

Can Topic Maps describe context for enterprise-wide applications?

Duane Degler
IPGems

Lisa Battle
Lockheed Martin

Abstract

Topic maps provide exciting opportunities not just to make information easier to find, but to increase the usability of software. In order to provide users with the information that applies to their particular situations, in forms that they can use, software must be aware of a user's context (in a broad, multi-dimensional sense). Topic maps can serve as the language for linking information to software applications and for sharing information about context among applications. Using topic maps in the design of user-centered software applications for the U.S. Social Security Administration, we have encountered several interesting issues that are not necessarily found in the design of stand-alone information resources. In the future, there is an increasing potential role for topic maps in the design of flexible applications. The infrastructural investments required to achieve richly contextual applications can be made gradually, over the course of many small and successful projects, as the role and benefits of Topic Maps becomes increasingly visible.

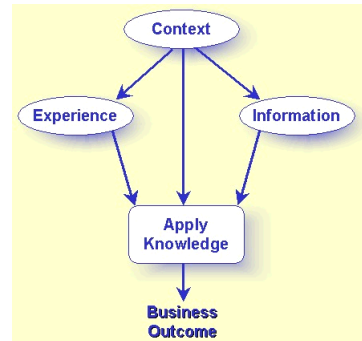
Can Topic Maps describe context for enterprise-wide applications?

Table of Contents

Introduction: our preoccupation with context.....	1
The user-centered approach to software design.....	1
What do we mean by context?.....	2
A potential new role for Topic Maps as the “common language” for context.....	3
Using Topic Maps to provide contextually relevant information.....	3
Case study: a government legal research application.....	3
The first stage: “I know what I’m looking for — get me to it!”.....	4
The second stage: “I don’t know what I need to know”.....	5
The third stage: “Just tell me whatever I absolutely must know”.....	9
The fourth stage: “I didn’t know I needed that”.....	12
Benefits of proactive information.....	13
Building and maintaining the Topic Map.....	13
Topic creation by end users.....	13
Can a map be both generic and situation-specific?.....	14
Using Topic Maps to create contextually relevant applications.....	15
Providing an API/Web service to information resources.....	15
Creating flexible, modular applications.....	15
Storing and maintaining business rules.....	16
Translating between business and technical language.....	17
Conclusion.....	17
Bibliography.....	18
The Authors.....	19



Figure 1: Illustration of the relationship between context and knowledgeable performance



Can Topic Maps describe context for enterprise-wide applications?

Duane Degler and Lisa Battle

§ Introduction: our preoccupation with context

The user-centered approach to software design

We did not come to the emerging topic map paradigm as technologists or information scientists. That is to say, our goal is not to simply organize information or develop XML. What interests us about topic maps is their potential in helping users perform work tasks effectively and have a positive experience when using software. Our professional practice is in user-centered design, organizational communication, and knowledge management. User-centered design is a methodology for designing usable software by focusing on user needs, the goals of the tasks being performed, and the context in which they are performed. It evolved from innovations in the disciplines of human factors, usability engineering, and software engineering, to become a recognized and widely accepted practice within the software development industry [Norman and Draper, 1986] [Winograd, 1996] [Landauer, 1995] [Beyer and Holtzblatt, 1998].

Why do people use software applications? The goal is to achieve a successful outcome, typically to help in achieving business goals. A successful outcome can be understood through quantitative measures (for example, the user successfully processed 15 orders, wrote 800 words, illustrated an idea) and through qualitative measures (for example, the user made no serious errors, or was satisfied with the experience). Businesses are typically interested in improving the quantity or quality of outcomes, and this desire is what prompts most large-scale systems development initiatives. When we are called in to help in the design of software, we always analyze user and business performance and look for ways in which it can be enhanced through better design. We find that the best designs — those that support successful performance — are those that fit well into the context in which the task is performed, provide the appropriate level of support for the user's level of experience with the subject domain, and provide relevant information at the time of need. Conversely, the worst designs — where the software actually becomes a barrier to successful outcomes — are those that fit poorly into the work context, fail to anticipate the user's level of experience, and provide irrelevant information or no information to the users. The relationship between these factors is summarized by Figure 1 below [Degler and Battle, 2000]. In order to enhance business outcomes and support successful user performance, we need to find ways of bringing context, experience, and information together into business applications.

Since late 1980's, we have been looking for ways to make task performance easier and the supporting information more immediately relevant, whether those tasks are operating complex machinery,

processing purchase orders, or changing an address online. A key design challenge for us has been to capture and use context in order to create flexible, adaptive application interfaces and knowledge-supported environments. We have often custom-designed approaches for managing and relating context data. Now, the introduction of topic maps appears to offer a new and better way to address this challenge.

What do we mean by context?

Context is the collection of relevant conditions and surrounding influences that make a situation unique and comprehensible. The human cognitive and perceptual systems are designed to identify and use context automatically as we go about our daily lives [Anderson, 1995] [Hasher and Zacks, 1984]. It is harder than originally imagined to teach a machine to do the same. If software could adapt to context as easily as people do, it could dynamically change its sequence, the style of interactions, and the type of information it provided to different users under different circumstances. The ideal application should be able to provide information that is both accurate and relevant without requiring the user to actively seek this information and determine its relevance.

A few years ago, we identified several broad categories of information that can be used to describe the current context [Degler and Battle, 2000], as shown in Figure 2 below. The reason for doing this was to begin to apply a series of filters for relevance. Consider the category of “location,” for example, and assume that an organization has some policy information that applies to employees in Florida but not to employees in New York. If the current user is based in Florida, that information may be relevant, but if the current user is based in New York, it isn’t. If a staff member in Florida is asked to work temporarily in New York, then a different body of content is required (relevant to New York), provided at some level appropriate to the person’s experience with the subject domain (high) and the specific content (low). If the organization has some information that is targeted to attorneys and other information that is targeted to salespeople, the user’s job role may be an important category of context. If the current user is an attorney, the information targeted to salespeople can be filtered out, and vice versa. The categories of context are represented as overlapping layers, because as each layer is added the filtering becomes progressively more refined. We anticipated that filters like these could limit the universe of available information to only the most the relevant information and even determine the appropriate behaviors for applications.

Figure 2: Illustration of some of the key categories of context



We also described two approaches to identifying context in business applications. The first approach, which we called **context description**, occurred when the human end-user of the system provided a description of details that were outside the computer environment. The second approach, which we called **context recognition**, occurred when the automated system used the data already available about the user, the business process, other software applications, and the subject domain to identify attributes of context. The advantage of the first approach was that the person was in control, but there were several important disadvantages, including the burden on the user to seek information, and the potential that the user might have an incomplete understanding of the subject, the problem, or the current situation. The advantage of the second approach was removing the burden from the user, but there were several

possible disadvantages if the application were to incorrectly interpret context and push unwanted or irrelevant information to the user (or worse, provide inadequate decision support). We recommended that the best solution was often to combine those approaches so that some attributes of context are drawn from the individual (description) and some from the computer (recognition).

A potential new role for Topic Maps as the “common language” for context

According to the user-centered design perspective, to make applications more useful and usable, they have to know as much as possible about the user’s context — and they have to be able to share that information with each other. To share that information, *they need a common language*, or at least common reference points to the same subject. One of the challenges in implementing our vision of the layers of context has been the need for a machine-readable way to describe context that can be shared among different applications and easily maintained by humans. We believe that common language can be the language of topic maps. However, as with any technology, our exploration and evaluation helps us identify whether it will prove to be feasible and appropriate for the task.

Although we recognize that topic maps were originally conceived as a flexible answer to “the traditional back-of-book index in the world of electronic information” [Pepper, 2000], we see opportunities to extend their use into the realm of software applications. There is a potential to use topic maps as the “glue” that connects software applications to each other and to supporting information resources. According to the international standard [ISO, 2002], topic maps may be used for the “delivery of partial views depending on user profiles and/or knowledge domains.” This need is as least as vital in the design of software user interfaces as it is in the design of standalone information resources. There is even a further possibility of using topic maps to create more flexible application architectures. Topic maps provide an opportunity to offer a consistent framework for software developers to create the tools that we need to realize our dream of flexible, adaptive application interfaces, and knowledge-supported environments.

§ Using Topic Maps to provide contextually relevant information

According to Steve Newcomb [Newcomb, 2003], topic maps are commonly used for one of three reasons: to make information available, to require people to look at information, or to prohibit people from looking at information. Most implementations of topic maps so far have focused on making information available on stand-alone information sites. Because our primary concern is supporting user job performance, and we typically work with information resources that are closely connected with job tasks and software applications, we encounter several interesting issues that are not generally a concern for stand-alone information resources. The first is that it isn’t good enough to make information readily available; the information must be *relevant* to what the user is doing and not distract from that (relevance is based on fitting in with the user’s context, as described above). The second issue is that there is a wide spectrum of *user awareness* of their need for information. On one end of the spectrum, a user might be aware of a specific need and able to clearly articulate the problem. Users in the middle of the spectrum have awareness of some type of need but vary in their ability to articulate what it is. On the other extreme, a user may have no awareness at all of a need for information. Sometimes information is required for successful task completion, and so must be provided proactively if the user is not aware that they need it. This situation is as likely to occur for “expert” users — who may already feel they know something and won’t take time to look for information [Klein, 1998] — as it is for inexperienced users who don’t recognize that they are encountering a situation that requires additional information.

Over the past two years, we have had the opportunity to influence the design of a range of applications and information resources within a large government agency. Through this work, we have explored more deeply the interaction between applications and collections of information resources, and the implications on large organizations. In this section, we will illustrate how we have begun to create and use topic maps to provide contextually relevant information in support of job performance.

Case study: a government legal research application

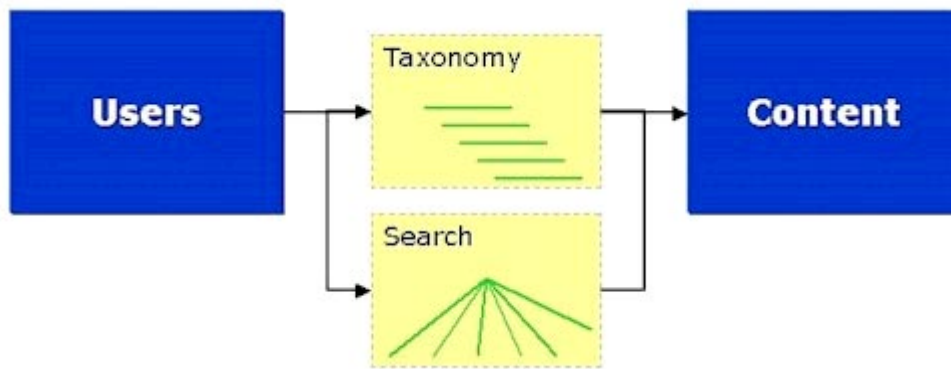
When the Benefits Improvement and Protection Act of 2000 legislation mandated that Medicare appeals case processing time be reduced to less than 90 days, the OHA [Office of Hearings and Appeals] (which is a part of the Social Security Administration) saw the need to help people access the wealth of information about Medicare that was only available on many different external web sites and internal hard disks.

When the project team began to analyze the OHA’s information needs, they uncovered several challenges. Most of the 1,000 Administrative Law Judges in OHA will hear a Medicare case very rarely, as their primary duty is Social Security disability cases. Medicare policy/law is a broad complex subject that changes frequently, so traditional training is not feasible. The people who need to use the Medicare information range from 21-year-old case technicians to 80-year-old judges. In addition, most Medicare information is on external web sites. As a result, case research is very difficult and time-consuming.

The first stage: “I know what I’m looking for — get me to it!”

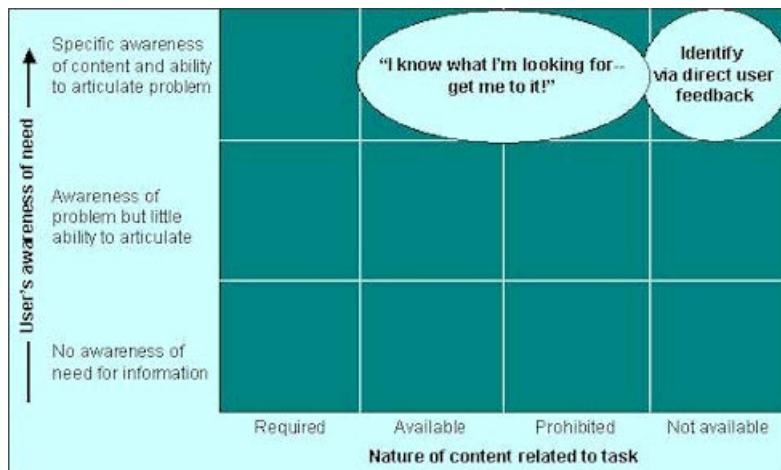
At the start of the project, the stakeholders described their ideas for taxonomy-based navigation, and they also expressed the desire to be able to search for information. This is consistent with traditional approaches to connecting users and the content that they need. As shown in Figure 3 below, there are commonly two ways for users to get to their content: drilling down through a hierarchical taxonomy structure or using a search engine.

Figure 3: Common existing approaches that are used to locate information



Both of these approaches to locating information are well-suited to situations in which the users know what they need, can articulate it accurately, and are motivated to go looking for it. However, as shown in Figure 4 below, this situation represents only a small part of the overall picture of information need.

Figure 4: The scope of the more traditional approaches to locating content



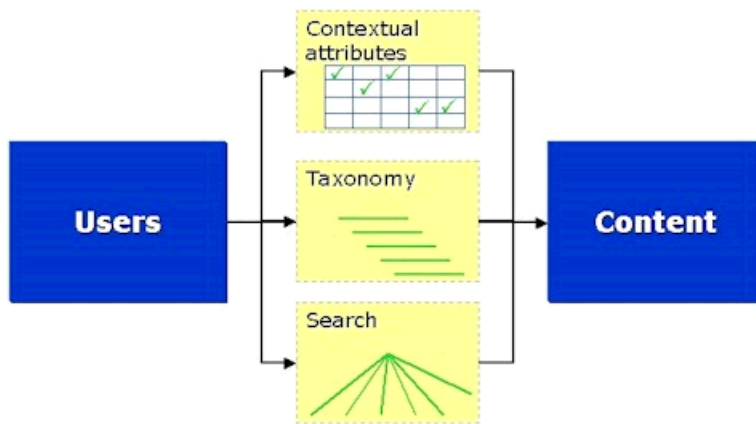
Note about “Required” information: The label is presented this way for the sake of brevity. It should be interpreted as “required for successful task completion,” because the requirement may come from the rules of the organization or the specific context of the task being performed, rather than what the user believes to be required at the time.

We were concerned about the limitations of taxonomy and search. To successfully use taxonomy for navigation, people must be familiar with the hierarchical structure of the information, and they must know or at least recognize the terminology. There are also typically few ways to research and filter across the hierarchy, to explore different aspects of a subject. Using search is sometimes considered faster because it eliminates the need to know and navigate the structure, but it requires the users to know the terminology (without the benefit of seeing and recognizing terms, as often happens in the taxonomy). People who are not familiar with the terminology struggle to use search because they are not able to formulate appropriate queries. Given our user population’s inexperience with the subject, the highly technical nature of the complex subject, and the lack of consistency in the information resources they needed to use, we believed that these approaches would not meet the organization’s performance needs to significantly reduce the research time required, nor would they provide a positive user experience.

The second stage: “I don’t know what I need to know”

So if users don’t know the subject well, what could we do? The approach taken was to focus on what users would know, no matter what their level of experience or familiarity with the case issues in front of them.

Figure 5: Approaches applied for Medicare Online to support users



A third method of navigation was developed that allowed users to quickly describe the context of their need for information, based on three questions:

- What is the case about? (the subject)
- What are you doing now? (the task)
- Do you need only specific types of information?

Figure 6: Simple selection of topics supports most common case research

Quick Start

Select any of the options below to go directly to information relevant to what you are doing.

[Legal Considerations](#)

Choose your case type:
Part B

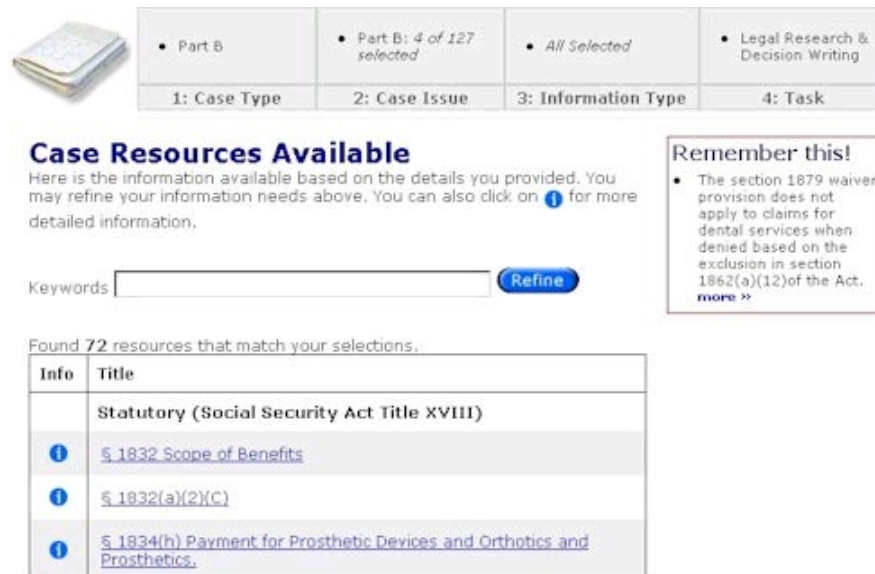
Choose your issue:
Chemotherapy (Drugs/Biologicals)

Choose the information type:
-- show all information types --

Choose your task:
-- show all tasks --
-- show all tasks --
Docket a Case (Tasks)
Transfer a Case (Tasks)
Pull a File (Tasks)
Create an Exhibit List (Tasks)
Obtain Additional Documentation (Tasks)
Issue Identification (Tasks)
Party Identification (Tasks)
Legal Research & Decision Writing (Tasks)
Schedule a Hearing (Tasks)
Notices (Tasks)

The result is specific information relating to the case they're working on, requiring no special knowledge of the information landscape that they are traveling. Information is presented in a ranking based on the organization's identification of the value to the task — legally binding information that is mandatory for citation in a legal decision is presented near the top of the list, with other supporting resources (such as training or medical references) presented further down the page. The selections made by users within the three broad topic areas — with the ability to refine those choices — remains on the page for the users' convenience.

Figure 7: Results based on the situation described by the user's topic selections



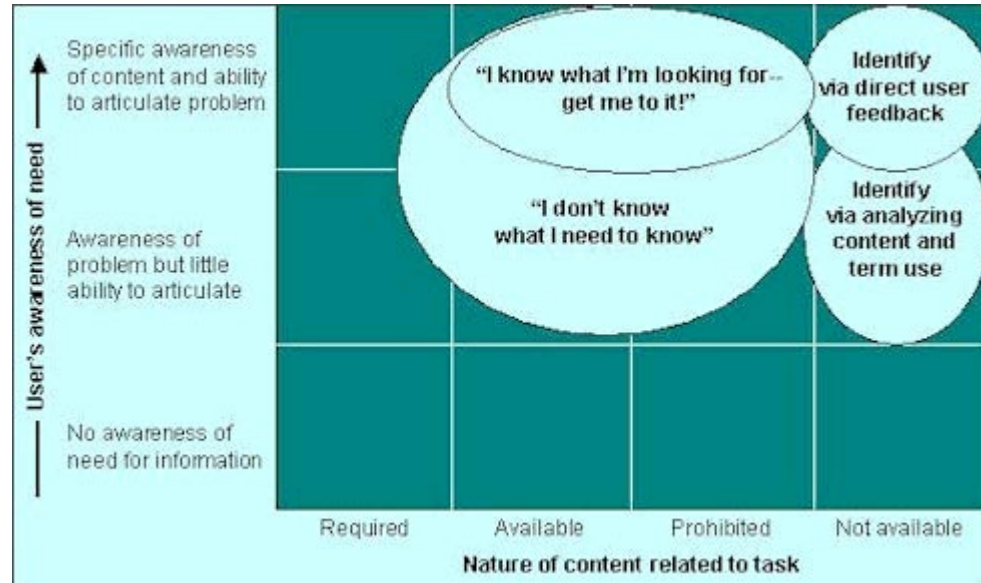
The screenshot shows a search interface with the following elements:

- Filters:** A horizontal bar with four categories:
 - 1: Case Type (selected: Part B)
 - 2: Case Issue (selected: Part B: 4 of 127 selected)
 - 3: Information Type (selected: All Selected)
 - 4: Task (selected: Legal Research & Decision Writing)
- Case Resources Available:** A section with a sub-header and a paragraph of text. Below it is a "Keywords" input field and a "Refine" button.
- Remember this!:** A callout box containing a bullet point: "The section 1879 waiver provision does not apply to claims for dental services when denied based on the exclusion in section 1862(a)(12) of the Act. [more >>](#)"
- Search Results:** A table with the following data:

Info	Title
	Statutory (Social Security Act Title XVIII)
i	§ 1832 Scope of Benefits
i	§ 1832(a)(2)(C)
i	§ 1834(h) Payment for Prosthetic Devices and Orthotics and Prosthetics.

This topic-based approach to navigation that we have described is reflective of current theoretical perspectives of knowledge organization, categorization, and representation [Fisher, 2003] [Sigel, 2003] [Hackos, 2002] [Berners-Lee, Hendler, and Lassila, 2001]. We have applied a design approach to a user interface that increases usability and supports user task performance because it removes the need for users to know the structure of the taxonomy and to know the appropriate terminology. They can start with whatever they do know and still get useful results. The topic-based approach to navigation helps to address a broader range of user information needs, as illustrated below.

Figure 8: Topic maps improve ability of less experienced users to locate information



The issue remains, however, that we still have to rely on users recognizing their need for information and being motivated to seek relevant information. The information resource is completely passive, waiting for the user to initiate a request for information.

The third stage: “Just tell me whatever I absolutely must know”

Given the constraints and pressures of typical work environments, where there are backlogs of work to be processed and limited time available, it is not unusual for people to occasionally forget things, skip a step, or rely on their own memory rather than turning to reference sources. As application designers, we need to balance the needs of the users with business goals and legal requirements. Sometimes there are compelling reasons to make sure the users are aware of specific information.

In our case study, one important feature of the interface was dubbed by the users “Remember This.” Certain information was deemed to be important enough that it was to be presented to the user directly on the screen without relying on the user selecting a link. To ensure relevance, this information was categorized so that it would only appear in the appropriate context. The “Remember This” feature was used primarily to provide hints about avoiding known misconceptions and errors in case handling, or to present important new information that may change the way a case is decided.

Figure 10: Illustration of how a required AR [Acquiescence Ruling] maps to the user's task

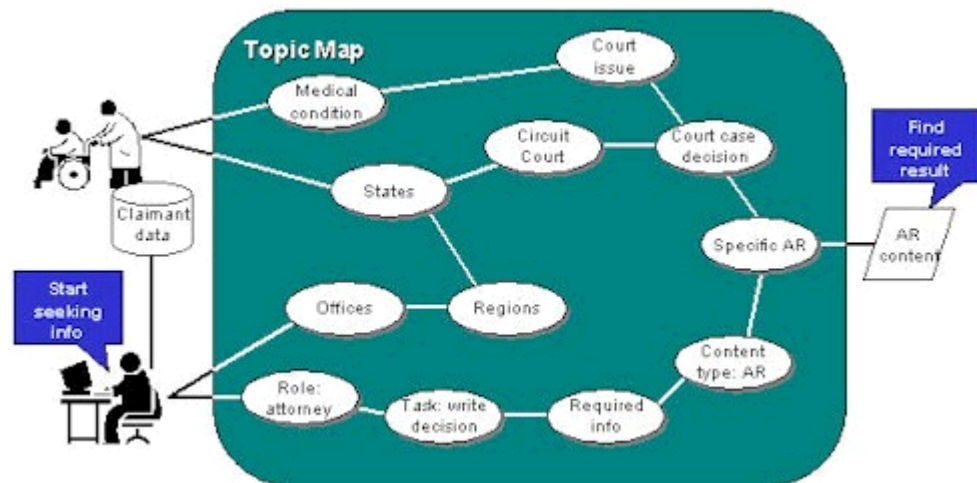



Figure 9: "Remember This" was created to display required information



In subsequent work, we have further developed the approach that we have started calling "situation-based filtering," looking for ways to reduce the amount of information a user would have to provide while still getting highly relevant results and reminders. We developed a proof-of-concept topic map application for the complex environment of disability claims determination, *i.e.*, identifying whether an individual is disabled and also qualifies for financial support. One scenario explored how the system could notify the user of a court ruling that affects their decision on a case, in particular a court ruling that relates to the claim but is based in a different Circuit Court jurisdiction than the user is familiar with. At a very basic level, the topic connections between the user and the content looked like this:

Using traditional table of contents navigation, a user would typically require 6-8 different steps to get to the list of rulings, and then would have to read them to determine which one was appropriate to the case (document titling is typically based on the case litigants and the Court references, not the subject matter). Using a site search, the relevant documents would rarely appear in the ranking higher than number 10, and would appear below other similar documents that were not relevant to the key facts in the case.

Using the situation-based filtering approach, with the application navigating to content via the topic map, the required information was retrieved quickly and successfully by the user making only four selections from a set of drop-down boxes on a single page. All selections were based on information clearly available in the summary of the case itself.

Figure 11: User interface to support accessing required information efficiently

Quick Start

Select the topics relevant to your situation and case.

Office: X94
Region: 4
Your role: Decision writer
Task: Drafting decision

Claimant location:

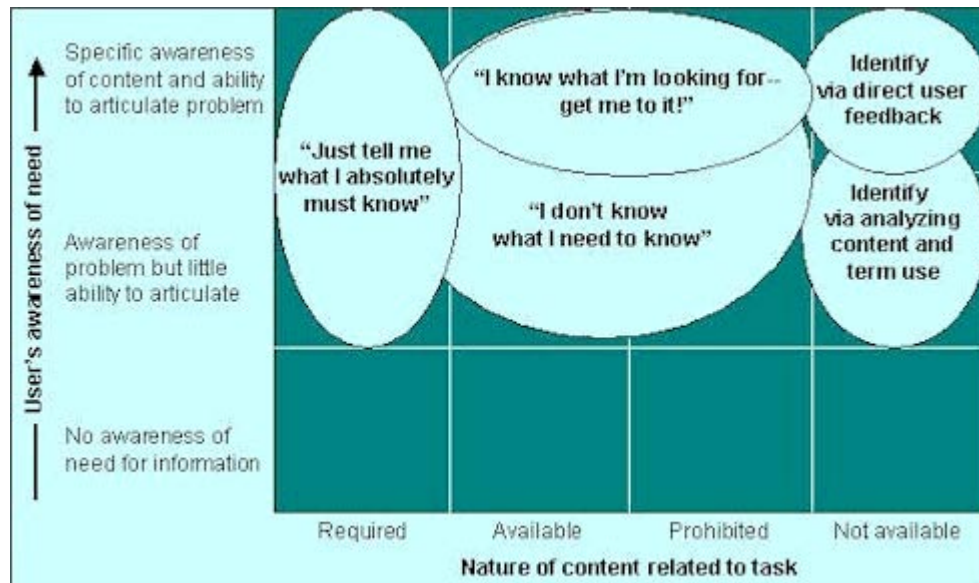
Other state involved:

Disability type:

Decision being written:

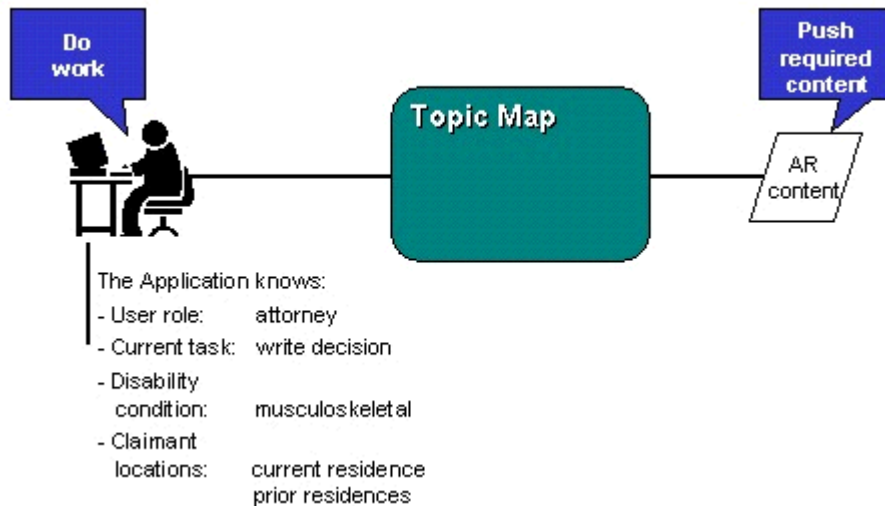
[Find Resources](#)

Figure 12: Topic maps used for locating required information



This example moves beyond the idea of making information readily available to users into the realm of providing information that is essential to the successful completion of the task.

From our perspective, this represents an important step toward our overall vision. It involves simplifying the interaction for end users, based on working with what users know about their situation, rather than

Figure 13: The application describes context, supporting proactive information retrieval


assuming they have “learned” a formal taxonomy or can “speak” a formal organizational language. More importantly, it draws the users’ attention to information that can improve their job performance and successful outcomes. However, it still remains “above the line” of user awareness of their need for information.

The fourth stage: “I didn’t know I needed that”

A significant area of poor performance can arise when users don’t know that they need help or don’t know that information is available that is required for their successful completion of a task. This is sometimes true for experienced users who don’t look for information because they believe that they already know the answer. It can certainly be true for novice users who do not know the subject domain well enough to identify problems and/or to find the supporting information that could help them.

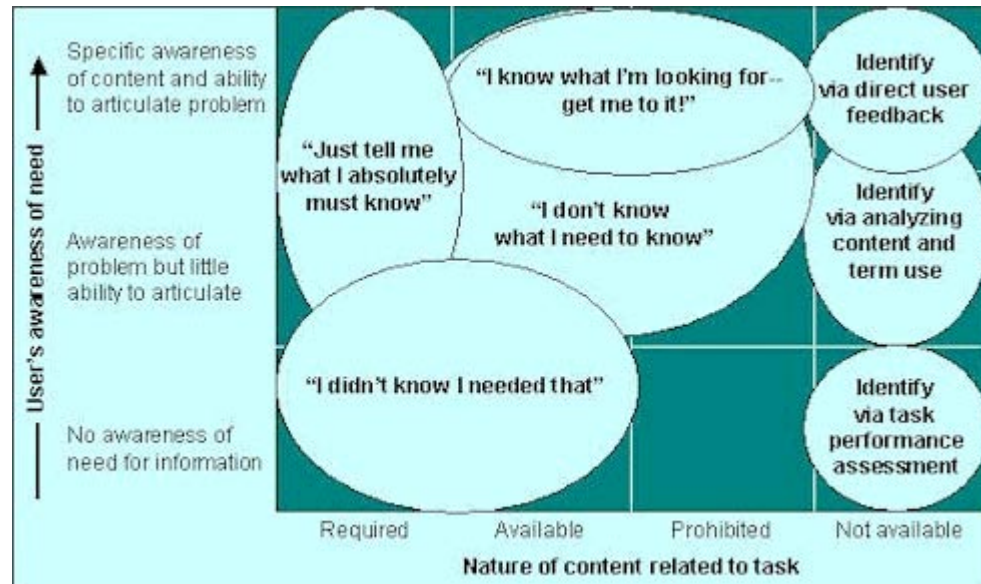
In these circumstances, the passive methods of providing information (as discussed above) offer no value to improving user performance. Something in the environment — the application that a user is using to perform the task — needs a mechanism to recognize the gap between anticipated levels of performance and the actual performance exhibited by the user, and then the ability to proactively request information resources that the user needs. To do this, the application needs two things that we believe can become part of the enterprise-wide topic map construct:

- An expression of required task performance (a form of task validation or business rules)
- A common language to request information or process support from an information repository, or an action to be performed by another application

Let’s return for a moment to our example of the attorney writing a decision regarding a claimant who, while living in one state, was also covered by a Circuit Court judgment that was only applicable in another state — a judgment that the attorney was unlikely to know about. What information did the user provide in order to get that information? Two state names, the nature of the disability, and the task being worked on. *All these things are known by the case processing application.*

Similarly, the “Remember This” feature that provided hints about case processing and reminders of recent updates was designed to eventually be incorporated into the case processing application. From a user-centered design perspective, it is far more effective to provide the hint directly on the relevant screen of the case processing application so that the users will see it while they are doing their normal work. This prevents any interruption of the work, reduces the business risk of the users not seeing the information, and removes the burden from the user to go seeking information.

Figure 14: Proactive information delivery via applications



Once applications begin to provide information to support performance problems that may not be recognized by the user, information becomes more valuable, based on the fact that it is both proactive and relevant to the user's situation. Thus, the picture of the user/content relationship now begins to look significantly different.

Benefits of proactive information

The value that topic maps bring to this more integrated information environment can be illustrated as follows:

§ Building and maintaining the Topic Map

This section explores some of the issues and decisions faced when creating the infrastructure that facilitates both a general enterprise topic map and specific applications' use of the topics.

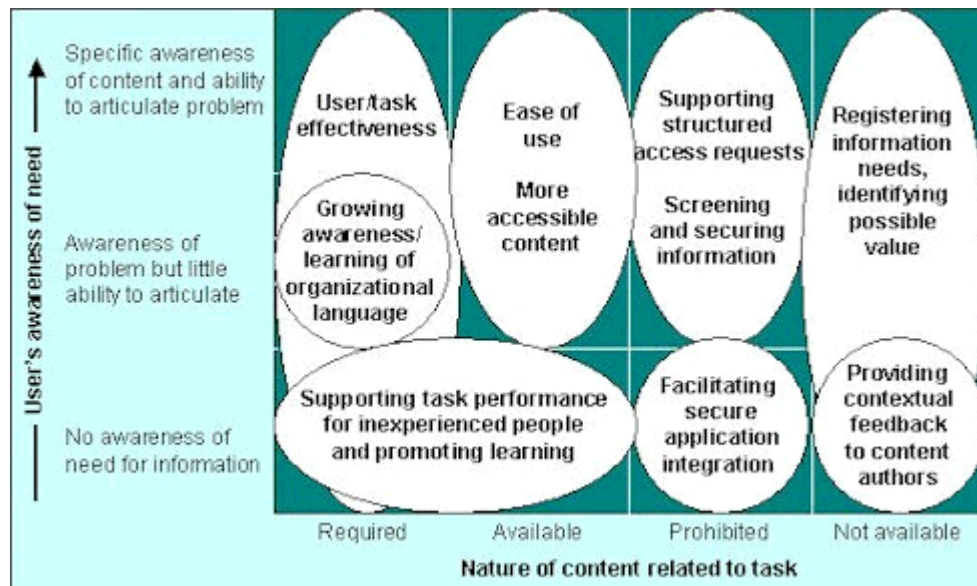
Topic creation by end users

The first, and probably most important, decision was that topics need to be created and maintained over time by domain subject experts within the organization (not solely by information specialists or technical markup staff). The topic map should grow in a systematic-organic way in order to meet the immediate needs of particular information projects and still extend to other subject areas (of unknown scope or subject focus) in the future. Having many people create topics and associations in a structured environment allows the library of available topics to grow quickly as well as maintain relevance to the organization. The long-term goal is to build an organizationally maintained, enterprise-wide, dynamic resource. Different applications can then begin to share one topic reference source for their indexing and context mapping.

An alternative would have been to create application-specific topic maps and then merge them when required. Topic map merging has not been pursued at this stage because we felt that it would probably have required a significant future manual effort to review and decide on merges, or it would have had to rely on a high degree of terminology (and spelling) consistency among tens — or hundreds — of subject specialists.

What are the implications of this approach?

Figure 15: Some benefits gained from topic mapping



- Adding and changing topics must be easy for a “non-specialist” to do (don’t burden the user with arcane topic map syntax and behaviors)
- Risk must be reduced: adding topics, associations, and scopes must not break existing sites and uses
- Topic association syntax must be standardized to allow the map to be processed and queried automatically by individual applications, and yet flexible to allow new association types in the future

Simple user interface designs have been developed to allow subject experts to add “keywords” (topics) and make connections between them (associations, roles, and scopes) based on defined templates. This is consistent with literature on using templates as a basis for creating common associations [Biezunski and Newcomb, 2001] [Vatant, 2002]. In existing user content, we have encountered common categorization and vocabulary patterns that form the basis for some pre-defined templates that map taxonomy organization, context filtering, synonym terms, acronyms, reference numbering, and replacing one topic with another as vocabulary or business circumstances change (this is likely to be particularly useful during departmental reorganizations).

It is expected that in the future the topic map templates and authoring facilities will expand on our basic foundation. For now, we are building the first draft set of enterprise-wide topics in a way that we believe will minimize rework on tens-of-thousands of topics and associations in the future as technology changes.

Can a map be both generic and situation-specific?

Will it constrain the design of future applications to use a common topic map, rather than develop vocabulary and relationships that are specific to the needs of each application? Of course, the benefits to the users from having more consistent, predictable terms and interactions are clear and desirable. Some classification schemes already exist in parts of the organization, but they are currently duplicated in many publications and systems, increasing the risk of error and the problem of updating content. We believe that the benefits to the organization and the systems development community of having an enterprise-wide “common language” outweigh any possible limitations that might arise. The consistency and expandability of the topic map will reduce duplication of effort and increase value across projects, as well as allowing the greater richness of context that is needed by applications.

As described above, we've started with simple association structures that mirror existing approaches to categorizing information. The topic map application presents a simple way to create hierarchies of information based on existing classifications already prepared by subject experts (for example, detailed classifications of medical conditions and their relationship with disability claims). The topic hierarchies can then interact across multiple dimensions (for example the specific sub-classifications for adult disabilities and for childhood disabilities).

To support the need for contextual relevance, users can define where certain topics are *only* relevant in the context of other topic areas. For example, detailed disability claim information is not relevant to all types of benefits claims (such as retirement, some widows benefits, etc.). Users assign particular associations which declare that medical tests for a type of disability are specific to processing disability claims (meaning they are not generally relevant to retirement claims). Scope appears often to be used for this type of context-setting [Pepper, 2001]. In the enterprise topic map, we chose to implement context-setting connections as associations (rather than scopes) because it allows some flexibility for alternative uses of the topics in the future — otherwise, we could foresee a situation where some topics could have dozens of scope statements in order to be recognized (and not ignored) by various applications.

However, scope will be used to support context in specific applications. The general context-setting associations can be inefficient to program at the individual application level, and also may impact performance of the application for the end user when the number of topics gets very large. In those cases we transform the associations within a specific application — it “reads” the association from the enterprise topic map and then “writes” a scoped topic for use by that single application. This transformation process helps us be specific to particular situations where required, while retaining the generic nature of the enterprise topic map.

Lessons continue to be learned through end user experiences, helping to clarify the best ways to represent and manage the topic map and to provide the most effective balance between general topic availability and the use of topics in specific situations and applications.

§ Using Topic Maps to create contextually relevant applications

Any discussion about context always returns to the idea of “common language” — or at least common ability to reference something — among computer applications and information resources. Increasingly, the data that describe user context is captured within the applications they are using.

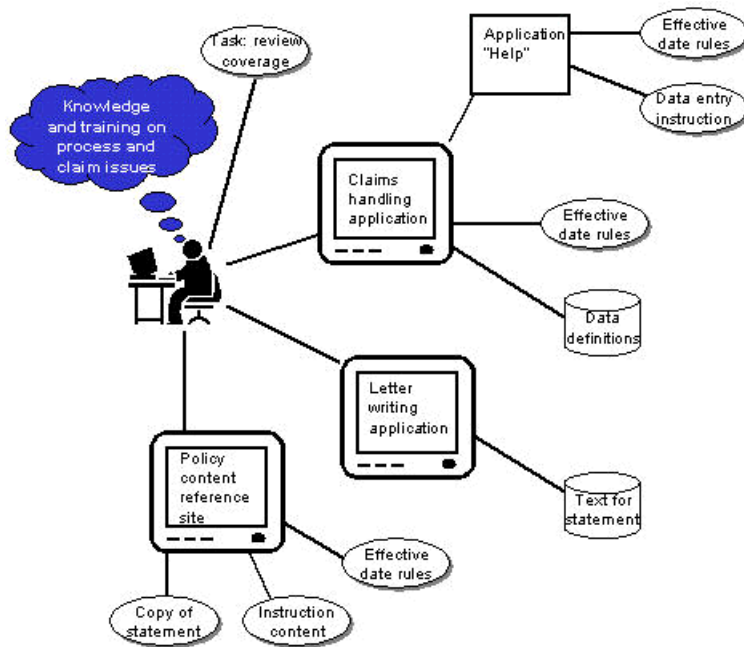
At the XML 2002 conference, Jean Paoli of Microsoft made the point that there is no longer any difference between documents and data — in XML, it's all data [Paoli, 2002]. This is a significant perspective shift, and it gives us food for thought. As designers of business applications who see the tremendous potential of topic maps, we are asking ourselves what roles topic maps play in the data world. In this section, we will briefly discuss ways in which we envision that topic maps might one day be used to create more flexible application architectures.

Providing an API/Web service to information resources

Imagine the efficiencies that can be gained when programming new applications if it becomes possible to adopt a standard form of API [Application Programming Interface] for sharing data and information resources, rather than having to define, develop, and publish one with each development effort. This is in part what the promotion of concepts like XML DTD/Schema and web services has been about. However, a genuine interface or service goes beyond standardizing meaning, because it also standardizes the approach to requesting information and validating that request. Applications that would rely on an API or web service driven by topic map definitions could be maintainable by business subject experts, while also being understood equally by software developers and information content providers.

Creating flexible, modular applications

If a topic is a representation of any subject that you can talk about [Vatant, 2001], then shouldn't topics be a useful way of representing common subjects that appear on many different government and business forms or computer screens? Is an “address” really a different object with a different definition and form from one computer application to another? What about “the date you became disabled”? One area that we have been exploring for some time is a more object-oriented approach to application and online form development, *i.e.*, an application where each interaction builds dynamically based on the context derived from the data that are known or have been entered previously, providing the appropriate level of data capture and user support as required at the time. A key feature of this is to be able to

Figure 16: Illustration of current relationships around a single topic


identify common topics across applications and to establish organizational ownership at the topic level rather than the application or page level.

Currently, there is a tremendous amount of redundancy within legacy applications and data entry forms, as illustrated below.

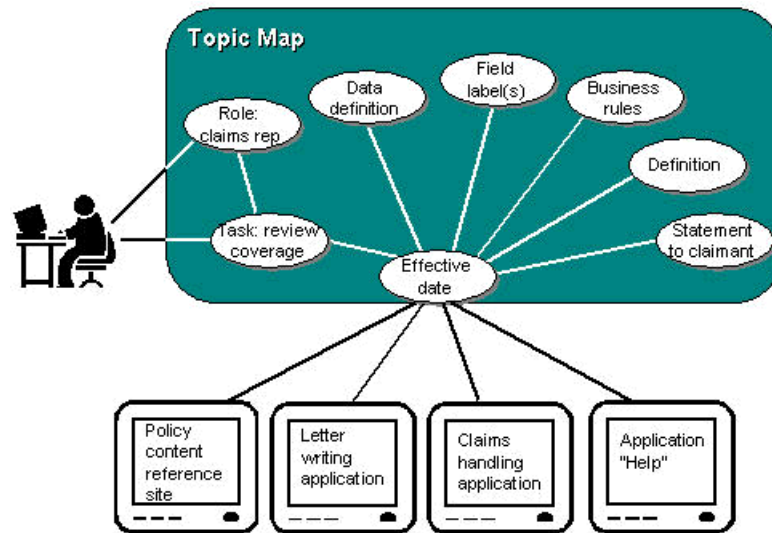
There are already established ways for depicting data flow and data entity relationships, but the language that they use is technical and is different from the language business uses. We're looking at how topic maps can provide a valuable way to begin to "normalize" the data, information, question flow, and interactions that users experience. They could potentially simplify the development of applications, as well, since the rules, branching logic, data definitions, and instructions could be managed from the same source.

Storing and maintaining business rules

Having worked on many applications in both large and small organizations, we are still surprised by the amount of reinvention and duplication that can be found in the coding of business rules and business logic. The IT industry faces a chronic problem with storing business rules in a way that reduces the risk of error when rules change and with the effort spent in development and maintenance. Approaches, models, and proposals have been written on rule language and schemes for databases and requirements [Ross, 1994] [Long, 1999].

One of the issues with maintaining business rules is that rules are, as the name implies, the province of the *business*, and should be maintained by business subject-matter experts to coincide with established business policies. Unfortunately, once they have been coded into applications, business rules typically are maintained by technologists and may not remain in synch with established policies. For example, consider the business rule that a claimant has 60 days in which to file an appeal. This business rule is related to the policy about what constitutes a valid filing of an appeal and to procedures that describe how it should be handled within and after the 60-day period. A subject-matter expert who updates the procedure may not know or be able to change the business rule coded in one or more applications. System architects generally agree that it is better to keep business rules separately maintainable from applications, rather than hard-coding them. As long as they are separate, why not associate them in a

Figure 17: Possible topic map approach to a single topic



topic map? This could allow the business subject matter experts increased control and visibility of business rules, and provide the advantage of linking the policy content to the application algorithms that implement it.

Even if the business rules themselves were not stored in a topic map, it could be useful to use topic maps to catalog the business rules, the applications that use them, and the documents that mention them to support impact analysis and change management.

Translating between business and technical language

Is it possible to use topic maps to bridge the gap between technologists and business people? The real issues are linguistic, not technical.

MI [Management Information systems] are often a classic case study of the data represented by the systems not matching the questions the business wants to ask. At the same time, the data warehouse industry continues to struggle with merging similar data from different database sources. Recently, we have encountered some projects looking at defining ontologies specifically for mapping multiple MI applications. Connecting related topics, merging topics, and providing different definitional views is something that topic maps do well. Beginning to harmonize the language required for MI is, in fact, a likely by-product of the other uses of topic maps discussed in this paper.

§ Conclusion

As the vision begins to take shape, it is clear that there are many things that remain to be worked out, not only within the organization but also within the tools and definitions of topic maps themselves. There are a range of subjects we believe need to be explored and considered among the topic map community generally.

- The challenge of visualizations, interactions, rules, constraints, and relationship management to make topic maps easier to construct and more maintainable by business subject domain experts
- The challenge of distributed ownership among people of various skill levels and awareness of the impact of changes to organization-wide topic maps
- The challenge of where to start

- The challenge of keeping projects small, while allowing an infrastructure to be developed that can grow with the abilities and interest within the organization
- The challenge of being understood, particularly by those stakeholders who are not familiar with (and shouldn't need to be familiar with) topic maps

Topic maps are potentially valuable to support high-level organizational consistency, *i.e.*, a “reference” ontology for the enterprise. At lower levels of detail, subject expert-specific and process-specific maps that associate or merge with the enterprise map, and yet have the required amount of flexibility to adapt to the user business needs, could be an answer to that delicate balance between taxonomic flexibility and powerful drivers for applications. Tying this syntax for organizational communication to the management information framework as a “lens” to view task performance, we begin to see the potential for a revolution in performance assessment and improvement.

Bibliography

- [Anderson, 1995] Anderson, J.R. (1995). *Cognitive Psychology and its Implications*, 4th edition. New York: W.H. Freeman.
- [Berners-Lee, Hendler, and Lassila, 2001] Berners-Lee, T., Hendler, J., and Lassila, O. (2001). “The Semantic Web.” *Scientific American*. May 2001. <http://www.scientificamerican.com/2001/0501issue/0501berners-lee.html>.
- [Beyer and Holtzblatt, 1998] Beyer, H., and Holtzblatt, K. (1998). *Contextual Design*. San Francisco, CA. Morgan Kaufmann Publishers.
- [Biezunski and Newcomb, 2001] Biezunski, M., and Newcomb, S. (2001). Specializing Occurrences in Topic Maps by Association Template Subclassing. In *Proceedings of Extreme Markup 2001*. Montréal, August 2001. <http://www.idealliance.org/papers/extreme02/html/2001/Biezunski01/EML2001Biezunski01.html>.
- [Degler and Battle, 2000] Degler, D., and Battle, L. (2000). “Knowledge Management in Pursuit of Performance: the Challenge of Context.” *Performance Improvement Journal* (EPSS Special Edition). ISPI, 39(6), July 2000.
- [Fisher, 2003] Fisher, Kathleen M. (2003). Prediction: A Profound Paradigm Shift. In *XML Topic Maps*. Park, J. (ed). Boston, MA: Addison-Wesley.
- [Hackos, 2002] Hackos, J.T. (2002). *Content Management for Dynamic Web Delivery*. John Wiley and Sons.
- [Hasher and Zacks, 1984] Hasher, L., and Zacks, R.T. (1984). “Automatic processing of fundamental information: the case of frequency of occurrence.” *American Psychologist* 39(12), 1372-1388.
- [ISO, 2002] International Organization for Standardization, ISO/IEC 13250:2000-2002, Information technology - SGML Applications - Topic Maps (ISO, Geneva 2000/2002).
- [Klein, 1998] Klein, G. (1998). *Sources of power: How people make decisions*. Cambridge, MA: MIT Press.
- [Landauer, 1995] Landauer, T.K. (1995). *The Trouble with Computers: Usefulness, Usability and Productivity*. Cambridge, MA: MIT Press.
- [Long, 1999] Long, J. (1999). A new notation for representing business and other rules. In *Semiotica Special Issue: Notational Engineering*. Long, J. (guest editor). Volume 125-1/3.
- [Newcomb, 2003] Newcomb, S.R. (2003). Conversation. March 20, 2003.
- [Norman and Draper, 1986] Norman, D. A., and Draper, S. W. (Eds.) (1986). *User centered system design: New perspectives on human-computer interaction*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- [Paoli, 2002] Paoli, J. (2002). Bringing the XML Vision to the Desktop With “Office 11.” In *XML 2002 Conference*. Baltimore, MD, December 2002. http://www.xmlconference.org/xmlusa/2002/plenaries_paoli.asp.

- [**Pepper, 2000**] Pepper, S. (2000). The TAO of Topic Maps. Ontopia: <http://www.ontopia.net/topicmaps/materials/tao.html>.
- [**Pepper, 2001**] Pepper, S., and Grønmo, G.O. (2001). Towards a General Theory of Scope. In *Proceedings of Extreme Markup 2001*. Montréal, August 2001.
- [**Ross, 1994**] Ross, R.G. (1994). *The Business Rule Book: Classifying, Defining and Modeling Rules*. Boston, MA: Database Research Group, Inc.
- [**Sigel, 2003**] Sigel, A. (2003). Topic Maps in Knowledge Organization. In *XML Topic Maps*. Park, J. (ed). Boston, MA: Addison-Wesley.
- [**Vatant, 2001**] Vatant, B. (2001). Binding Points for Subject Identity: The case for standard Published Subject Indicators. In *Proceedings of Extreme Markup 2001*. Montréal, August 2001.
- [**Vatant, 2002**] Vatant, B. (2002) From Implicit Patterns to Explicit Templates: Next Step for Topic Maps Interoperability. In *Proceedings of XML 2002 Conference*. Baltimore, MD, December 2002. http://www.idealliance.org/papers/xml02/dx_xml02/papers/05-03-06/05-03-06.html (password protected site).
- [**Winograd, 1996**] Winograd, T. (1996). *Bringing Design to Software*. New York, NY: ACM Press/Addison Wesley.
-

The Authors

Duane Degler

IPGems
11022 Two Ships Ct.
Columbia
MD
USA
21044
tel: 301-523-8819
ddegler@ipgems.com
<http://www.ipgems.com>

Duane Degler has over fifteen years' consulting experience, supporting a broad range of organizational performance through user-centered design, information/knowledge strategy, process analysis, international system implementation, training, and multimedia. His recent activities include supporting the Usability Center at the Social Security Administration as a designer and strategist, as part of a team from Lockheed Martin. Prior to that, his interface/system design work received three US awards. He managed pioneering multimedia projects in the '80s and spent much of the '90s in the UK involved in knowledge management research and consulting, contributing to his multi-disciplinary approach to interaction design.

Lisa Battle

Lockheed Martin
3300 Lord Baltimore Drive
Baltimore
MD
USA
21244
lbattle@acm.org

Lisa Battle has specialized in user experience design for over ten years. Her career began with the design of information resources such as online help, instruction manuals, and computer-based training, and progressed to electronic performance support systems and then to user interface design. She has consulted on interface usability and led the design of software applications and web sites for a variety of government and commercial clients. In her current position at Lockheed Martin, she is responsible for user-centered design for web-based applications at the Social Security Administration and mentoring project teams in user-centered methods.

Extreme Markup Languages 2003

Montréal, Québec, August 4-8, 2003

*This paper was formatted from XML source via XSL
by Mulberry Technologies, Inc.*