligent tools: to explain clearly the basis of their decisions. If their actions are not understandable and inspectable, it is difficult to know whether they actually provide their intended service. The goal of a *perceptual classifier* then is to observe a real-world entity, create a data object to reference it, and extract the comprehensible semantic attributes that best describe it. Each object's attributes are weighted to reflect how descriptive they are; attributes above a certain threshold will be interpreted as *descriptive*, while the rest are *characteristic*.

Since documents are a "cognitively plausible" source material for a semantic system (Frank et al., 2003; Landauer, 2002; Lemaire and Denhière, 2004), they make an instructive example domain. Information retrieval is concerned with the classification, clustering, and recovery of documents in large collections. *Tfidf*, an acronym for "term frequency, inverse document frequency" (also written *tf-idf*), is one of the simplest and most illustrative classification algorithms (Salton and McGill, 1983). It extracts human-readable identifying attributes ("keywords") for each document in a corpus, and thus acts as a perceptual classifier for documents. The intuition behind *tfidf* is that frequently-appearing *terms* (i.e. words) in a document tend to be most descriptive of that document, and should be used as its keywords. At the same time, in a corpus of documents describing for example the many fates of fish, the term *fish* is itself clearly a poor choice for a keyword, since it will not help to differentiate between documents. If the corpus has many subtopics treating the different species of entrée, more specific terms such as *salmon*, *herring*, and *bream* would make useful descriptors.

With N the number of documents in a corpus, n_k the number of documents containing term k, and tf_{ik} the number of times term k appears in document i, the expression—

$$\begin{aligned} idf_k &= log\left(\frac{N}{n_k}\right) && \text{for some term } k \\ w_{ik} &= tf_{ik} * idf_k && \text{for some document } i \end{aligned}$$

guarantees, in the first line describing inverse document frequency, a bias against terms that appear in many documents in the corpus. The second line calculates the *weight* of each term in a given document; the number of times that the word appears in the document is tempered by the term's corpus-wide *idf*. Once all the words in all the documents in the corpus have been weighted, the top-weighted words in each document are used as that document's attributes.

Classifiers such as *tfidf* are consistent in principle with cognitive *global matching models* of memory such as SAM (Raaijmakers and Shiffrin, 1981) and MINERVA 2 (Hintzman, 1984), which describe a simultaneous matching of cues against all images in long-term memory. Unfortunately, the computational expense of global matching models is impractical for serial machines: in the case of *tfidf* all of the words in the corpus are counted before a document's keywords can be extracted. Clearly it would be costly to repeat such a search as new objects are added to a large corpus. Fortunately there are more parsimonious alternatives such as (Matsuo and Ishizuka, 2004), which uses the distribution of term clusterings in individual documents to extract their keywords, and shows comparable performance to *tfidf* without requiring a corpus-wide summation.

The implication of *perception* for the design of information management systems is that a homogeneous (e.g. document-only) data source requires just a single classifier tuned to its domain. Heterogeneous databases pose a greater problem: classifiers are necessarily data-dependent. For instance, the automatic extraction of attributes from images should cope with the specifics of pictorial data (e.g. Barnard and Forsyth, 2001). Thus a suite of classifiers will be needed to cover each of a range of expected data types.

Similarity — *Object similarity should be based on shared attributes.*

Objects are compared by the attributes that they share: the more two objects share attributes, the greater their similarity score. Similarity valuations between all pairs of objects are then used to build semantically-based associative knowledge structures. In humans the judgment of similarity is a real-time process that compares an object's attributes against those of other objects stored in memory. Objects are organized first by their sensory qualities, and as people learn more about them their mental representations gain additional attributes and become more strongly associated with similar objects. The *similarity* principle applies this process to the organization of information within information-management systems.

Objects that have all attributes in common are highly similar, while those with none in common are completely dissimilar. More accurate judgments are possible if attribute weights (such as provided by *tfidf*) are considered. For example, a document with descriptive attributes "fish, depletion, conservation, migration" would be related to another document with attributes "recipes, baking, cream, fish" at least to some degree, since both documents elaborate further on the fate of fish. Many different similarity equations are possible, depending on the desired ontology and how it should be tuned by weighting its components (see e.g. Tversky, 1977)

For example, the following assumptions would be congruent with P-MAK's principles:

- Attributes that are more strongly weighted are the *dominant* attributes of the object (analogous to the *defining features* of Feature Set Theory; Smith et al., 1974).
- Two objects that share dominant attributes are more similar than are two objects that share weaker attributes.
- Two objects that share dominant attributes but whose weaker attributes diverge might be two different descriptions of the same thing.
- Where one object's set of attributes is larger than that of another object, the attributes of the latter may simply be incomplete.

From these constraints, the following simple similarity measure can be derived. Given two objects defined by attribute vectors N and M of arbitrary lengths—

$$similarity(n, m) \quad = \quad \psi \left(\sum_{i \in N \cap M} w(N_i) + w(M_i) \right)$$

adds the weights w of all shared attributes i, and ψ is a function normalizing the output score to the range [0,1]. Clearly we could make further refinements, for example by penalizing non-matching attributes or by biasing toward supra-threshold attributes, but the determination of similarity is dependent on *what similarity means*—it could be context-dependent or include asymmetric relations between objects (such as *chili* being more closely related to *red* than vice versa), to better capture the nuanced inter-dependencies of language.

Explicitly storing the results of similarity calculations between objects, thereby building an associative data structure, yields significant advantages: once determined, information of relatively low volatility—such as overt similarity—does not need to be recalculated. The strength of relationship between two objects can be immediately determined, and an item's nearest neighbours can be immediately retrieved, giving an amortization of searching that obeys the principle of *parsimony*. By contrast, one problem with vector-space representations is that all vectors must be scanned to find the closest match, resulting in greater time costs. Since information management often deals with large data sets, high dimensionality, and expensive similarity measures that would be impractical to recalculate in real time (Moreno-Seco et al., 2003), the storage of cumulative results in an associative knowledge structure can lead to significant time savings.

Navigation — *Objects and categories should be organized to allow efficient retrieval from a knowledge space.*

Since information retrieval in the broadest sense represents an *interaction* between user and information space—involving judgment, learning, and change of state of both user and space—the principle most associated with retrieval is called *navigation*. The interaction may involve only a single step, as in direct querying, or it may involve iterative searching based on the local and global attributes of the information space.

Direct Querying — The simplest form of retrieval is *direct querying*, which takes users to their final goal in one step by matching search criteria against a corpus to find the best matches. Direct querying can be performed as a spreading activation process, such as in the memory model of (Anderson, 1983) where highly activated items are retrieved as most relevant to the query, with the single most highly activated item being the target. Similarly, since objects in P-MAK are indexed by descriptive attributes, a search using those attributes will retrieve the most relevant objects. While direct queries are fast when the target is well-known, they may retrieve too many objects when search criteria are too vague. Conversely, an overly-specific query may return few or no results, and exclude some that may have been highly relevant. Direct queries will also definitely fail if the query's target attributes (or its synonyms) do not appear in the knowledge structure.

When direct querying fails, a more costly *search* process must be used. The amount of time required can be considerable if objects are not associated by some useful property such as similarity or co-occurrence. Moreover, if only partial information is available (in either objects or query), additional time-consuming reasoning processes may be necessary (Russell and Norvig, 1995). Searching is also inevitable when the size of an information space grows to exceed a critical size as documents are added, while the capacity of users to enlarge and refine their queries is ultimately constrained by the limits of human patience and attention. In the worst case, the system may need to check each item in turn.

A faster, more effective form of *navigation* through an information space is possible given two complementary sources: (i) local cues (or *signposts*), based on attributes visible in an object's most similar "neighbour" objects, and (ii) global overviews (or maps), based on a summary description of a large knowledge structure. These search strategies are also found in human wayfinding, both in determining the next move with respect to immediate surroundings, and in reviewing available resources from a situating overview such as a map. Seeking behaviours seem to come naturally to humans, possibly because they are similar to the foraging and hunting behaviours of our ancestors (Sharps et al., 2002).

Local Search with Signposts — In human wayfinding, local *sign-following* is based on nearby cues in the environment, and occurs quickly since these cues are usually present (Hunt and Waller, 1999). Cognitive models such as (Anderson, 1983; Raaijmakers and Shiffrin, 1981) posit that recovering a memory begins with a best guess that retrieves a set of items. The best of these items is activated to retrieve its cohorts, and then an item selected from among those cohorts is activated, repeating until a satisfactory item is found or the search is abandoned. Similarly, users of an information-management system are likely to pick a next step that appears to bring them closer to their intended target, based on attributes visible in currently-retrieved objects (Teevan et al., 2004). During the search, a document's nearest neighbours will share a significant number of keywords, but will have some distinguishing keywords closer to what is being sought—the likelihood of a successful search is improved if local attribute "signage" is adequately complete and non-ambiguous. Such iterative searching approaches items of interest, in a semantic gradient descent that quickly narrows the search for relevant items. Ideally, the system will assist in suggesting next steps, such that items are retrieved based on their relation to previous user choices.

Such semantic traversal is more efficient if objects are not too highly inter-related, and show the necessary *small-world* property of many short-range connections and a few long-range ones (Milgram, 1967; Watts and Strogatz, 1998). Small-world systems (typically represented as networks) tend to have diameters exponentially smaller than their size, so that on average a short path may be found between any two items. For example, when the Web contained a billion pages, only 18.59 clicks on average separated any two (Albert et al., 1999). The organizing principle of the Web (its "classifier") involves authors linking their pages to others based on subjective similarity judgments. This property of local self-organization has positive implications for improved navigability in other semantic information spaces, such as the sort of automatically-constructed similarity-based knowledge structures described by P-MAK.

<u>Global Search with Maps</u> — Global *map-following* occurs where users coordinate their searches using high-level summary overviews that act as guides (or substitutes) to exploration (Hunt and Waller, 1999). By itself, local sign-following does not reveal the big picture, and with too many choices at every step users can quickly become overwhelmed and disoriented. To judge quickly the relevance of a corpus to one's information needs, a summary of its offerings would be useful.

Users unfamiliar with the contents of a corpus will have difficulty developing retrieval strategies since they will not know what comprises a good cue. They will also have difficulty in judging the relevance of what is retrieved relative to other items in the system. The principle of *abstraction* is applicable here, by providing a high-level summary that exposes points of interest and imposes some inspectable order upon them, permitting users to jump to dominant items and use local sign-following from there. Users can then ascertain how the information space may be useful to them, and orient themselves wherever they happen to be within it.

To create such a *<u>semantic index</u>*, the attributes of the corpus's objects can be used to build a tree of descriptors, with the most common at the top, connected to major sub-descriptors, each of which are connected to their sub-descriptors, etc. down to the leaf level of individual objects (cf. Koller and Sahami, 1997).

Objects that are general enough in content can appear at higher levels of the tree. The resulting map exposes the main topic, as well as describing subtopics at increasing levels of detail, in a general-to-specific containment hierarchy of categories—such an organized summary of the semantics, in human-readable form, offers a comprehensible mental model that improves navigation through the information space (Hunt and Waller, 1999; Rainsford and Roddick, 1999).

Because they highlight *configural* information, maps are in some ways superior to exploration based on local cues. First, they show the topic categories available in an information base, how they are related (by semantic distance, or by intervening topics), the relative importance or amount of material in certain topics, the biases of over- or under-represented topics, and how the knowledge structure has been organized. Second, maps support orientation by showing users the best entry points from which to start local sign-following, where the users are currently situated within the knowledge structure, their semantic "bearing" (the topics that they appear to be approaching or leaving), and wayfinding clues such as well-used pathways, neighbourhoods, and landmarks.

Local signposts and global maps are complementary; humans use either according to need and preference. Information management systems should purposely exploit them as well, since systems become more useful if they store information in a way that promotes quick contextual retrieval, and their interfaces become more useful if they include human-centred searching, browsing, and indexing facilities (Bertel et al., 2004).

5.2 *Situational Principles*: Capturing Context

Situational principles describe how an information management system should interact with the world—and the users that inhabit it—by encoding the temporal and environmental *context* that humans find important. Since computers are not inherently sensitive to such properties, the most pragmatic approach is to examine how humans incorporate the statistical regularities of their experiences into their memory structures. Human memory is by definition time-based and dynamic, requiring the formation of new memory traces, the fading of less-important traces, and the formation of new associations based on discovered similarities. Cognitively-inspired information systems should therefore take into account the observed *persistence* of items. The dynamics of persistence is applicable, with the addition of modality-specific attribute types, to the temporal and sensorial coding of the regularities of events. The situational principle of *temporal cueing* encodes temporal patterns of observed behaviours which can then be used to trigger reminders, while *sen-*

sorial cueing triggers reactions based on real-world stimuli captured by sensors. An *event* is then defined as the context of interaction with an information object that encodes combinations of temporal and sensory data in an *event convergence*.

At a functional level, the dominant cognitive theory of episodic memory is *encoding specificity* (Tulving and Thomson, 1973), in which memories of events are better retrieved if the cues that were present as an event was encoded are also present at retrieval. This description has formed the basis of several computational memory models, such as the *Search of Associative Memory* (SAM) (Raaijmakers and Shiffrin, 1981): given a set of attributes as cues, both semantic memory images and episodic memory of events and can be retrieved.

Although *instance theories* of memory assume that *each* encounter with an object is stored individually (Medin and Schaffer, 1978; Hintzman, 1984; Nosofsky, 1984), and that cues are matched against *all* of these images in real time for a best fit, such a large and detailed amount of information would overwhelm an information management system. Instead some form of summarization is necessary. With reference to human memory models, the approach should be more *constructive* (e.g. Kintsch, 1974) than *reconstructive* (e.g. Loftus and Palmer, 1974): storing trends that emerge from the details, but not the details themselves. Human memory tends to integrate the details of experience, particularly those of repeated, similar experiences, as efficiently-stored summary heuristics for rapid reasoning Bransford and Franks (1971).

Once these patterns are encoded, they can be used to trigger *prospective* episodic memory, which cues recall of intended tasks at particular times and under particular circumstances. Three types of prospective memory have been identified (Einstein and McDaniel, 1990): time-based, which refers to actions that will occur at a particular time of day, such as taking medicine at 0900h; event-based, such as returning an item the next time you see its owner; and activity-based, such as remembering to return a book the next time you go to the library. Event-based and activity-based prospective memory are relatively straightforward: if the appropriate contextual cues are strongly enough stimulated, then memories of appropriate actions will gain activation and rise to consciousness. Time-based prospective memory is more problematic: although episodic memory models describe the retrieval of specific snapshots, the explicit coding of temporal information in episodic memory has been little studied in cognitive science (Tranel and Jones, 2006). However, the contexts of time and space are important for human-centered information retrieval. The situational principles capture this context as follows:

***Persistence** — human memory emphasizes events that are recently, consistently, and concurrently experienced.*

Items that appear regularly, those that show some degree of *persistence*, represent the key factors of our environment. Persistence is closely bound to the phenomenon of human forgetting: items that rarely appear in the environment are considered less useful, and memories for them are less highly activated, than items experienced more recently and regularly (Anderson and Schooler, 1991). A bias for items that regularly co-occur is also useful: *Hebbian learning* describes the cognitive process whereby two items that appear together consistently become associated regardless of their similarities; retrieving one then automatically retrieves the other (Hebb, 1949). Persistence appears to reflect a deeply ingrained aspect of human intelligence: even the extremely young are good at learning the statistical regularities of their environment, even when events are not semantically related (Munakata, 2004).

An information management system that models persistence should similarly encode such regularities. Persistence records the degree to which objects have been used, and used together. Encoding the prominence of an individual object is achieved simply by assigning the object a weight; the greater the value, the more that object is used. As the object is ignored in preference for other objects, its weight falls. Thus over time, persistence weights indicate the most useful and important objects in a corpus, as some must rise in activation over others that are "pushed further back" as their usage diminishes. The co-occurrence of objects is similarly encoded: when two items are used together, the system should infer provisionally that there is some meaningful, unseen relation between them, even if they are not otherwise related. A weight is assigned to the pair of co-occurring objects, and as objects continue to be used at the same time, the weight between them strengthens, until eventually as one is retrieved the other will be also. However as any one object is used without its peers, the association between them weakens, and an association may disappear altogether if the correlation that it represents is not repeated.

Co-occurrence and similarity are fundamentally different. Two objects will remain similar to the extent that they share attributes. How objects are *used* is something more fluid: there is a categorical difference between items related by overt similarities (such as cats and dogs), and items related by co-occurrence (such as leashes and dogs). Similarity relations allow the retrieval of declarative knowledge, while co-occurrence relations reflect events experienced over time—how objects have been visited and combined, and what trails have been followed through an information space. Objects clustered based on usage indicate a context of activity, as per the ad hoc categories of (Barsalou, 1983). Thus P-MAK uses two fundamental types of association: those that represent the (relatively static) similarity of objects, and those that represent the dynamic concurrence of objects and cues.

The principle of persistence is the basis for all adaptive learning in the P-MAK framework: *temporal* and *sensorial* cueing use the same kind of dynamic association weights to represent how objects persist at—in fact *co-occur* with—particular times and stimuli.

Temporal Cueing *— humans encode the temporal regularities of past experiences.*

Temporal cueing describes the human tendency to retain images in episodic memory that encode when events experienced first-hand have typically occurred, for instance "every morning", or "at the solstice". Episodic memory often includes defining details and context, along with a subjective impression of when the memory was formed (such as minutes, days, or years ago). Human awareness of time is built-in at the cellular level: the *suprachiasmatic nucleus* is a cell assembly at the base of the hypothalamus comprising some 10,000 neurons; it triggers the daily secretion of melatonin that induces sleep (Yamaguchi et al., 2003). This sensitivity to time—common among living things—appears to be based on the natural periods to which humans are exposed, such as heartbeats, daylight, lunar phases, and seasonal variations. New temporal information is encoded by adjusting neuronal time delays to particular temporal combinations of inputs (Cariani, 2001); the *neocortical microcircuit* appears to maintain a virtual continuum of timescales, with time constants ranging from milliseconds to years (Denham and Tarassenko, 2003).

Encoding specificity can be used to infer the *recurrence* of events in temporal terms. Memories of one-time, unremarkable incidents are clearly less useful than memory for events that have repeated, ongoing significance in our lives. Such temporal trend-based encoding is enormously useful in information management: it extends retrieval beyond basic similarity judgments to enable users to ask such questions as, "when do I usually do X?", or "what do I regularly do at time T?" Although this seems like a useful function, we know of no information management systems purposely designed to answer these sorts of questions.

Such time-based indexing of events requires the definition of *temporal attributes*. These should not be confused with time-related terms used as semantic keywords. For example, a document that contains repeated references to a particular temporal epoch (e.g. *Monday*) may include it as an attribute, but this says nothing about how the document itself has been used. Rather, an implementation of temporal cueing must provide *a* *priori* a scale of permanent, weighted time-unit attributes (or rather *cues*) appropriate to the expected time scales that the system will encounter: for example the minutes, hours, days, etc. of the calendar.

Temporal cues representing particular times become associated with information objects if the objects are used at those times. Just as with *persistence*, the cues gain activation the more their corresponding information objects are used, and the association weights between each object and its cues grow stronger. The cues, and their associations with objects, will weaken if the objects are not used at the expected times. The strongest cues then represent the times that see the most activity, and the strongest associations indicate which objects are used most at which times.

***Sensorial Cueing** — humans are sensitive to their surroundings, and react appropriately in the presence of contextually relevant stimuli.*

The use of sensors is clearly related to the epistemic principle of *perception*. Perception can occur both abstractly, such as when "sensing" the word content of a document, and physically in quantifying motion, temperature, illumination, etc. In this case however, the "classifier" is the array of sensors connected to the system. Sensors act as inputs to a probabilistic decision process. Given a particular set of stimuli, each sensor determines to what degree its conditions have been met; its activation is then proportional to the strength of the stimulus. Sensor correlations are learned by associating objects and sensors if the objects are used while the sensors are active. With an appropriate combination of sensors an information system can learn correlations in the environment, enabling queries such as, "under what conditions are object O used?" and "which objects are used under condition C?"

As with temporal cueing, sensorial cueing is based on the principle of *persistence* in its use of dynamic weights and associations. The more a sensor is stimulated, the stronger its weight becomes. The more an object is used while a sensor is stimulated, the stronger the association between them becomes. A sensor's weight diminishes if it is idle while other sensors are active, and if an object is used while an associated sensor is idle, the association between them weakens and may eventually disappear. The strongest sensors represent the dominant actions in the environment, and the strongest associations indicate objects that are used most in particular contexts.

Sensorial cueing is critically dependent on the choice of sensors, which may be too specific or lack sensitivity. If important properties can be identified in advance, a sensor with perfect *alignment* can be designed to pick out the desired attribute cheaply and without fail. Perfect alignment requires no intelligence: it will respond immediately to a stimulus and trigger a reflex (as per Skinner, 1977). An inappropriately chosen or poorly-aligned set of sensors may lead to perceptual biases that result in *concept blindness*, where important themes go undetected due to gaps in the sensor array. A system may diagnose such problems, for example by detecting events that are registered in the absence of a consistent stimulus pattern, and may then assume that its domain has been too narrowly specified. If greater generality of domain is assumed, then more *intelligence* will be required to learn a pattern of activations across a set of noisy sensors to uncover new relationships and to maintain reliably correct classifications, and the necessary computations may not be bounded by time. *Expert systems* reduce this burden by functioning in a well-described restricted domain, where efficient algorithms can be written given prior knowledge and assumptions about the data (Thornton, 2000).

Event Convergence *— humans experience an event as a discrete entity comprised of a dynamic set of cues.*

An *event* represents an occurrence in the world, specifically an encounter with information objects at the specified time(s) and under the specified condition(s). Although temporal and sensorial cues may be associated individually and directly with an object, when more than one cue is used to describe a compound context they must be represented as a *conjunctive set* to avoid ambiguity. An "event" is precisely this conjunctive set of some combination of both temporal and sensorial cues. A single event representing multiple cues may then be associated with one or more information objects that occur in that context, and the relative typicality of the pattern's components is reflected in their individual weights. For example, both weather-report and bus-schedule objects may be consulted according to the conjunctive set "every Tuesday and Thursday when it's cloudy."

This we call the ***cue-event-object*** (CEO) model: temporal and sensorial cues are combined in an event that mimics cognitive *convergence zones*, which similarly synchronize perceptions and concepts (Moll et al., 1994; Amedi et al., 2005). We represent those concepts in the P-MAK framework as information objects, and event convergence is clearly related to *associationism* in its dynamic relation of disparate elements, and also *abstraction* in representing dynamic situations as single events, and *quantization* in summarizing an entity from a set of cues.

Concept drift can then be easily simulated with the CEO model in the associations between each cue and the event, and between the event and its objects. These associations strengthen or weaken as objects conform to or deviate from the event's pattern. An event's conjunctive set can fragment as necessary: if a user's schedule changes, then the event associated with the bus-schedule and weather objects can migrate to just "every Tuesday when it's cloudy". The determining factor is the strength of the event's associations with its component time units. An event may be represented as *rare* by its weak association strengths, but it will remain stable as long as all its cue associations are uniformly supported within some tolerance.

Reminders of upcoming, contextually relevant events are modeled on human prospective episodic memory, based on the notion that a structure's past usage can predict its current usage Anderson (1989); Anderson and Schooler (1991). Prospective memory has two principal components: *cue identification* that recognizes the appropriate context, and *intention retrieval* of the appropriate reaction to the cues Simons et al. (2006). In our model, cue identification is performed by stimulating the timers and sensors of current conditions, and if all the components associated with an event are active, intention retrieval is performed by spreading activation to objects associated with the event. If the retrieved objects are used following retrieval, the associations become stronger, otherwise they weaken. Thus objects are retrieved at appropriate times by temporal cueing, mimicking time-based prospective memory. Objects are also retrieved under appropriate conditions by sensorial cueing, mimicking event- and action-based prospective memory. In the sense that objects that have been previously used in particular contexts can be retrieved automatically when the same conditions recur, the CEO model is similar to models of stimulus-response learning (e.g. Skinner, 1977). In the sense that if certain preconditions are observed, then certain responses must follow, the model's timers and sensors also act as inputs to production rules represented as events (cf. Taatgen et al., 2006).

Although rare, information management systems that use temporal and sensorial cues as triggers include the experimental CybreMinder (Dey and Abowd, 2000) that generates reminders when user-specified temporal and situational conditions are satisfied, and the wearable Forget-Me-Not (Lamming and Flynn, 1994) that records a user's wanderings and interactions for later analysis.

6 Associative Network Representation

We have referred repeatedly to "associative knowledge structures" without committing to any particular representation. Here we commit to a graph-based network design; we believe this to be the best representation, although others are possible. For example, relations between objects could be represented in a matrix, with similarity scores in the cells indexed by object identifier labels. This representation is fine if the pattern of relationships is dense within a corpus: a large proportion of cells will then be filled. But there is evidence that *semantic* knowledge structures tend to be sparse, and may generally exhibit small-world properties (Barabási, 2002; Steyvers and Tenenbaum, 2005), as found in patterns of inter-word relations within languages (Motter et al., 2002; Sigman and Cecchi, 2002). In this regard, matrices do not appear to be the best choice for representation since resources would be wasted on unused cells, violating the principle of *parsimony*. Since matrix-based representations cannot accommodate new relation types without the addition of new tables, their inflexibility and rapid growth also violate *plasticity* and *scalability*.

6.1 The Advantage of Networks

Networks on the other hand seem ideal in many ways. They are the most parsimonious associative knowledge structure, since they only use resources to represent what is actually there. [6] Using the rules of graph theory (see e.g. Harary, 1969) networks are simple to define: each object appears as a node, with nodes linked by edges if they are related. Relationships of different kinds may be represented by typed links. Networks are easy to reconfigure and may even be nested inside the nodes of other networks to create more complex structures, such as compound concepts and dynamically organized categories. Networks are also inherently graphical: they are straightforward to draw, and people often use network diagrams to clarify and communicate their ideas. Humans are good at wayfinding, and since networks are navigable from node to node along links, they encourage explorative browsing to supplement more typical query-based searches. As we've seen, networks are also popular in associative models of human knowledge and memory. [7]

Although networks are *de facto* the most flexible means of organizing data, they are also potentially costly since any node may be directly connected to any other: a fully connected network of n nodes would have $\left(n^2 - n\right)/2$ links connecting them—implying a quadratic increase in links as new nodes are added. However if a domain exhibits a small-world distribution in the associations between its objects, an explosion of links will be averted by emphasizing local connections. There are other possible benefits. A network's small-world parameters could be monitored to maintain its navigability (Kleinberg, 2000); a skewed distribution of similarity links would indicate a poorly-tuned classifier for a given corpus, since the classifier extracts the attributes that determine connections. Maintaining a network's small-world property would be important for avoiding the *phase transitions* that occur as node connectivity increases, causing non-optimal changes in behaviour such as a dramatic expansion of the number of nodes activated by the spreading-activation "event horizon" (Shrager et al., 1987). But despite these potential advantages, information management systems do not test for or exploit small-world properties in their domains of application (Perugini et al., 2004).

A network-type topology offers further benefits through graph-theory analyses. Highly-connected *authority* (or *hub*) nodes indicate semantic trends by their many links. In a cognitive sense authorities are like concepts that come readily to mind, since they are related to so many things, while from the point of view of graph theory, the distance from authorities to other members of their group is on average a minimum (Adamic, 1999). As with landmarks in human wayfinding, authority-nodes may act as meaningful entry points to the network.

Meanwhile, clusters of nodes that are closely connected represent semantic trends; they can be organized into a *clique graph* (Harary, 1969) that summarizes clusters into semantic "neighbourhoods" connected by "highways". *Cut-points* in the clique graph indicate natural divisions between larger sections of the network. Topology can also serve as a diagnostic metric: for instance a large number of unconnected nodes could also indicate a poorly-tuned or inappropriate classifier. Conversely, if too many nodes are clustered by the same set of attributes, the cluster could be subdivided into smaller distinct groups by increasing the sensitivity of the classifier until additional differentiating attributes emerge. This is equivalent to the mental processes that provide us with a sense of *distinctiveness*: the more people know about something, the less it seems like other things (e.g. Rabinowitz and Andrews, 1973).

6.2 Basic Network Elements

In the P-MAK framework, different types of nodes and links are used to represent different types of elements. There are two types of nodes:

- **semantic nodes** represent the meaningful entities of the world: objects, ideas, and events. They are defined by the descriptive attributes that they contain :
  - ▷ **object nodes** represent an individual data object, such as a document or image.
  - ▷ **event nodes** represent events in the environment, such as a user's behaviour or the stimulation of environmental sensors.

    - ▷ **action nodes** are activated by the system to trigger an effect in the environment, such as sounding a notice or dispensing medication.

- **index nodes** encode the contexts in which entities described by semantic nodes occur. In particular :
    - ▷ **temporal nodes** index the particular times at which events occur, and contain only a time-stamp.
    - ▷ **sensor nodes** index the observable circumstances under which events occur.
    - ▷ **conjunctive nodes** combine two or more temporal or sensor nodes into a more specific context representation.

Although research into semantic networks has identified many different link types, mostly with respect to linguistic relationships between words (e.g. Woods, 1975), for fundamental information management it is enough to begin with bi-directional edges that represent strength of *relatedness*. For the basic operations of information management, two types of relatedness are necessary:

- **similarity links** represent the degree to which two nodes share attributes. The more attributes they share, the more they represent similar things.
- **usage links** represent the degree to which two nodes are activated at the same time.

With these node and link types, we are able to represent knowledge in several useful ways: by the similarity between objects, by the objects with which they are typically used, and by the temporal and environmental context in which objects occur.

6.3 Networks for Similarity, Usage, and Situations

Three network types conform to the basic operations of information management and P-MAK's principles, using the node and link elements just described. The *similarity network* is used to search for items based on their content; as items are retrieved, similar items can be found by navigating to them along connected links. Apart from semantics, items can also be retrieved from a *usage network* that links items to the degree that they co-occur, and from a *situational network* that indexes items by when and under what conditions they are activated.

Similarity Networks — As the epistemic principles are used to induct a corpus, objects (such as documents) are represented by ***object nodes*** created according to the principle of *perception*, such that each node contains

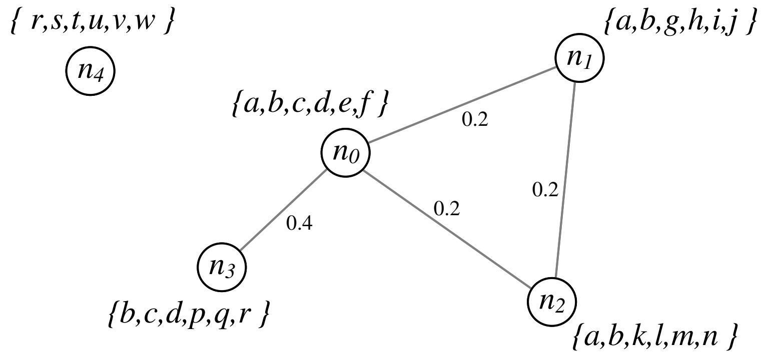

Figure 3: An example of a simple similarity network. Each node is defined by discrete attributes, here represented as letters in lists. Nodes are linked to the degree that they share attributes. Thus object nodes n_0, n_1, and n_2 are all equally connected for sharing the same two attributes; n_0 and n_3 are more strongly connected for sharing three. n_4 is unconnected—it shares no attributes with any node but n_3, and in this example one shared attribute that they share is not enough to overcome the similarity function's weight threshold; such tuning of a similarity measure can prevent nodes from being linked when their relationship is too weak, and be used to preserve large semantic network's small-world property.

a list of attributes that describe its semantics. [8] A *similarity network* is formed as nodes are connected by weighted ***similarity links*** —according to the *similarity* principle nodes are more strongly linked the more that they share attributes and thus represent similar things. The calculated similarity value can be used as the link's weight; pairs of nodes with low similarity scores would not be linked. The network thus organizes and stores the results of the classification process in the pattern of connections between nodes (Figure 3); relations do not then have to be recalculated in real time.

There is strong evidence from brain scans (Habib et al., 2003) that semantic and episodic memory are largely processed in different parts of the brain. It therefore seems reasonable to disentangle episodic and similarity information into separate representations that reference the semantic object nodes separately.

Usage Networks — As items represented by existing object nodes are used concurrently and consistently over time, a new ***usage link*** can be created to connect them, based on the principle of *persistence*. The more the two objects occur together, the stronger the link grows, otherwise it decays and disappears. Once linked, spreading activation can be used to activate nodes and retrieve objects that tend to co-occur, whether or not they are semantically similar. The combination of all usage links forms a *usage network*.

To adhere to the principles of *parsimony* and *scalability*, activation values used for information management should be finite and bounded, so that as the activation of some links and nodes in the usage network rises, normalization limits these elements and decays all others in the network, in a process analogous to human forgetting. For an information system this is particularly useful: nodes with a long-term activation approaching zero have a low apparent utility. Depending on the information-management protocol, these nodes and the entity that they represent could be deleted, or the nodes retained but the entity moved to long-term storage, freeing up real-time resources for more important, recently-activated items.

Similarity and co-occurrence links are grouped into separate networks to avoid entanglement: usage should not affect similarity valuations, since the same object may be used in different ways, in different contexts, and by different people (nonetheless, the semantic and usage networks can work together to recover, say, all similar items that also tend to be used together). While the weights of usage links change with every user interaction, similarity links are virtually static: their weights are not recalculated unless the classifier's similarity metric changes, or unless more detailed knowledge is gained about the objects.

Situational Networks — Situational networks simulate human episodic memory by describing the times and conditions where objects are used. Parallel episodic memory models such as (Medin and Schaffer, 1978; Hintzman, 1984; Nosofsky, 1984; Raaijmakers and Shiffrin, 1981) would be impractical to implement on a serial machine, violating the mechanistic principle of *parsimony*; this motivates our use of networks to store relationships that would otherwise be expensive to recompute. Situational networks encode the occurrence of events using both temporal and sensorial index nodes representing timers and sensors external to the system. Once an event's temporal and environmental patterns have been encoded, the network structure can be queried, reminders can be generated, and actuators can be triggered.

An ***event node*** represents a real-world occurrence that involves an information object, such as a document retrieval, and may contain attributes that describe the occurrence. Systems are programmed *a priori* to respond to a finite number of input types, each of which can be described by attributes, and thus systems can assign attributes to events as they occur; users could edit these and also add their own 'tags' as attributes. Events can then be clustered by their similarities. An event node is created the first time a particular event type co-occurs with a given information object. Following the CEO model, the index nodes of timers and

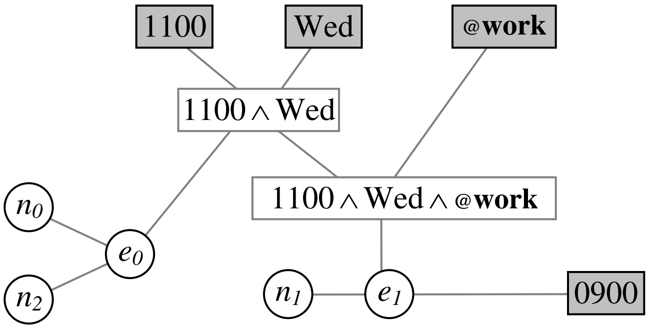

Figure 4: An example of the *cue-event-object* model implemented as a network. Temporal and sensorial index nodes typify observed events. Grey rectangles represent time and sensor nodes; white rectangles are conjunctive nodes. In this example, event node e_0 is associated with usage on Wednesdays at 1100h, when objects n_0 and n_2 are used. Object n_1 is also used on Wednesdays at 1100h, but *only* when the at-work sensor is active; n_1 is also used every day at 0900h, regardless of location. The network can be used to retrieve objects automatically by spreading activation from index nodes when their temporal and sensorial conditions are repeated.

sensors that correspond to the event are linked to a new conjunctive node that is then linked to the event node. The event node is then linked to the information-object node.

As the event re-occurs at other times, more conjunctive nodes are added to describe unique circumstances. Where the units of two conjunctive nodes agree, those conjunctive nodes may be ***aggregated*** to produce a more compact representation. For example, an event that occurs every day at 0900h can be reduced to a single time unit, and connected directly to the temporal node that represents 0900h. A conjunctive node's well-formed-formulas can also include disjunctions such as $0900h \wedge \{Mon \vee Wed\}$ to represent partial adherence to a particular scale, in this case the days of the week. As events re-occur at the time specified by a pattern, the weights of the nodes involved increases asymptotically, as do the weights of the links that connect them. Thus the weights indicate the frequency with which a pattern as a whole and each of its components are true. Events must receive support in order to persist, otherwise they are "forgotten": if an event does not re-occur as expected, then the weights of its corresponding elements decay.

If an event pattern becomes only partially supported, then a conjunctive nodes is ***disaggregated*** as the entropy of its link weights exceeds a threshold, fragmenting a single pattern into two or more with different levels of support. Together, the processes of aggregation and disaggregation are essential for modeling

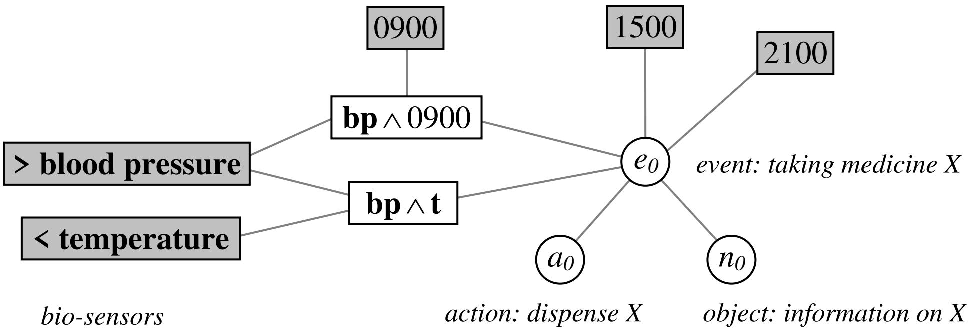

Figure 5: Index nodes for triggering an event node. In this case, medicine will not only be dispensed at given times, but also if biomedical sensors detect a potential problem. If the user's blood pressure rises above a certain level and their temperature falls, dispensation is triggered strongly-connected links. Sensor and temporal index nodes can also be combined: dispensation will also occur if blood pressure is high at 0900. The event node can also be linked to the object nodes of documents that describe the medicine's properties and dosage. Each of the nodes can be queried to determine its role with respect to other nodes, following connected links to ask such questions as what happens at 0900h, when medicines should be taken, and which medicines are available for dispensation.

concept drift. Like human memory, the situation network encodes unique events, but also integrates similar events into a gist-like summary.

Once event patterns have been encoded into the situation network, they can be queried to determine when events tend to occur, what combinations of sensors tend to be engaged, and what times are busiest, by examining the stronger weights in the network. For specific queries, activation flows from objects through event nodes to index nodes to find out *when* they occur, and activation flows from time or sensor nodes through conjunctive nodes (if any) and event nodes to determine *what* occurred.

Queries can also be automated to act as reminders: as time moves forward, index nodes corresponding to the current time and stimuli are activated, and the activation spreads to connected conjunctive nodes and events. Conjunctive nodes whose input links (including any sensors) are all activated will pass activation on to connected events. The events activate any objects that correspond, retrieving them in the appropriate context. Support is then increased for objects that are used after recall, while it decays for those that are ignored. Users may set their own reminders by connecting events to desired conditions. Here high link weights act as alarms by guaranteeing forceful retrieval, and users can also program the system to produce real-world effects by setting their own patterns and connecting them to ***action nodes*** that trigger actuators (Figure 5).

Although we are unaware of previous work that uses associative networks and spreading activation for temporal encoding, symbolic networks have been used in some models that include sensory processes. The *perceptual maps* of ***Convergence-Zone Episodic Memory*** (Moll et al., 1994), use a handful of perceptual stimuli to retrieve full episodic memories, and the ***Memory Extender*** (Jones, 1986) uses "context" nodes that serve a sensory role by biasing node activation according to changing circumstances. Adding sensors to such models would allow the *embodiment* of associative networks in a machine, with potential applications in robotics and contextual computing.

7 Conclusion and Future Work

P-MAK describes a framework at the intersection of cognitive science and informatics that fulfills the basic operations of information management—specifically the collection, identification, classification, and retrieval of meaningful objects. These processes are analogous to cognitive memory faculties of learning and recall, while for information systems these processes are the foundations of machine learning and information retrieval. Using cognitive properties in system design provides users with increased familiarity of function, which leads to improved usability. The assertion is that the functional processes of existing highly-refined and powerful biological systems—such as the brain—may suggest new and more efficient models of human-machine interaction.

Based on the constraints and properties of memory and machines, and on the basic operations of information management, P-MAK inducts information as discrete objects that are described by discrete attributes, the objects are associated to the degree that they share attributes, and their use is recorded as *events* indexed by context. Objects are thus described both by *internal* attributes and *external* cues. Although P-MAK's principles are each already well-known, their assembly into an information-management framework is novel.

An associative knowledge structure is most simply and flexibly represented as an associative semantic network. Networks are a cognitively plausible representation, forming the basis of popular cognitive memory models; they are also well-understood with respect to exploitable topological properties—such as clustering, hubs, and the small world property—that appear in various real-world semantic domains. A network representation is also compatible with neurophysiological models that equate neural cell assemblies with semantic nodes, and neural synchronization pathways with links that represent strength of similarity.

P-MAK's symbolic associative model avoids the computational complexity of a fully parallel cognitive model by storing similarity evaluations as links for rapid retrieval. Once similarity associations are stored, an object's nearest neighbours can be immediately retrieved; all retrieval in P-MAK occurs using a constrained spreading-activation model. P-MAK's network can in principle accommodate other association types; as an illustrative example a second link type is introduced that encodes Hebbian-type learning due to usage. Objects may also be indexed by *temporal* and *sensorial* index nodes that register when an object is used and under what circumstances. Objects can then be retrieved automatically at a particular time, or in the presence of particular stimuli.

Together, this defines a set of principles for information management that responds in a plausible, user-comprehensible way to changing conditions, derived from well-known cognitive models and usefully applied to automated information management. P-MAK is widely applicable and domain-independent.

Implications

P-MAK's generality suggests that it can be extended into epistemology, dynamic behaviour, data visualization, and applications. Further ideas from cognitive science could be applied to *abstraction*—

- Any subset of interest, such as categories implicit in node clusters, could be represented by a generated *meta-node* (as per the *prototypes* of Posner and Keele, 1970) or by an existing centroid node (as per the *exemplars* of Brooks, 1978). More generally, an *ad hoc* context—such as personalized starting points for navigation, or the set of nodes retrieved by a query—could be saved for re-use in a *template* of relevant node and link activations (as per *schema theory*, e.g. Mandler, 1984).
- To make high-dimensional semantic spaces more comprehensible, the network could be summarized with the cognitive abbreviations found in human wayfinding, such as *landmarks*, *paths*, *districts* and *boundaries* (Vinson, 1999).
- P-MAK could be expanded to include rule-mining and rule inference, nodes that act as processors, and a wider assortment links as found in the semantic networks appropriate to semantic networks and expert systems (see Sowa, 1991); this implies in due course a graph-based programming language.

Several possibilities relate to learning. Based on the context of a user's interests and actions —

- Semantic link strengths could vary dynamically, reflecting the notion that meaning is not abstractly fixed, but contextually dependent on the link pattern among a number of features (Lakoff, 1987).
- P-MAK's perceptual classifier could adapt to the composition of a corpus by increasing its sensitivity within subtopics; it could similarly adjust to the drifting bias of a corpus as items are added.
- Rates of link and node decay could vary to reflect greater utility and persistence.
- To index popular subjects, a dynamic *usage map* could be generated as a minimum spanning tree of often-used nodes and link pathways (cf. the *semantic index* of the *retrieval* principle).
- To extend P-MAK from personal to group information management (GIM) and Computer-Supported Collaborative Work (CSCW), users could be matched by similar interests to share usage maps; this points to possibilities in collaborative filtering and resource administration.

Some improvements are fundamental —

- An explicit *activation strategy* could define—for different circumstances—how activation spreads between nodes, how its strength is calculated, how far it should reach, and how *forgetting* is modeled. (see e.g. Crestani, 1997).
- To refine semantics, attributes could be represented as nodes related through networks of their own, both within individual nodes and across an entire corpus.

P-MAK-based systems may be implemented on devices of various types. Its *temporal* and *sensorial cueing* are useful for context-sensitive portable devices and ubiquitous computing; for example with the inclusion of metabolic and affective sensors, it could support medical applications and critical tasks that demand concentrated user attention. The cognitive plausibility of P-MAK's various components suggests that it could be applied to cognitive modeling simulations as well as to information systems, as an empirically-based testbed. As a personal assistant, if run continuously a P-MAK-based system can reflect ongoing trends, and as such form the basis of a prosthetic cybernetic system, a true *memory extender*.

Notes

[1]However the archealogical record suggests a stoic state of mind, with little evidence of imaginative artistry. Apparently the lesser working memory of *Neanderthalensis* eventually allowed the incremental innovation of *Homo Sapiens* to pull ahead.

[2]The *mechanistic* principles also apply to the constraints of brain function, as suggested by Marr (1982), but such an analysis is separate from our focus on improved human-machine interaction.

[3]Information management is primarily concerned with *long-term* storage and retrieval. Thus we do not directly discuss short-term *working memory* (Baddeley and Hitch, 1974), which has more to do with transient attentional processes that assemble cues to process retrievals.

[4]As "an enlarged intimate supplement" to memory, Memex is the earliest system design to embody the crucial human-centred principles of *associationism* and *persistence*.

[5]A possible method of coding asymmetry of relation between words is mentioned briefly in (Landauer and Dumais, 1997), although not elaborated.

[6]Different knowledge structures are optimal for different types of representation. For instance, as Bayes nets provide a compact representation of joint probability distributions, so the network representation of P-MAK provides a compact representation of semantic and contextual relations.

[7]"Network thinking is poised to invade all domains of human activity and most fields of human inquiry. It is more than another helpful perspective or tool. Networks are by their very nature the fabric of most complex systems, and nodes and links deeply infuse all strategies aimed at approaching our interlocked universe." (Barabási 2002, p. 222)

[8]Attributes could themselves be represented as nodes, as in (Jones, 1986) and following on our use of index nodes to represent cues. However for illustrative purposes we use a simpler formulation of nodes with self-contained semantics here; a simple device such as an *inverted index* (Salton and McGill, 1983) can then be used to retrieve all nodes that contain a given attribute.

Acknowledgements

We are grateful to Joel Lanir and Heidi Lam for their thoughtful insight and helpful comments. This paper was supported in part by NSERC postgraduate scholarships PGS B - 267320 and IPS 2 - 268129.

References

Adamic, L. A. (1999). The small world web. In Abiteboul, S. and Vercoustre, A., editors, *Proceedings of the European Conference on Digital Libraries (ECDL99); Lecture Notes in Computer Science*, volume 1696, pages 443–452. Springer-Verlag.

Albert, R., Jeong, H., and Barabási, A.-L. (1999). Internet: Diameter of the world-wide web. *Nature*, 401(6749):130–131.

Amedi, A., von Kriegstein, K., van Atteveldt, N. M., Beauchamp, M. S., and Naumer, M. J. (2005). Functional imaging of human crossmodal identification and object recognition. *Experimental Brain Research*, 166(3-4):559–572.

Anderson, J. R. (1983). A spreading activation theory of memory. *Journal of Verbal Learning and Verbal Behavior*, 22(3):261–295.

Anderson, J. R. (1989). *A Rational Analysis of Human Memory*, chapter 11, pages 195–210. Varieties of Memory and Consciousness: Essays in Honor of Endel Tulving. Lawrence Erlbaum Associates.

Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebière, C., and Qin, Y. (2004). An integrated theory of the mind. *Psychological Review*, 111(4):1036–1060.

Anderson, J. R. and Bower, G. H. (1973). *Human Associative Memory*. V.H. Winston.

Anderson, J. R. and Schooler, L. J. (1991). Reflections of the environment in memory. *Psychological Science*, 2(6):396–408.

Autonomy (n.d.). Autonomy Corporation plc (LSE: AU.) . Corporate home page accessed on the World Wide Web; Retrieved August 31, 2006, from http://www.autonomy.com.

Baars, B. J. (1993). How does a serial, integrated and very limited stream of consciousness emerge from a nervous system that is mostly unconscious, distributed, parallel and of enormous capacity? *CIBA Foundation Symposium*, 174:282–290.

Baddeley, A. D. and Hitch, G. (1974). *Working Memory*, volume 8 of *The Psychology of Learning and Motivation*. Academic Press.

Baecker, R., Grudin, J., Buxton, W., and Greenberg, S. (1995). *Readings in Human-Computer Interaction, 2nd Ed.* Morgan Kaufmann Series in Interactive Technologies. Morgan Kaufmann.

Bahrick, H. P. (1984). Semantic memory content in permastore. *Journal of Experimental Psychology: General*, 113(1):1–29.

Barabási, A.-L. (2002). *Linked: The New Science of Networks*. Perseus Publishing, Cambridge, MA.

Barnard, K. and Forsyth, D. (2001). Learning the semantics of words and pictures. In *Proceedings of the Eighth IEEE International Conference on Computer Vision, ICCV 2001*, volume 2, pages 408–415.

Barsalou, L. W. (1983). Ad hoc categories. *Memory & Cognition*, 11(3):211–227.

Barsalou, L. W. and Sewell, D. R. (1984). *Constructing representation of categories from different points of view*. Emory Cognition Project Report No.2. Emory University Press.

Bertel, S., Obendorf, H., and Richter, K.-F. (2004). User-centered views and spatial concepts for navigation in information spaces. Technical report. SFB/TR 8 Spatial Cognition.

Blank, M. A. and Foss, D. J. (1978). Semantic facilitation and lexical access during sentence processing. *Memory & Cognition*, 6(6):644–652.

Bransford, J. D. and Franks, J. J. (1971). The abstraction of linguistic ideas. *Cognitive Psychology*, 2:331–350.

Bransford, J. D. and Johnson, M. K. (1973). *Consideration of some problems of comprehension*, volume 2 of *Visual Information Processing*, pages 331–350. New York: Academic Press.

Brewer, W. F. and Treyens, J. C. (1981). Role of schemata in memory for places. *Cognitive Psychology*, 13:207–230.

Brooks, L. R. (1978). *Nonanalytic concept formation and memory for instances*, pages 170–211. Cognition and Categorization. Lawrence Erlbaum Associates.

Burgess, C. and Lund, K. (2000). *The Dynamics of Meaning in Memory*, pages 117–156. Cognitive Dynamics: Conceptual and Representational Change in Humans and Machines. Lawrence Erlbaum Associates.

Bush, V. (1945). As we may think. *Atlantic Monthly*, 176(1):101–108.

Cariani, P. (2001). Symbols and dynamics in the brain. *Biosystems*, 60(1–3):59–83.

Carroll, J. B. and Whorf, B. L. (1956). *Language, Thought, and Reality: Selected Writings*. MIT Press.

Chomsky, N. A. (1965). *Aspects of the Theory of Syntax*. MIT Press.

Collins, A. M. and Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, 82(6):407–428.

Collins, A. M. and Quillian, M. R. (1969). Retrieval time from semantic memory. *Journal of Verbal Learning and Verbal Behavior*, 8(2):240–248.

Cowan, N. (2000). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, 24:87–185.

Crestani, F. (1997). Application of spreading activation techniques in information retrieval. *Artificial Intelligence Review*, 11:453–482.

Deerwester, S., Dumais, S. T., Furnas, G. W., Landauer, T. K., and Harshman, R. (1990). Indexing by latent semantic analysis. *Journal of the American Society for Information Science*, 41(6):391–407.

Denham, M. and Tarassenko, L. (2003). Sensory processing. Technical report. Foresight Cognitive Systems Project Research Review.

Dey, A. K. and Abowd, G. D. (2000). *CybreMinder: A Context-Aware System for Supporting Reminders*, volume 1927 of *Handheld and Ubiquitous Computing; Lecture Notes in Computer Science*, pages 172–186. Springer-Verlag.

Dumais, S. T., Cutrell, E., Cadiz, J. J., Jancke, G., Sarin, R., and Robbins, D. C. (2003). Stuff I've Seen: a system for personal information retrieval and re-use. In *SIGIR 2003: Proceedings of the 26th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, July 28 - August 1, 2003, Toronto, Canada*, pages 72–79. ACM.

Einstein, G. O. and McDaniel, M. A. (1990). Normal aging and prospective memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 16(4):717–726.

Ekman, P. (1971). Universals and cultural differences in facial expressions of emotion. In Cole, J., editor, *Nebraska Symposium on Motivation 1971*, volume 19, pages 207–284. University of Nebraska Press.

Fertig, S., Freeman, E., and Gelernter, D. (1996). Lifestreams: An alternative to the desktop metaphor. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI '96)*, pages 410–414. ACM Press.

Fodor, J. A. (1975). *The Language of Thought*. Crowell, New York.

Foltz, P. W. (1991). Models of human memory and computer information retrieval: Similar approaches to simiar problems. Technical report.

Frank, S. L., Koppen, M., Noordmana, L. G. M., and Vonk, W. (2003). Modeling knowledge-based inferences in story comprehension. *Cognitive Science*, 27:875–910.

Gemmell, J., Bell, G., Lueder, R., Drucker, S., and Wong, C. (2002). MyLifeBits: Fulfilling the Memex Vision. In *Proceedings of ACM Multimedia '02, December 1-6, 2002, Juan-les-Pins, France*, pages 235–238. ACM Press.

Gillund, G. and Shiffrin, R. M. (1984). A retrieval model for both recognition and recall. *Psychological Review*, 91:1–67.

Goertzel, B. (1997). *From Complexity to Creativity: Explorations in Evolutionary, Autopoietic, and Cognitive Dynamics*. IFSR International Series on Systems Science and Engineering. Plenum Press.

Habib, R., Nyberg, L., and Tulving, E. (2003). Hemispheric Asymmetries of Memory: The HERA Model Revisited. *Trends in Cognitive Sciences*, 7(8):241–245.

Harary, F. (1969). *Graph Theory*. Addison-Wesley.

Hebb, D. O. (1949). *The Organization of Behavior*. John Wiley.

Hintzman, D. L. (1984). Minerva 2: A simulation model of human memory. *Behavior Research Methods, Instruments, & Computers*, 16(2):96–101.

Hoffman, R. R., Klein, G., and Laughery, K. R. (2002). The state of cognitive systems engineering. *Intelligent Systems*, 17(1):73–75.

Hunt, E. and Waller, D. (1999). Orientation and wayfinding: A review. Technical Report N00014-96-0380, Arlington, VA. Office of Naval Research.

Huyck, C. R. (2001). *Cell Assemblies as an Intermediate Level Model of Cognition*, volume 2036 of *Emergent Neural Computational Architectures Based on Neuroscience: Towards Neuroscience-Inspired Computing*, pages 383–397. Springer-Verlag.

Jacoby, L. L. and Witherspoon, D. (1982). Remembering without awareness. *Canadian Journal of Psychology*, 36:300–324.

Johnson, T. R. (1997). Control in ACT-R and Soar. In Shafto, M. and Langley, P., editors, *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society*, pages 343–348. Lawrence Erlbaum Associates.

Jones, W. P. (1986). The Memory Extender Personal Filing System. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 298–305. ACM Press.

Kahneman, D. and Treisman, A. (1984). *Changing views of attention and automaticity*, pages 29–61. Varieties of Attention. Academic Press.

Kintsch, W. (1974). *The Representation of Meaning in Memory*. Halsted Press.

Kleinberg, J. M. (2000). Navigation in a small world. *Nature*, 406:845.

Koller, D. and Sahami, M. (1997). Hierarchically classifying documents using very few words. In *Proceedings of the 14th International Conference on Machine Learning (ML), Nashville, Tenessee, July 1997*, pages 170–178.

Labov, W. (1973). *The boundaries of words and their meaning*, volume 42 of *New Ways of Analyzing Variation in English*, pages 340–373. Georgetown Press.

Lakoff, G. (1987). *Women, Fire, and Dangerous Things: What Categories Reveal about the Mind*. University of Chicago Press.

Lamming, M. and Flynn, M. (1994). "Forget-Me-Not"—Intimate Computing in Support of Human Memory. In *Proceedings of FRIEND21 '94 International Symposium on Next Generation Human Interfaces*, pages 1–9. Rank Xerox Research Center.

Landauer, T. K. (2002). *On the computational basis of learning and cognition: Arguments from LSA*, volume 41 of *The Psychology of Learning and Motivation*, chapter 13, pages 43–84. Academic Press.

Landauer, T. K. and Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, 104(2):211–240.

Landauer, T. K., Laham, D., and Foltz, P. W. (1998). *Learning Human-like Knowledge by Singular Value Decomposition: A Progress Report*, chapter 10, pages 45–51. Advances in Neural Information Processing Systems. MIT Press.

Lemaire, B. and Denhière, G. (2004). Incremental construction of an associative network from a corpus. In Forbus, K., Gentner, D., and Regier, T., editors, *Proceedings of the 26th Annual Meeting of the Cognitive Science Society*, pages 825–830. Lawrence Erlbaum Associates.

Loftus, E. F. and Palmer, J. C. (1974). Reconstruction of automobile destruction: An example of the interaction between language and memory. *Journal of Verbal Learning and Verbal Behavior*, 13:585–589.

Mandler, J. M. (1984). *Stories, Scripts, and Scenes: Aspects of Schema Theory*. Lawrence Erlbaum Associates.

Marr, D. (1982). *Vision : A Computational Investigation into the Human Representation and Processing of Visual Information*. W.H. Freeman.

Matsuo, Y. and Ishizuka, M. (2004). Keyword extraction from a single document using word co-occurrence statistical information. *International Journal of Artificial Intelligence Tools*, 13(1):157–169.

McClelland, J. L. and Kawamoto, A. H. (1986). *Mechanisms of sentence processing: Assigning roles to constituents*, pages 318–362. Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Vol. 2: Psychological and Biological Models. MIT Press.

McRae, K., de Sa, V. R., and Seidenberg, M. S. (1997). On the nature and scope of featural representations of word meaning. *Journal of Experimental Psychology: General*, 126(3):99–130.

Medin, D. L. and Schaffer, M. M. (1978). Context theory of classification. *Psychological Review*, 85:207–238.

Meyer, D. E. and Schvaneveldt, R. W. (1971). Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology*, 90(2):227–234.

Milgram, S. (1967). The small world problem. *Psychology Today*, 1:60–67.

Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *The Psychological Review*, 63:81–97.

MindManager (n.d.). MindJet: Software for visualizing and managing information. Corporate home page accessed on the World Wide Web; Retrieved August 31, 2006, from http://www.mindjet.com.

Moll, M., Miikkulainen, R., and Abbey, J. (1994). The capacity of convergence-zone episodic memory. In *Proceedings of the 12th National Conference on Artificial Intelligence, AAAI-94*, pages 68–73. MIT Press.

Moravec, H. (1998). *ROBOT: Mere Machine to Transcendent Mind*. Oxford University Press.

Moreno-Seco, F., Micó, L., and Oncina, J. (2003). *Extending fast nearest neighbour search algorithms for approximate k-NN classification*, volume 2652 of *Pattern Recognition and Image Analysis; Lecture Notes in Computer Science*, pages 589–597. Springer-Verlag.

Motter, A. E., de Moura, A. P. S., Lai, Y.-C., and Dasgupta, P. (2002). Topology of the conceptual network of language. *Physical Review E*, 65(6):Art. No. 065102 Part 2.

Munakata, Y. (2004). Computational cognitive neuroscience of early memory development. *Developmental Review*, 24(1):133–153.

Nason, S. and Laird, J. E. (2005). Soar-RL: Integrating Reinforcement Learning with Soar. *Cognitive Systems Research*, 6(1):51–59.

Nosofsky, R. M. (1984). Choice, similarity, and the context theory of classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 10(1):104–114.

Osgood, C. E. (1952). The nature and measurement of meaning. *Psychological Bulletin*, 49(3):197–233.

Osgood, C. E., May, W., and Miron, M. (1975). *Cross-Cultural Universals of Affective Meaning*. University of Illinois Press.

Perugini, S., Gonçalves, M. A., and Fox, E. A. (2004). Recommender systems research: A connection-centric survey. *Journal of Intelligent Information Systems*, 23(2):107–143.

Posner, M. I. and Keele, S. W. (1970). Retention of abstract ideas. *Journal of Experimental Psychology*, 83:304–308.

Pulvermüller, F. (1999). Words in the brain's language. *Behavioral and Brain Sciences*, 22(2):253–336.

Quillian, M. R. (1969). The Teachable Language Comprehender: A simulation program and theory of language. *Communications of the ACM*, 12(8):459–476.

Raaijmakers, J. G. W. and Shiffrin, R. M. (1981). Search of associative memory. *Psychological Review*, 88:93–143.

Rabinowitz, F. M. and Andrews, S. S. R. (1973). Intentional and Incidental Learning in Children and the von Restorff Effect. *Journal of Experimental Psychology*, 100(2):315–318.

Rainsford, C. P. and Roddick, J. F. (1999). Database issues in knowledge discovery and data mining. *Australian Journal of Information Systems*, 6(2):101–128.

Rosch, E. and Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, 7:573–605.

Rumelhart, D. E., Hinton, G. E., and McClelland, J. L. (1986). *A general framework for parallel distributed processing*. Parallel Distributed Processing: Explorations in the Microstructure of Cognition. Vol. 1: Foundations. MIT Press.

Russell, S. and Norvig, P. (1995). *Artificial Intelligence: A Modern Approach*. Prentice Hall, Englewood Cliffs, NJ.

Salton, G. and McGill, M. (1983). *An Introduction to Modern Information Retrieval*. McGraw-Hill.

Schlögl, C. (2005). Information and knowledge management: dimensions and approaches. *Information Research*, 10(4):16pp.

Sharps, M. J., Villegas, A. B., Nunes, M. A., and Barber, T. L. (2002). Memory for animal tracks: A possible cognitive artifact of human evolution. *Journal of Psychology*, 136(5):469–492.

Shrager, J., Hogg, T., and Huberman, B. A. (1987). Observation of phase transitions in spreading activation networks. *Science*, 236(4805):1092–1094.

Sigman, M. and Cecchi, G. A. (2002). Global organization of the Wordnet lexicon. *Proceedings of the*

National Academy of Sciences, 99(3):1742–1747.

Simons, J. S., Schölvinck, M. L., Gilbert, S. J., Frith, C. D., and Burgess, P. W. (2006). Differential components of prospective memory? evidence from fMRI. *Neuropsychologia*, 44:1388–1397.

Skinner, B. F. (1977). Why I am not a cognitive psychologist. *Behaviorism*, 5:1–10.

Smith, B. C. (1996). *On the Origin of Objects*. MIT Press.

Smith, E. E., Shoben, E. J., and Rips, L. J. (1974). Structure and process in semantic memory: A featural model for semantic decisions. *Psychological Review*, 81:214–241.

Sowa, J. F. (1991). *Principles of Semantic Networks: Exploration in the Representation of Knowledge*. Mogan Kaufmann Series in Representation and Reasoning. Morgan Kaufmann.

Steyvers, M. and Tenenbaum, J. (2005). Small worlds in semantic networks. *Cognitive Science*, 29(1):41–78.

Taatgen, N., Lebière, C., and Anderson, J. R. (2006). *Modeling paradigms in ACT-R*. Cognition and Multi-Agent Interaction From Cognitive Modeling to Social Simulation. Cambridge University Press.

Teevan, J., Alvarado, C., Ackerman, M. S., and Karger, D. R. (2004). The perfect search engine is not enough: a study of orienteering behavior in directed search. In *CHI '04: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 415–422. ACM Press.

TheBrain (n.d.). TheBrain Technologies Corporation. Corporate home page accessed on the World Wide Web; Retrieved August 31, 2006, from http://www.thebrain.com.

Thornton, C. (2000). *Truth from Trash: How Learning Makes Sense*. MIT Press.

Todd, P. M. and Gigerenzer, G. (2000). Précis of simple heuristics that make us smart. *Behavioral and Brain Sciences*, 23:727–780.

Tranel, D. and Jones, R. D. (2006). Knowing "what" and knowing "when". *Journal of Clinical and Experimental Neuropsychology*, 28(1):43–66.

Tulving, E. (1972). *Episodic and Semantic Memory*, pages 381–403. Organization of Memory. Academic Press.

Tulving, E. and Thomson, D. M. (1973). Encoding specificity and retrieval process in episodic memory. *Psychological Review*, 80(5):352–373.

Tversky, A. (1977). Features of similarity. *Psychological Review*, 84(4):327–352.

Vinson, N. G. (1999). Design guidelines for landmarks to support navigation in virtual environments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems: The CHI is the Limit*, pages 278–285. ACM Press.

Wang, Y. and Liu, D. (2003). On information and knowledge representation in the brain. In *Proceedings of the Second IEEE International Conference on Cognitive Informatics (ICCI'03)*.

Want, R., Hopper, A., Falcao, V., and Gibbons, J. (1992). The active badge location system. *ACM Transactions on Information Systems (TOIS)*, 10:91–102.

Watts, D. J. and Strogatz, S. H. (1998). Collective dynamics of 'small world' networks. *Nature*, 393(6684):440–442.

Whittaker, S. and Hirschberg, J. (2001). The character, value, and management of personal paper archives. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 8(2):150–170.

Wickens, C. D. and Hollands, J. G. (1999). *Engineering Psychology and Human Performance*. Prentice Hall, 3rd edition.

Witten, I. H., Moffat, A., and Bell, T. C. (1999). *Managing Gigabytes: Compressing and Indexing Documents and Images*. Morgan Kaufmann.

Woods, W. A. (1975). *What's in a link: foundations for semantic networks*, pages 35–82. Representation and Understanding. Academic Press, New York.

Wynn, T. and Coolidge, F. L. (2004). The expert neandertal mind. *Journal of Human Evolution*, 46(4):467–487.

Yamaguchi, S., Isejima, H., Matsuo, T., Okura, R., Yagita, K., Kobayashi, M., and Okamura, H. (2003). Synchronization of cellular clocks in the suprachiasmatic nucleus. *Science*, 302(5649):1408–1412.

Zha, H. and Simon, H. D. (1999). On updating problems in latent semantic indexing. *SIAM Journal on Scientific Computing*, 21(2):782–791.