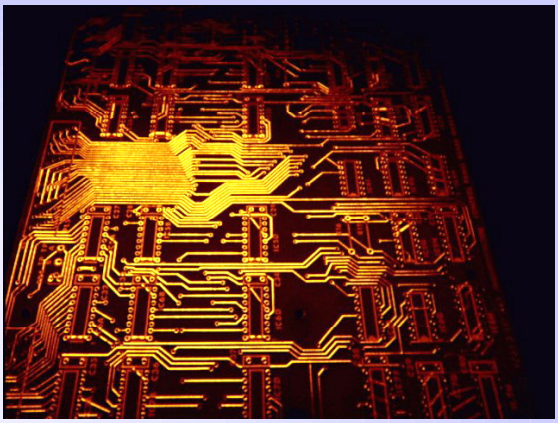

The Knowledge Navigator

An Auction Metaphor for the Brokering of Corporate Knowledge Assets

Machine Intelligence on the Edge: Java-Based Fuzzy-Neural Hetero-Genetic Object Oriented Knowledge-based Nano-Synthetic Reasoning Models: Throwing the Kitchen Sink at Problem Solving

Scianta Intelligence

1289 North Fordham Blvd. Suite A312
Chapel Hill, NC 27517

(919) 678-0477
www.scianta.com

PCAI

Nam et ipsa scientia potestas est.
For also knowledge itself is power.
Francis Bacon (1561-1626)
Meditationes Sacrae (1597)

Today's Web-centric distributed corporations are increasingly committed to the concept of knowledge management – the perspective first espoused by Francis Bacon, that knowledge is power. But they are rapidly discovering that warehousing knowledge is not the same as using that knowledge. For knowledge is not yet a commodity easily packaged for sale or distribution to the corporate consumer. Knowledge is a tangled web of interconnected resources – human experts, operational databases, historical archives, company policies, and the sometimes intangible knowledge gained through the company's interaction with its customers, its supply chain, and its own geographically distributed divisions.

In fact, the multi-faceted nature of knowledge places a significant burden on both the consumers of knowledge and the channels of corporate knowledge (its cadre of subject matter experts). Consumers are finding it more and more difficult to locate experts that have the time, skills, flexibility, and empathy to solve problems with increasingly higher priorities and increasingly tighter deadlines. The last constraint, an increasingly narrow window of opportunity to exploit knowledge, is tied directly to an underlying property of intelligence: All knowledge is perishable. Extremely important knowledge is highly perishable (otherwise it would be available for general use!) Thus consumers need to find experts that are available, empathetic, and have the ability to react within a compressed time frame.

A *Knowledge Navigator* bridges the gap between the consumers and the suppliers of knowledge. The navigator is not only a repository of expertise, but also an adaptive feedback facility incorporating client rating of the expert's knowledge and their abilities. More At root, the knowledge locator application provides flexible channels connecting knowledge consumers (the end users) with knowledge providers (the corporate technology experts as well as, in many cases, external consultants). As Figure 1 illustrates, this is often a difficult problem,

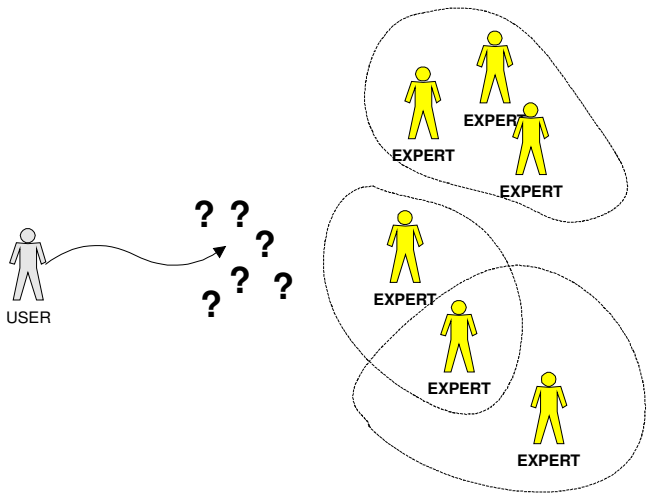

Figure 1. The Knowledge Navigator Problem

Finding the proper expertise for a particular user's problem is compounded by the often vague notions of what constitutes "expertise", the often equally vague statement of the problem by the end user, the qualifications and peer rankings of the expert, the usually high demands of a true expert's time, and the political, implicit as well as explicit costs associated with acquiring expert knowledge. And, as Figure 1 indicates, experts can be experts in more than one technological specialty.

When expertise is scattered across a globally diverse corporation, the difficulties in finding a required (and available) knowledge resource are intensified. Figure 2 schematically illustrates the way users and experts are distributed across international divisions, within corporate headquarters and other divisions in the continental United States.

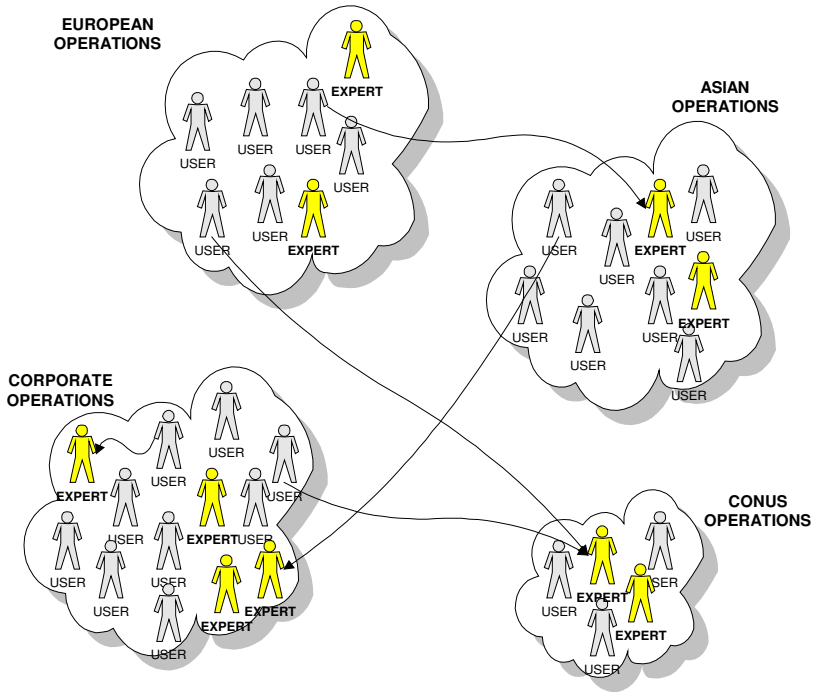

Figure 2. The Distribution of Knowledge Resources

In such a geographically (and consequently, culturally) diverse organization we are also faced with the prospect of identifying exactly what, within the organization, constitutes an expert or knowledge resource. While some divisions or affiliates may have full time and recognized technical staffs supporting end users in a wide spectrum of technologies (from waste disposal to molecular modeling to thin film chromatography to the principles of computing), some will have rather informal bodies of knowledge, experts will be members of the working staff, and quite often experts are identified more from the respect and interaction of peers and users rather from formal titles.

The Knowledge Navigator is an active, intelligent knowledge management system designed to address and remedy these problems trough an intelligent, adaptive feed-back system. The core functionality of the navigator maintains knowledge resources, logs problems and solutions (as a history of solved and unsolved cases), compiles client feedback as well as expert feedback, and, among other capabilities, maintains a matrix of associated expertise categories.

PCAI

Knowledge Management

We live in an age of high speed technological change. Not only is the depth of this Future Shock difficult to comprehend, but its endemic affects on the mechanics of modern business is hard to foresee and increasingly more difficult to manage. Managing change has always been a problem, but the accelerated global change carried into every office and consumer's home by the Internet has made accurately predicting even the direction of change highly suspect in most manager's eyes. Further, as we will discuss shortly, the disconnect between technology, the keepers of the technology, and the consumers of the technology continues to steadily grow.

Why is it important to connect technology providers and technology users in a coherent and reliable fashion? The nonlinear rate of change has a profound influence on corporations in several dimensions. Like ripples in a pond, the impact moves through the corporation at all levels, some in very subtle ways. There are many reasons for this new focus on technological impact.

The shared knowledge culture	Our society is not only moving from a service to a knowledge-based orientation, it is making huge quantities of this knowledge freely available. Organizations are caught in a paradoxical position of having to compete in a market of open intelligence while at the same time seeking ways to exploit their ever diminishing stock of proprietary knowledge.
The Move from physical to Web-based client/prospect management	The old store fronts are evaporating as new, on-line, Internet (Web) based sites provide secure transactions and highly visual interfaces with rapid navigation. Prospects and customers alike shop and order from the Web. Even non-retailing corporations use the web to promote their products, their corporate philosophies, their alliances, and their image among a vast sea of Internet browsers.
The Protection of intellectual property	The core sustaining asset in a corporation is its intellectual property. Innovation, client services, employee satisfaction, and longevity are all ultimately connected to the preservation and exploitation of a company's unique intellectual assets.
The value of that intellectual property	The value of knowledge rests on its economic utility and its availability (when needed and focused on critical corporate objectives). In this case knowledge becomes active intelligence. And intelligence is perishable, volatile, and difficult to quantify.
The ability to manage and understand change	Strategic management – deciding the company's products and missions in the next millennium – requires an understanding of technological change. Corporate management more and more draws on the aggregate expertise of formal and informal experts to steer its decisions in the right direction.
Knowledge Resource management -- Integration and leveraging of acquired	Understanding how working intelligence is distributed through the corporation is a critical component of a complete knowledge management program. In particular, this entails understanding the

expertise	dynamics of experts and their relationship to problem solving in the corporate culture. This issue is dealt with directly by Karl Wiig in his book, ***Knowledge Management Foundations***: *Thinking About Thinking – How People and Organizations Create, Represent, and Use Knowledge* (Schema Press, 1993).
Expertise Retention and Business Intelligence Codification	Corporations are learning that having experts is not the same as having a robust, business intelligence program. Without a way of encoding intelligence so that it can be deployed to knowledge consumers, the corporation's business process becomes fragmented – knowledge is available to those who are lucky enough to know the channels or fortunate enough to be in a site where knowledge has been (at least implicitly) codified. Creating a knowledge or intelligence broker formalizes the acquisition of expertise and makes it available on a demand (and need to know) basis throughout the enterprise.

Knowledge Management – the preservation and exploitation of the corporation's core intelligence – is *the* challenge for the next millennium. It is important for two crucial reasons. First, to discover and model what a company is really doing. Second, to change an organization's business model to meet the severe challenges imposed by the pervasive, over-whelming influence of the World Wide Web. Future corporations will prosper or stagnate depending on the effectiveness and scope of their knowledge management programs. One core component of this program, demanded more from external pressures then from the methodologies of formal knowledge management, is the identification and articulation of expertise sources in the corporation. Without this crucial interlock between end users and the scattered field of experts, the knowledge management infrastructure will be simply a static record of assets filling a few gigabytes of hypertext on the corporation's intranet service.

Knowledge Guardians, Providers, and Consumers

Coordinating the business intelligence of a corporation to address and embody both current and emerging technologies places a heavy burden on an organization. How can a corporation find and use its sources of knowledge? How can a corporation bring end users – the knowledge consumers (staff analysts, project managers, line officers, and technology workers) together with experts – the knowledge providers - who are globally and sparsely distributed? We also need to know who are knowledge providers and who are simply knowledge guardians. Well, we start with the basic relationship between the user and the expert. This is represented in Figure 3.

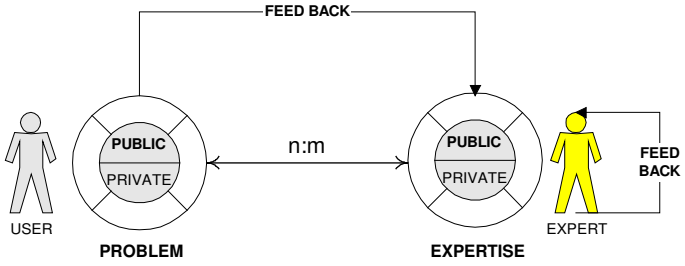

Figure 3. The Relationship Between User and Expert (Basic)

Figure 3 shows the mapping between a problem definition and a knowledge source or source of expertise (in this case a human expert). There are many users and many experts in an organization (which shouldn't be a surprise to anyone). But a user often needs or requests the expertise from multiple experts (in a one to many relationship). An expert also works for and responds to multiple end users (also in a one to many relationship). How an end user locates the available set of experts[1] is often a challenging and mysterious process. Thus, the Knowledge Navigator sits on the connecting edge between the User PROBLEM state and EXPERTISE or knowledge source. An integral part of the navigator is the feed-back loops ranking the quality of the expert from the end user's perspective and ranking the expert from a peer-to-peer perspective.

A Few Functional/Task Considerations

The Knowledge Navigator takes a client problem statement (or some description of required expertise with minimum confidence levels) and finds the resources that can be pulled in to satisfy the problem. At a very high level, the navigator works somewhat the schematic illustrated in Figure 4.

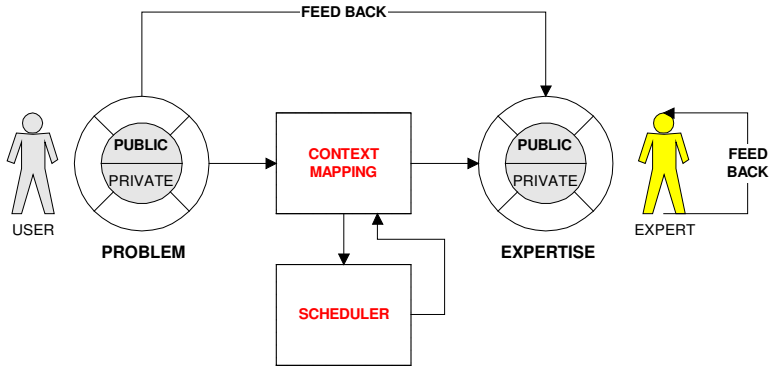

Figure 4. The Navigator Operation (High Level)

Context mapping translates the problem definition into a set of contextual maps (semantics, problem classes, expertise fuzzy sets, etc) which are applied to a global *Synopticon* (a multi-path knowledge index) to retrieve a rough collection of experts or other expertise sources (depending on the problem constraints).These sources are further reduced by the compaction expert system. Finally, the scheduler presents the client with a ranked list of experts. A possible ranking includes availability windows, probable costs, feed-back/peer ratings, and, when possible, alternates. We might also expect that the top candidates would display the expert's organizational affiliations and electronic mail address. Schedules always an expiration date.

[1] It is the function of the knowledge navigator, of course, to constrain the semantics of each end user problem specification. Thus, adaptively, we want the user to tell the navigator that it succeeded or failed in isolating the proper experts (actually we want the end user to indicate the *degree* to which an expert met or failed to meet the requirements of the problem statement).

Schedule versus Availability

A Knowledge Navigator schedule might **not** necessarily reflect the true availability of the resource. To do this, we would need to know the resource's concurrent and simultaneous commitments to on-going projects, off-site conferences, travel, vacations, and so forth. Such a perspective on scheduling real availability would entail integrating a complete project management component into the Knowledge Navigator. This is obviously not possible nor is it required in order to perform the primary function: identifying the location of specific experts throughout the enterprise. The schedule is based on two proposed attributes of the expertise resource: its shared opportunity vector and the resource's open dates calendar.

The shared opportunity vector (SOV) is a degree of share (in the interval [0,1]) across all the enterprise Locations (which are enumerated values in the Location container object). The value indicates whether or not the expertise resource is available to that location (as an example, an expert in Hong Kong might not be available to offices in Europe on the east Coast, thus the shared degree would be [0], the expert is definitely available throughout the Far East yielding a degree of [1], but may only be available on the U.S. West Coast under very special circumstances, giving a share degree of [.4] for west coast operations.) Unless otherwise specified, expertise resources are available globally.

The calendar is part of the resource class's Calendar class. Calendars are 6x7 compressed matrices one for each month and year combination. They provide a simple method of isolating an expertise resource from the knowledge navigator context manager. Figure 5 shows the prototype for an expertise resource's calendar.

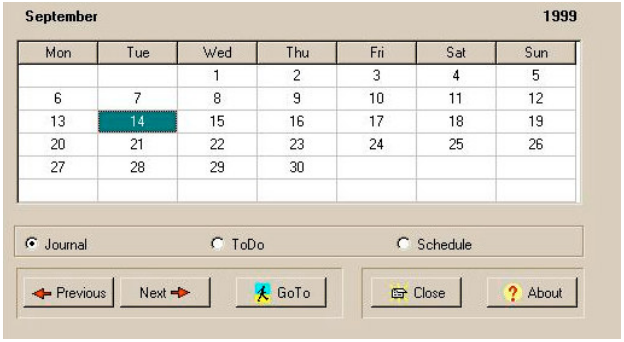

Figure 5. An Expertise Resource's Calendar

The resource simply clicks on a day (or a group of days or a group of months) and turns off the *Is_Available* attribute for that cell (which simply sets the matrix cell to [0]. All week ends and national holidays are automatically set to [0]).

We also supplement the Context mapping with a more generalized method of focusing on expertise that has been proven relevant to a particular problem. As Figure 5 illustrates, this is accomplished through the maintenance and use of a Case-Based Reasoning facility.

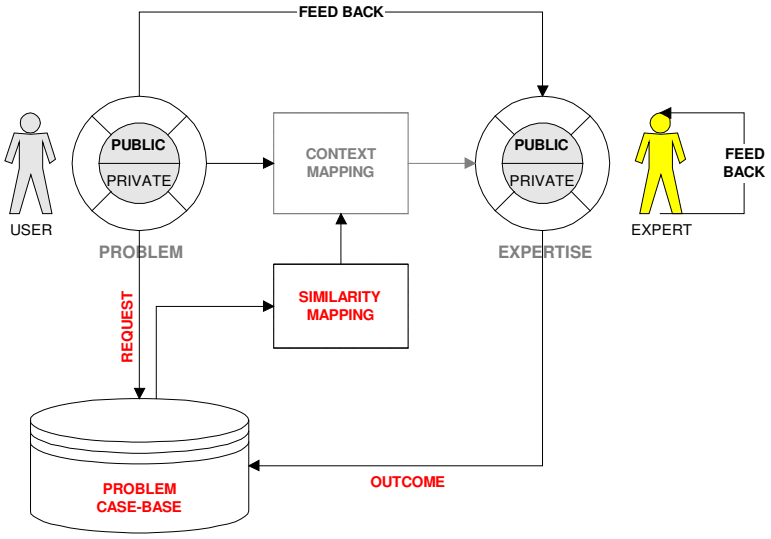

Figure 5. The Knowledge Locator CBR Component (High Level)

A Case-Base maintains a record of all the requests for corporate expertise (not all of these requests, of course, will actually be problems in the usually defined sense of the word). A case instance tracks the parameters and semantics (properties) of the problem request, the identifiers of the expertise elements brought to bear on the problem[2], and the outcome (if available). In this fashion, the Knowledge Navigator provides the Context Mapper with the reference identifiers of expertise resources that handled similar problem definitions in the past. As expertise resources are matched to problem states the case repository creates a weighted connectionist graph, establishing a correspondence between a problem state and a set of expertise resources. Figure 6 schematically illustrates how this graph might appear.

[2] When an expertise resource is removed from the Knowledge Navigator, the CBR repository is searched for problems that used this expertise. If the candidate expertise resource is used, then it is kept in the system as a TYPE-OF-EXPERT instance so that future CBR focusing activities can find the expertise and make a match to the most similar expertise (for a discussion of fuzzy similarity functions, see "*Coping with the Uncertainty Principle: Predictive Risk Assessment and Risk Classification using a Fuzzy case-Based Reasoning System*" by Earl Cox (PC/AI Magazine. Nov/Dec 1999).

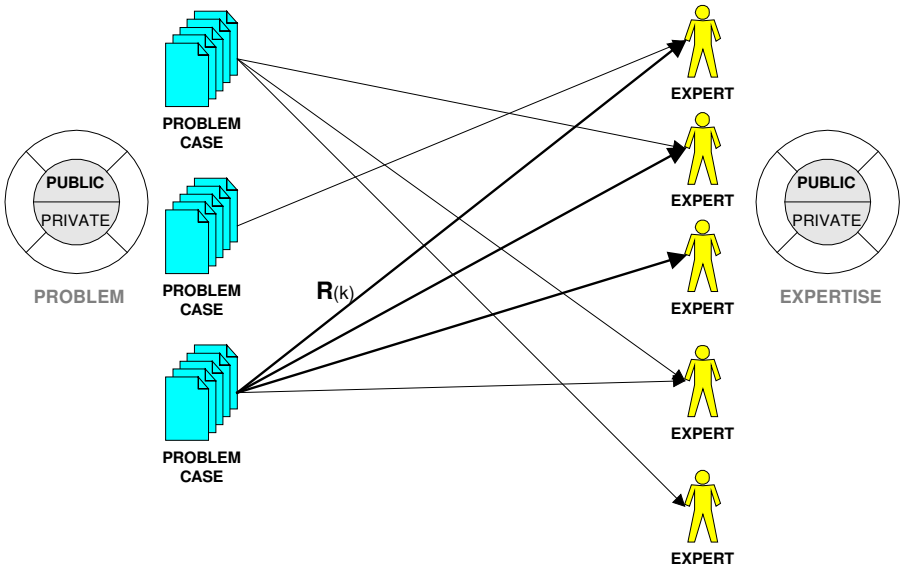

Figure 6. The Re-enforcement Graph

Each time a problem state selects a resource, the connection's weight – R(k) - is re-enforced. This is done using an adaptation of the Hebbian learning vector quantification scheme appearing in self-organizing neural networks (although this connection architecture is not a neural network). Thus, if the resource was a good match to the problem, then the weight is slightly increased; if the resource was not a good match then the weight is slightly decreased. Over time this re-enforcement approach will identify expertise resources that are strongly capable in resolving issues associated with a particular class of problem states (because problem states use fuzzy attributes, multiple similar states are often combined – this will have the affect of merging multiple connection edges in the re-enforcement graph. Such a proximity merge insures a high degree of robustness in the Context Mapping effort.)

Knowledge Navigator Infrastructure

We envision the knowledge navigator as a collection of persistent classes maintained in an object database A class (or object) represents a conceptual element in the navigator code space. Without delving into the nature of object-oriented design and programming, this approach caches attributes, methods, and data management into a single, cohesive entity. While not complete, Figure 7 illustrates some of the attributes and methods associated with the Expertise class.

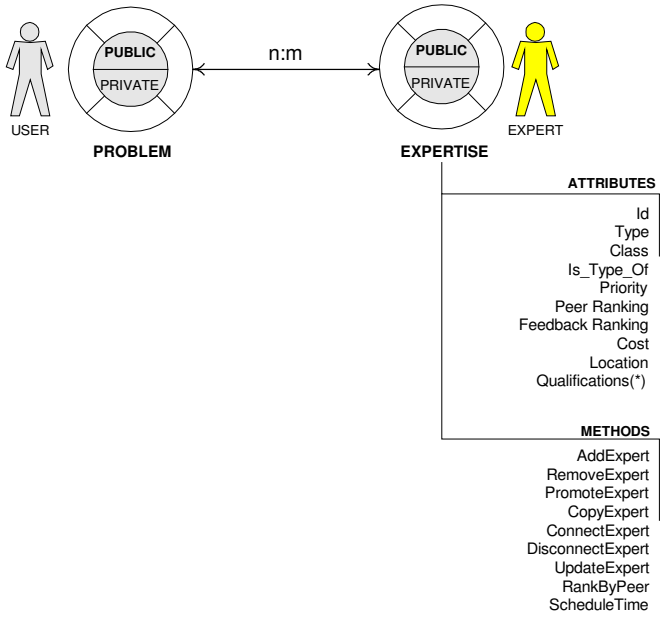

Figure 7. Expert Class Definition (High Level)

Each Expertise has a unique (global) identifier along with the such qualifiers as the resource type and the class. Peer and Feedback rankings are weighted averages of the rating vectors (not shown). Many different methods or procedures are associated with the Expertise class; these methods handle the management of the internal data structures (all object attributes have the private or protected attribute and are accessed only through an associated method). Of course the actual composition of an expertise resource is a little more complex, involving several associated and container classes. As a composite entity, Figure 8 shows many of the important classes connected to the expertise resource.

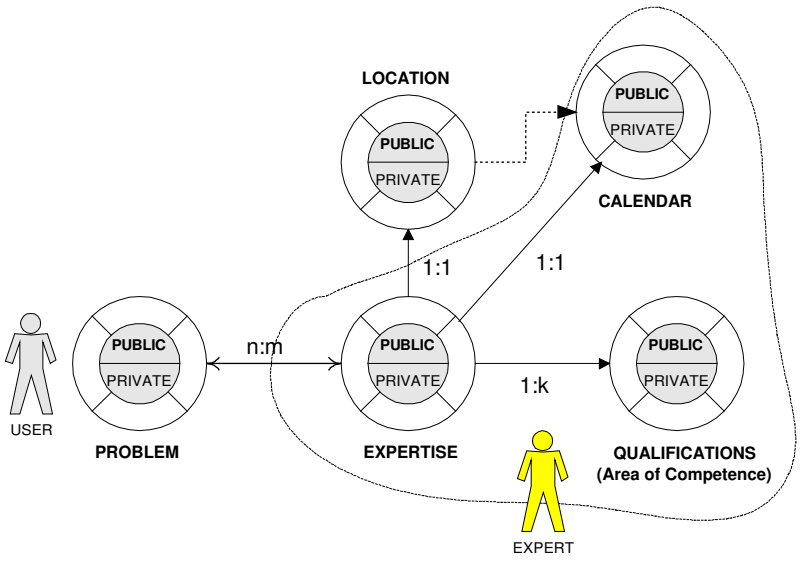

Figure 8. (Some) Classes Associated with the Expertise resource

As we can see from this schematic, a resource has a single location (its "residence") and a single calendar. However, a resource has many qualifications. These qualifications describe the degree and range of competence in the resource' skill set. The (Expertise + Qualification) instance is also indexed by the *Synopticon* to provide rapid access to experts with sets of qualifications or to experts with specific type and class attributes.

The Knowledge Navigator

Putting all the components together and highlighting the major system parts, Figure 9 shows the over-all architecture of the navigator. The core case based reasoning facility (CBR) keeps track of previous problem cases, the customer and the expert sources as well as the resolution of the problem (in fact, the degree of resolution is actually stored.) Rating histories are maintained separately from the selection process in order to avoid preference bias in the navigator.

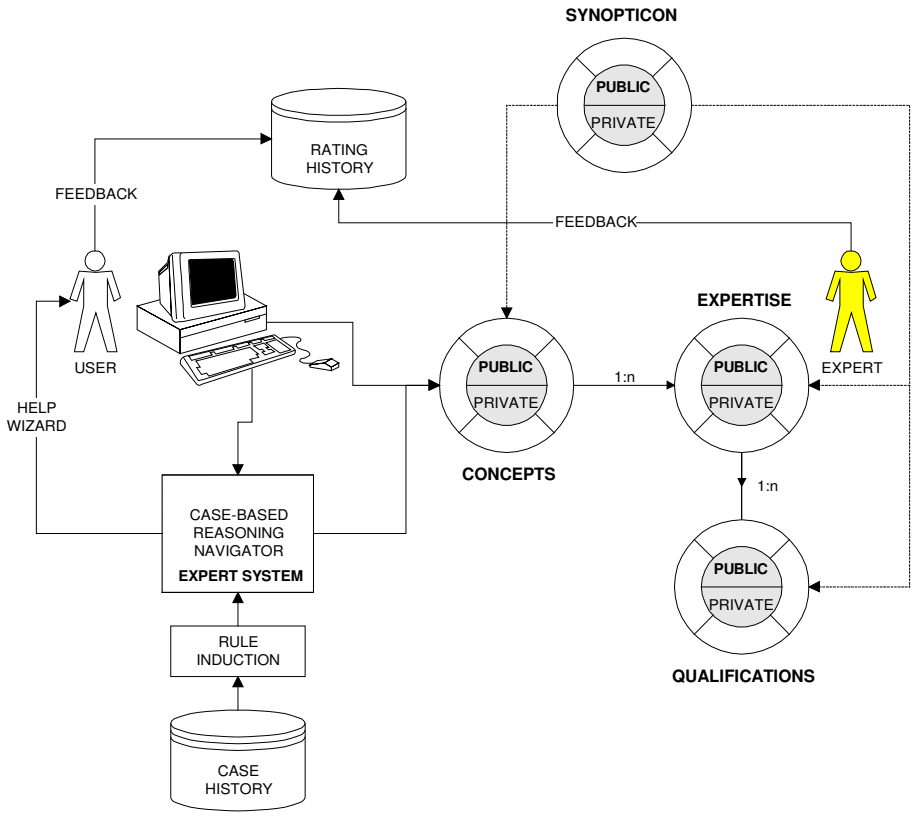

Figure 9. The Knowledge Navigator Architecture (Basic)

One addition design note is worth discussing – the use of the Concepts multi-dimensional object. The Concepts entity reflects all the types of expertise available in the enterprise. It is used internally by the Context Mapping function and the *Synopticon* to focus on the expert resources identified in a single request. However, the external representation of the Concepts object is a visual representation of the expertise centers throughout the organization. This is supported by a Self-Organizing Map (SOM, but also

called a Kohonen Net). This neural network-like structure provides a color-coded density map of the expert knowledge in the company. By clicking on a center of expertise, a user can visually navigate to the resources and their locations and well as discover related or similar centers of expertise. A second SOM encodes textual resources (journal articles, corporate technical documents, etc.) These are also expertise resources in the navigator (but the user can, restrict searches to human experts only, text resources only, or both).

The Knowledge Navigator provides a systematic way of coupling problems with experts (or, in a more general sense, with expertise.) The navigator has the virtue of dynamically reformulating its case base as problems are introduced and resolved. The feed back mechanism allows both consumers and experts to rank each other (thus providing a clear indication of the working behavior of experts and clients.) Since expertise is a rare commodity in an organization, the over-all navigator architecture provides not only a feedback ranking of services (such as we have on many auction sites) but also provides for a bidding mechanism – based on the complexity and feedback ranking of the end user. The Navigator concept thus provides a corporation with a flexible, robust and extensible method of activity engaging its knowledge. Through a proactive marriage of client with expert, the Navigator not only acts as a store house for critical corporate knowledge assets but also insures that these assets are brought to bear on critical problems.

PCAI

AR-PA-011