Interactive Information Retrieval Systems with Minimalist Representation

Eric Domeshek, Smadar Kedar, and Andrew Gordon

Institute for the Learning Sciences
Northwestern University
1890 Maple Ave.
Evanston, IL 60201
{domeshek | kedar | gordon } @ils.nwu.edu

Abstract

Almost any information you might want is becoming available on-line. The problem is how to find what you need. One strategy to improve access to existing information sources, is *intelligent information agents* — an approach based on extensive representation and inference. Another alternative is to simply concentrate on better *information organization and indexing*. Our systems use a form of conceptual indexing sensitive to users' task-specific information needs. We aim for minimalist representation, coding only select aspects of stored items. Rather than supporting reliable automated inference, the primary purpose of our representations is to provide sufficient discrimination and guidance to a user for a given domain and task.

This paper argues, using case studies, that minimal representations can make strong contributions to the usefulness and usability of interactive information systems, while minimizing knowledge engineering effort. We demonstrate this approach in several broad spectrum applications including video retrieval and advisory systems.

1. Introduction and Motivation

1.1. Data, Data Everywhere

In the Information Age, what the average user should be able to expect from a computer is useful information. What they currently get is a hard drive crying out for archeological excavation, an Internet connection with the discrimination of a fire hose, and WWW entree to a maze of twisty little passages all alike. What everyone wants is a bright knowledgeable personal secretary and research assistant. While some seek to approximate that end using statistical techniques, the more traditional AI approach is to construct a rich underlying representational and inference system which can reason about information needs and how those needs might be satisfied — in common parlance, intelligent information agents.

Work on information agents reflects certain assumptions: that the information glut is going to increase; that entropy being what it is, information disorganization is also going to increase; and that it is not our business to try and change the game. We can build competent agents for narrow tasks like finding people or organizations on the Internet, pulling statistics from known databases, and scheduling meetings. We can probably also build agents that broker services of simple agents. What remains intractable is building agents that solve broad problems using wide-ranging knowledge.

Instead of asking agents to tame the information explosion we might start by simply improving the organization and indexing of existing on-line information. All of us have it in our power to tidy our piece of the infosphere, and many organizations controlling larger resources have significant interests in making their information more accessible for particular purposes. Our initial observation, then, is that it is far easier to construct representational systems to be intelligible to a user and helpful in managing information, than to build them to support an autonomous information agent. Furthermore, given the limitations on our current ability to represent the nuances of what people really need when they seek information, interactive systems are likely to be more usable and more useful than autonomous agents.

1.2. Conceptual Indexing

Information retrieval systems based on keyword search have been shown to successfully search text databases. However, users often would like to search for concepts that are not directly in the text (e.g. the point of a story). Other times, a user may search for multimedia information items that are opaque to the computer (current approaches to parsing video or text (Hauptmann and Smith, 1995) are still quite limited in their ability to automatically and accurately reflect the content and meaning of the material). In such situations, *conceptual indexing* enables retrieval based on the underlying meaning of the information item (the semantics) rather than the multiple ways these concepts may be represented in the document (the syntax).

Conceptual indexing is an attempt not to fully *represent* an information item, but rather to *index* it. A conceptual index then is a specially designated label formally describing the most important aspects of the content or use of a memory item. Only selected aspects of an item are indexed for future retrieval — those aspects that capture the central meaning of the item, or that best indicate when the item is likely to be useful. Which aspects of an item should serve as its indices depends in large part on what kind of an information item it is, who is likely to retrieve it, and what purpose they are likely to be putting it to (McDougal, Hammond, and Seifert, 1991). As an example, in retrieving past experience to predict the outcome of new situations, it is usually good strategy to index on features of situations that are causally relevant to producing the outcome, and that are distinctive for that outcome. Likewise, how much detail is required in a conceptual index depends in large part on how many items are to be stored and how discriminating the consumer is likely to be.

Our goal in using conceptual indexing for interactive information retrieval is to take advantage of having a human user in the loop, and to invest in just the right amount of representation and inference to have a useful system. As a result, not only should the knowledge engineering effort be a lighter chore than representing the content of entire item, but matching should proceed more quickly and with greater assurance of relevance to the core meaning of the item. The question is: What degree of representation is required to meet these conditions? In our case, we have found that an extreme point on the representation spectrum — minimal representation and inference — can facilitate useful interactive systems.

This paper argues, using case studies, that simplified representations can reduce knowledge engineering costs yet still aid construction of useful interactive systems. Our strategy is to design representations whose level of detail and complexity is primarily driven by the demands of some immediate task. The value of this enterprise lies first in the immediate systems created, second in improved methods and tools for tailoring appropriate minimal representations, thirdly in the expanding repertoire of representational conventions, and eventually in shareable representations.

1.3. Case Studies

We illustrate the points in this paper using three sets of systems constructed over the last several years. Section 2 discusses ASK systems. Over the last six years at least two dozen of these *conversational hypermedia* systems have been built (Bareiss and Osgood, 1993). The key point about ASK systems is that they employ very small and simple representations. Generally each ASK system requires a simple domain model emphasizing the steps in accomplishing some task. In addition, all ASK systems share a simple cross-domain model of conversation (eight types of questions and answers that shape a conversation) (Schank, 1977). The representation is small, yet sufficient for the purpose of guiding interactive end-users for the task of 'talking' to the experts, and guiding knowledge engineers for the task of constructing these systems.

Section 3 discusses **Deja-Vu**, an on-line retrieval system for an archive of stock video. **Deja-Vu** is still a work in progress. In contrast to the many ASK systems with their many tiny domain models, the single **Deja-Vu** system has one very large (but still quite simple) domain model; its domain is "stuff happening in the everyday world", divided into five categories: places, people, things, activities and time. As with ASK systems, the main use for the minimal representation is to guide interactive users to find the information (in this case, the stock video footage) they need. In **Deja-Vu**, however, representations also support some limited inference.

Section 4 briefly discusses a set of systems collectively referred to as case-based aiding systems. Our chosen exemplars are case-based design aids (CBDAs) built using the **Design-MUSE** shell, and the **Abby** lovelorn advising system. While the focus of representational design is again conceptual indexing, and the approach is again minimalist, these systems provide examples of how this minimalist strategy can produce advances in the state of the art including some quite detailed representational systems.

Together the suite of systems cited in this paper illustrate how the representational approach advocated here can be carried on to varying depths, and contribute to the argument for the approach's general applicability and usefulness.

2. ASK-Systems: The Value of Very Small Representational Theories

ASK systems are collections of digitized video clips in which experts answer questions about experiences they have had in their area of expertise. The goal is to make it easier for other practitioners, especially novices, to benefit from the opportunity to converse with these experts. The two key insights underlying ASK system then are (1) that the "war stories" of experts are a valuable resource, and (2) that a natural way of organizing access to those stories is to loosely simulate the properties of a normal conversation. It is in achieving this organization that ASK systems depend on some minimal representation.

2.1. How ASK systems work

To a first approximation, there are two possible moves in conversation: you can establish a new topic for discussion, or you can follow up on an already active topic in one of several logical ways. ASK systems have an interface built around these two operations. In the ASK interface design, these are handled by two different kinds of screens called *zoomers*, and *browsers*.

The purpose of zooming is to establish a topic, initially assuming a null context. The way zoomers work is to offer a limited set of choices from a model of the domain. Usually zoomers are nested two deep before a user is given a choice of relevant expert stories. Since ASK systems are most often built to support a particular task (or tasks), the zoomer choices are usually structured around a task model. Zoomer screens typically dress up the domain model in graphic form to take advantage of the extra visual cues afforded by good graphic design.

The purpose of browsing is to allow normal conversational follow-up or to let the user home in on an answer precise enough that they can get on with their work. Each ASK system has a standard browser screen that organizes all possible follow-up stories into a stable and limited set of categories based on a simplified theory of conversational coherence. Again, the browser screen dresses these choices in a graphic form that spatially codes semantic relationships between the current focus story and possible follow-ups.

ASK systems are hypermedia systems — statically linked networks of stories. All of the connections between stories are established by human indexers as the system is built. The role of the representational theory is to delineate the kinds of connections that are worth offering.

2.2. Representation in ASK Systems

Three sorts of representation play a role in ASK systems. We have mentioned the domain model underlying the zoomer screens and the conversational model underlying the browser. In addition, during system construction, indexers evolve a domain topics hierarchy to aid them in constructing appropriate links. The topic space grows incrementally in response to the particular stories found in the corpus being indexed. Indexers repeatedly traverse and elaborate the hierarchy as they try to categorize each new story. Some elements of this categorization may eventually find their way into a visible zoomer domain model.

As mentioned earlier, the kind of representation appearing most often in the zoomer interface is a task model. Other types of models that have been used include organizational structures, physical object component part breakdowns, and historical time-lines. Since the primary purpose for any of these models is to guide users to appropriate stories, they need not be terribly detailed, and they need not support any autonomous inference. For instance, task models may only be elaborated enough to indicate typical sequencing and nesting of steps. One ancillary use for the zoomer models is worth noting: in training applications constant exposure these models serves to passively indoctrinate novices into a particular expert-sanctioned view of the domain.

Finally, the representation of conversational structure is built into the browser screen. The screen's graphic layout depicts eight conversational associative categories (CAC) (Schank, 1977) related to the story in the center of the "lotus." These eight have become the standard for ASK systems.[1] Logically they cluster as four pairs: *context* and *specifics*, *causes* and *results*, *analogies* and *alternatives*, *opportunities* and *warnings*. In a conversation focused on carrying out a task, these categories serve well to organize possible follow-up questions a user might want to raise. For any particular focus story, each category organizes a set of questions that other stories in the system answer.

2.3. Lessons from ASK Systems

After six years of development, ASK systems are now a replicable commercial product. Several evaluation studies have been completed. Evaluation indicates that users from high-school students to novice military planners can navigate ASK systems and find information successfully. Perhaps more importantly, customers rank them as effective corporate memories. For the purposes of this paper, the important point is that the role of representation in these systems is nothing like what is normally expected in an AI system. Yet minimal representation turns out to offer quite significant leverage. An effective system for an important class of domains and tasks is built by combining a small amount of systematic organized modeling with people's ability to interpret and reason.

These simple models examplify a reasonable and effective way to start development of more elaborate representations. For example, while ASK systems are already a success, current construction techniques do not scale. Ongoing research is aimed at automating the story linking process, and that research is looking at elaborated task models as a key component. Such elaboration, however, relies on data drawn from hand crafted ASK systems, receives funding only because of the success of the earlier ASK systems, and is builds on the initial simple models developed for existing ASK systems. These are all important senses in which representation work can be cumulative.

3. Deja-Vu: The Value of Broad but Shallow Representation Theories

The **Deja-Vu** system addresses a problem faced by video production facilities that turn to on-line archives to store their video holdings. Given an effective retrieval system, video producers can save thousands of dollars per clip by finding and reusing previously produced material. This section describes a solution to the video retrieval problem based on appropriate use of minimal conceptual indexing representation.

[1]Though the exact labels on these categories may vary a bit between systems (to better fit each task and domain), the underlying relations remain essentially the same. Within any one **ASK** system, the labels are completely stable.

The expected users of a stock video archive are video and multimedia producers who require the ability to perform fast and effective retrieval from an archive that may contain thousands of clips showing a broad range of everyday experiences. This retrieval task requires a representational vocabulary broad enough to describe the diverse contents, yet detailed enough to discriminate relevant retrievals from the huge number of clips that do not meet the user's needs. The size of the required representational vocabulary would pose an insurmountable challenge if we also required deep semantic representation of each term, but all we really need, as with ASK systems, is a way for users to reliably navigate through the available content to retrieve the clips they are looking for. **Deja-Vu**, offers both a representational vocabulary with minimal, unstructured categories, and a simple point-and-click interface for managing the storage and retrieval of stock video clips.

3.1. How Deja-Vu works

A professional video producer can see many different things in a video clip. A scene of a sailboat in a bay on a summer evening can be described by the movements of the camera, the quality of the recording, the emotions that it evokes, or even the points that the clip can be used to illustrate. For retrieval purposes, however, the single most useful representation for this clip is the concrete visual scene, including the *object* sailboat, its *place* on a bay, the *time* (a summer evening), and the *activity* of sailing. In **Deja-Vu**, each clip is represented using concepts drawn from five categories: things, places, time, activities, and people. The total representation for a clip is simply a set: e.g. {sailboat, bay, summer, evening, sailing}. This set acts as the index for the clip; if the user's request is a subset of the clip's index set, then the clip will be retrieved..

The representational vocabulary of the **Deja-Vu** system consists of thousands of concepts for the different people, places, things, times, and activities that make up the visual content of our testbed video archive. Much effort has gone into developing a semantic network to organize these terms. Each concept in the **Deja-Vu** system is a node in a directed graph linked by several meaningful relationships. Aside from hierarchical taxonomic and partonomic relationships, the network captures the associations between concepts of different classes, which are designed to be navigable by users during the retrieval process.

The principle which guided the development of the **Deja-Vu** interface is that it is better to offer choices than to use the Query-and-Search paradigm of traditional information retrieval. **Deja-Vu** borrows the ASK system Zoom-and-Browse model of information retrieval: first direct users to some broad topic in the video archive and then allow them to browse through the options presented to them. However, **Deja-Vu**, introduces a new twist on the Zoom-and-Browse interface: rather than browsing through video clips, users browse through the space of index terms, i.e. the semantic network of people, places, things, times, and activities.

As in ASK systems, **Deja-Vu** manages the complexity of the browsing space by structuring the browsing interfac in a manner that is congruent with the users' understanding of the domain. In particular, **Deja-Vu** structures the browsing interface around expectations about activities, based on the theory of Memory Organizational Packets (Schank, 1982). Each activity node in the semantic network is linked via associative links to the people, places, things, and times that people would expect to be involved in the activity. For example, the activity of sailing is linked directly to places like oceans and harbors and things such as sailboats and nautical charts. These types of associations, along with standard taxonomic and partonomic relationships, offer the user a stable and intuitive conceptual neighborhood around any concept which is the current focus of attention. Figure 1 shows a screen shot of the browsing interface with the concept *hospitals* as the item in focus.

As in ASK systems, one of the key concerns is developing appropriate ways to initially get users to zoom into some relevant part of the browsing space. **Deja-Vu** incorporates several different graphical zoomers which can be used to select an initial topic as the focus for browsing. Graphical zoomers for geographic place, classes of places, people types, and different times are provided. Some of these zoomers have two levels to them, e.g. selecting a country on the top level graphic of specific geographical places presents the users with a close up graphic showing the country's regions and major cities. In the design of these graphics, the goal is to present the user with many options per graphic screen by capitalizing on spatial organizations of each of the different zoomer's concept type. The zoomers provide an excellent means of getting quickly into a region of the conceptual space that is only a few clicks away from any target concept. This is particularly true of the zoomer for general types of places; because of the close

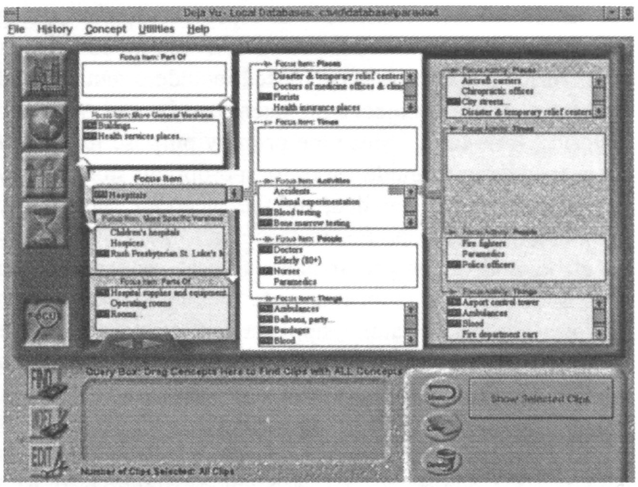

Figure 1: The Deja-Vu Browser Screen

connection between various activities and the places in which they are typically situated, this zoomer provides quick access to the majority of activity concepts in the system.

Quick access to index concepts through **Deja-Vu**'s zoom-and-browse interface allows for an effective form of incremental search. Each concept that is presented to the user on a zoomer or browser screen is displayed in a manner that indicates whether selecting the concept would result in the retrieval of any video clips. For example, in Figure 1, there are clips available for all terms with videotape icons next to them, including *ambulances, accidents,* and *doctors*. As the user selects concepts to flesh out their request, these availability indicators are recalculated so that only concepts that further reduce the retrieved set are available for selection. In effect, the selection of index concepts constructs a logical *and* function over a set of query terms that is always guaranteed to result in a positive number of retrieved clips. **Deja-Vu** constantly displays the number of clips selected by the current set of query concepts so users can incrementally decrease the number of retrieved clips by continuing to select different concepts. For example, selecting the concept *city streets* from the browsing interface shown in Figure 1 causes the system to report that 70 clips are assigned this index. Furthermore, videotape icons are removed from those concepts which are not represented in the 70 selected clips. Subsequently, selecting concepts which are marked with videotape icons can reduce the 70 clips down to a more manageable number.

3.2 Lessons from Deja-Vu

The **Deja-Vu** system has been implemented and is currently being used to archive the stock video library of the project sponsor, Andersen Telemedia. **Deja-Vu** was built using Borland's Delphi for Windows PCs. The current version supports multiple users working in a client-server environment, with concurrent retrieval, indexing, and editing of the archive contents. It currently has 3000 indexed video clips. As part of our iterative design, we performed a formative evaluation on a previous version of the interface with users who have done video production. The users found the system innovative and with potential to simplify the job of retrieving stock footage, and provided suggestions on the graphic design and labels used in the interface. Some prospect exists to move **Deja-Vu** out of the lab into a real video production environment for evaluation. Meanwhile, work focuses on developing Deja-Vu's representational vocabulary so that it will cover the majority of new video clips as they come in. The goal is a stable core semantic network that end-users will find easy to extend if necessary.

Deja-Vu's conceptual network was developed in a highly data-driven way. The major of concepts in the network were created by content analysts while indexing the first thousand clips of the sponsor's stock video library. The indexing effort has led to the development of effective indexing and concept-editing tools that share the point-and-click interface of the retrieval system.

4. Other Systems: Useful Component Theories for Complex Domains

So far, with ASK systems we have seen an example of very simple representations being developed to structure interactive systems, and with **Deja-Vu** we have seen broader, but still minimal, representational categories employed in a similar way (though with some facility for inference as well). In this section we illustrate how the minimalist approach to representation is also exemplified by several case-based aiding systems, and demonstrate how in two complex domains (artifact design and social interaction), that approach contributes to identifying novel, and useful representational categories.

4.1. Design-MUSE broadens the range of expressible design issues

Design-MUSE, a shell for constructing Case-Based Design Aids (CBDAs) (Domeshek, Kolodner, and Zimring, 1994), has been used to build several information repositories on different kinds of conceptual design. The most notable systems are **Archie** (Domeshek and Kolodner, 1993) — a CBDA for conceptual design of buildings — and **MIDAS** (Domeshek, Herndon, Bennett, and Kolodner, 1994) — a CBDA for early design of aircraft. Just as all ASK systems share a simple conversational model, so too all CBDAs share a simple model for browsing documentation and critiques of complex artifact designs. A CBDA organizes artifact design documents in a part and view network. It graphically links those *documents* to evaluative *stories* about the performance of the artifact. Those stories are in turn linked as exemplars to general discussion of recurring design *problems* and *response* strategies.

Beyond commitment to these classes of presentation (documentation, stories, problems, and responses) and the described connections among them, CBDAs also require a formalism for describing the possible interests a designer might have during conceptual design — what they might like to hear stories about. This is an area of representation that has not been much explored in AI applications to design. Construction of several CBDAs has demonstrated the general utility of characterizing designers' interests in terms of artifact *parts* (differentiated both physically and functionally) artifact *issues* (either structural, behavioral, or functional issues), and a further characterization of those issues in terms of when in the artifact *life-cycle* they arise, and which *stakeholders* they concern.

One interesting outcome of this work has been extension of the design issue vocabulary beyond concerns prototypically fitting the now standard trinity of structure, behavior, and function (SBF). For example **MIDAS** includes stories about the design of hydraulic lines that focus on the issue of contamination or gas inclusion; such deviations from expected operation are difficult to capture in a behavioral model. Likewise, there are stories focusing on the desirability of non-flammable hydraulic fluids, given a potential saving of lives to be traded off against increased weight and cost; these are difficult concepts to fit into a functional model of the aircraft. In **Archie** too, you find stories about issues that do not even remotely qualify under SBF: consider all the cultural, social, emotional, and aesthetic intentions that so actively shape the early stages of building design. Describing the full range of issues that designers care about during conceptual design of complex artifacts requires elaborating an extensive new vocabulary.

4.2. Abby deepens representation theory for the social world

As our final case study, we look briefly at the **Abby** lovelorn advising system (Domeshek, 1991). **Abby** is another case-based aiding system. Given a problem description it retrieves "Dear Abby" stories of similar problems that others had suffered (and often resolved) in the past. The main point of work on **Abby** was to detail the representations required to describe problematic social situations. In many ways, its representational distinctions for the social domain are among the most detailed yet developed in AI for this domain.

Abby elaborated a set of social actions and states, with particular attention to discriminating the fine structure of relationships between people and among people and organizations. Its detailed analysis of social relationships — both instantaneous snapshots and extended sequences — fed into detailed descriptions of the connection between such social context and particular goals and plans. All of this was complemented by descriptions of agents' attitudes towards goings on around them: how they felt about the existence of goals in which they were implicated, their volition with respect to the performance of actions, their expectations about the outcomes of their actions, and their intentions to affect motivationally important states. It was the average users' high degree of expertise in discriminating social situations that drove **Abby** to these extremes of representation. The result not only served as a conceptual indexing system for lovelorn advice, but has influenced the design of succeeding social simulations such as the **Yello** system (Burke, 1993).

However, **Abby**'s representational theory was exclusively a *component* theory: that is, it committed only to a taxonomized set of symbols and a set of combination rules for using them in accord with their intended semantics. It did not specify a full axiomization of the representation. To the extent that **Abby** elaborated a descriptive vocabulary but did not define a detailed inference structure, it fits the model of minimal representation described here.

5. Related Work

We look at three classes of related work: trends in knowledge representation and acquisition, the described systems' ancestors and contemporaries in the case-based reasoning community, and alternate approaches to multimedia information retrieval.

5.1. Knowledge Representation and Ontologies

Knowledge Representation (KR) is traditionally central to symbolic AI. However research in KR often focuses on the formal properties of syntactic systems, or on the detailed semantics of very fundamental concepts (e.g. time, space), rather than the study of particular content actually needed to perform some particular task (objects and their properties in specific domains). Concern for representing specific knowledge falls to those interested in Knowledge Acquisition (KA).

In the KA community there has been a growing trend towards exploring method-specific structures, called role-limiting methods (McDermott, 1988). Although less general than generic KR efforts, these methods provide more power in the representation by linking each piece of knowledge to the role it serves in the method. For example, SALT (Marcus and McDermott, 1989), a KA tool for generating "propose and revise" systems, uses the role of a "constraint" to guide the acquisition of knowledge from the user, to detect problems in the KB, and so on.

Even more recently a community has grown up around the idea of *sharable ontologies* (e.g. Gruber, 1993) whose goal is to represent content, but in a general enough way that the results can be reused. If everyone who had to represent some domain did a good enough job on the generally useful levels and made their results available, then over time people could start borrowing old solutions rather than rebuilding them from scratch. But it is exactly at the general levels that it is hard to recognize crucial inferences or code them in a way that will be useful across applications. The potential benefit of amortizing development costs across many applications tends to be lost because it is harder to do the job generally, and there may not be that many applications that actually accept the high-level axiomization.

There are some interesting ideas about how to improve the odds on building these sharable ontologies (Fikes, *et. al.*, 1994). Examples include developing 'seed' (starter) ontologies, providing authorship incentives for creating ontologies, providing tools to ease ontology creation, allowing for semi-formal specifications of ontologies (a mixture of free text and formal representation), and supporting on-line consensus-forming using the WWW to build a collaborative environment for distributed researchers (Farquhar, Fikes, Pratt, and Rice, 1995).

Although there is merit to the effort to share ontologies in the long term, in the short run ontologies will still be difficult to build, to agree upon, and to specify in directly reusable fashion. Thus, our bet continues to ride on the incremental development of systems each with a minimalist approach to representation. Such systems are easier to engineer, and produce tangible usable results. In the longer term, with the accumulation of such minimal ontologies both within a project, and across projects, more general ontologies may naturally emerge.

5.2. CBR / CB-Aiding

Case-Based Reasoning (CBR) (Schank, 1982; Kolodner, 1993) has been the inspiration and framework for all of the systems described here. CBR suggested two fundamental insights: First, if people often reason from experience and learn best from past experience, perhaps what really ought to be available on-line are stories of interesting experiences — records of how things happened and how things turned out; Second, a good way to think about organizing a corpus of such experiences is in terms of conceptual indexing. The first observation suggests that much of the potentially most useful information is not yet routinely available. The systems we cite here point the way towards getting that missing experience on-line.

In all these systems the kind of information being managed has the grain-size and character of a vignette capturing a situation usually to make a point about how to cope with related situations. Most often that means the unit of storage and retrieval is a couple of paragraphs of text perhaps with a picture or two, or thirty seconds of audio or video recording. We are not primarily interested in managing the sorts of simple facts that are tabulated in databases, nor are we aiming to catalog entire books, films, or transcripts as a library might. These are essentially new information repositories, and as such, focusing on building them right — on organization and indexing — makes even more sense than for the masses of random data already out on the net.

This is a common strategy in the CBR community where the emphasis has traditionally been on identifying features critical to determining the relevance of possible remindings. Systems have varied in the extent to which they concerned themselves with how indexes and retrieval cues would be generated. Relatively few systems have attempted to model a complete autonomous case-based reasoner — one that not only retrieved cases to help it cope with a new problem, but that also modified old cases and analyzed how well its solutions worked so as to store them appropriately for future reference (perhaps the best example would be **Chef**, (Hammond, 1989)). That sort of analysis requires much deeper understanding of a domain than is typically attempted. Many more systems, particularly those that have been built to really be used, have worked with minimal representation facilitated by reliance on a human in the loop.

5.3. Information Retrieval Techniques

The Information Retrieval (IR) community has made impressive strides in the development of statistics-based text retrieval systems. Notably, the INQUERY system (Callan, *et. al.*, 1992) has demonstrated that probabilistic models can be used for robust Query-and-Search text retrieval systems. Other researchers have developed interfaces that are analogous to the Zoom-and-Browse approach using solely statistical techniques (see Hearst and Pedersen, 1995). The success of these systems raises questions whether there is a real pay-off to hand-analysis and indexing of individual cases in a case library. While we recognize the great value of statistics-based techniques in the applications to which they have been applied, we argue that some circumstances will require more explicit representations of individual cases.

For some type of information, the text is not what is important. IR systems typically make the assumption that the appropriate retrieval conditions can be found within retrieved text itself. Aside from the obvious situation where the archive contains non-textual data (e.g. video or sound files), there are a number of instances where the text of the query is not likely to have any syntactic relationship to the text of the appropriate retrieval cases. Consider the difficulties involved in building a text-based retrieval system for the cases typically found in an ASK System. Often the most appropriate query for one of these stories would be a question that needs to be answered or a point that needs to be made, neither of which is likely to be explicitly stated in a story's text. Hand-analysis and indexing of ASK System stories allows system designers to sharpen the distinction between the contents of a story and the conditions under which it should be retrieved.

6. Conclusion

This paper advocates an extreme point on the representation spectrum: minimal representation and inference can facilitate construction of useful interactive systems. We demonstrate this using case studies: **Trans-ASK** uses a very minimal representation of domain topics and conversational categories to provide access to war stories from warriors responsible for logistics planning; **Deja-Vu** uses a minimal representation of the concrete contents of stock footage and standard expectations about the organization of such concepts to provide access to clips portraying scenes of the everyday world; **Design-MUSE** relies on a classification of design issues to index evaluations of existing artifacts; **Abby** relies on a detailed set of social and intentional descriptors to offer lovelorn advice by telling stories of other people's problems.

Representations in our system pay for themselves in the following ways: (1) Representations provide a controlled vocabulary both for indexing and for queries; the close correspondence between external form and internal representation removes most problems of interpretation. (2) Representations provide hierarchic organization(s) to make it easier for users to find appropriate descriptive terms. (3) In addition, other intuitive forms of conceptual organization can sometimes be made explicit and provided to the users as guides. (4) Finally, even limited representation can support certain simple classes of inference, and to ensure they execute efficiently.

For interactive systems where a human is in the loop, simplified representations can reduce the knowledge engineering cost yet still facilitate useful interactive systems. Our strategy is to design representations whose level of detail and complexity is primarily driven by the demands of some immediate task. The value of this enterprise lies first in the immediate systems created, second in the improving methodology and tools for tailoring appropriate minimalist representations, thirdly in the expanding repertoire of representational conventions which is an indirect form or sharing. Eventually it may lead to the sharable representations.

Acknowledgements

The Institute for the Learning Sciences was established in 1989 with support of Andersen Consulting. Special thanks to all the ASK system builders, and especially Ray Bareiss; to the VID group aided by Andersen Telemedia; and to GA Tech's Design Cognition team headed by Janet Kolodner.

References

Bareiss, R. and Osgood, R. (1993). Applying AI Models to the Design of Exploratory Hypermedia Systems. *Proceedings of Hypertext '93*.

Burke, R. (1993). *Representation, Storage, and Retrieval of Tutorial Stories in a Social Simulation*. Ph.D. Thesis. Technical Report #50, Institute for the Learning Sciences, Northwestern University.

Callan, J., Croft, W., and Harding, S. (1992). The INQUERY Retrieval System. *Proceedings of the 3rd International Conference on Database and Expert Systems Applications*, pp. 78-83.

Domeshek, E.A. (1991). Indexing Stories as Social Advice. *Proceedings of the Ninth National Conference on Artificial Intelligence.* pp 16-21.

Domeshek, E. A. and Kolodner, J. L. (1993) Finding the Points of Large Cases. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*. 7(2), 87-96.

Domeshek, E.A., Kolodner, J.L., and Zimring, C.M. (1994). The Design of a Tool Kit for Case-Based Design Aids. In J.S. Gero and F. Sudweeks (Eds.), *Proceedings of the Third International Conference on Artificial Intelligence and Design.* The Netherlands: Kluwer Academic Press, pp 109-126.

Domeshek, E.A., Herndon, M.F., Bennett, A.W., and Kolodner, J.L. (1994). A Case-Based Design Aid for Conceptual Design of Aircraft Subsystems. *Proceedings of the Tenth IEEE Conference on Artificial Intelligence for Applications.* Washington: IEEE Computer Society Press, pp 63-69.

Farquhar, A., Fikes, R., Pratt, W. and Rice, J. (1995). Collaborative Ontology Construction for Information Integration. Knowledge Systems Laboratory Department of Computer Science, KSL-95-63.

Fikes, R., McKay, D. Ballard, R., MacGregor, D., Saraswat, V., and Skuce, D. (1994). Tools Working Group Slides, *Working Notes of the Ontology Workshop*.

Gruber, T.R. (1993). A Translation Approach to Portable Ontology Specifications. *Knowledge Acquisition,* 5(2), pp. 199-220.

Hammond, K. (1989). *Case-Based Planning: Viewing Planning as a Memory Task*. Boston: Academic Press.

Hauptmann, A.G. and Smith, M.A. (1995). Text, Speech, and Vision for Video Segmentation: the Informedia Project. *Working Notes of IJCAI '95 Workshop on Intelligent Multimedia Information Retrieval*, pp. 17-22.

Hearst, M. A. and Pedersen, J. O. (1995). Revealing Collection Structure through Information Access Interfaces. *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, pp. 2047-2048.

Kolodner, J.L. (1993). *Case-Based Reasoning*. San Mateo, CA: Morgan Kaufmann.

Marcus, S., and McDermott, J. (1989). SALT: A Knowledge Acquisition Language for Propose-and-Revise Systems. *Artificial Intelligence*, 39(1), pp. 1-37.

McCarthy, J. (1968). Programs with Common Sense. In M Minsky (ed.), *Semantic Information Processing*. Cambridge, MA: MIT Press. pp 403-418.

McDermott, J. (1988) Preliminary Steps Towards a Taxonomy of Problem-Solving Methods. In S. Marcus (ed.), *Automating Knowledge Acquisition for Knowledge-Based Systems.* Boston: Kluwer Academic Publishers, pp. 120-146.

McDougal, T., Hammond, K., and Seifert, C. (1991). A Functional Perspective on Reminding. *Proceedings of the 1991 Workshop on Case-Based Reasoning*, pp. 63-76.

Schank, R.C. (1977). Rules and Topics in Conversation. *Cognitive Science*, Vol. 1. pp. 421-441.

Schank, R.C. (1982). *Dynamic Memory: A theory of reminding and learning in computers and people.* Cambridge, England: Cambridge University Press.