Seed Ontologies: growing digital libraries as distributed, intelligent systems

Peter Weinstein
Department of Electrical Engineering and Computer Science

Gene Alloway
University Library

NSF/ARPA/NASA Digital Library Project
University of Michigan
{peterw, cerebus}@umich.edu

Abstract

Ontologies are more than a particularly elaborate approach to the description and classification of information. They can be used to support the operation and growth of a new kind of digital library, implemented as a distributed, intelligent system. We describe the design and use of ontologies in the University of Michigan Digital Library. These ontologies will model all aspects of the digital library, including content, services, and licenses. We have refined and extended the IFLA hierarchy for the realization of work, and are starting to use ontologies to support reasoning about content search. We have also used the ontologies to classify the capabilities of computational elements of the system (agents), in a dynamic way that sustains functionality as new agents are added to the system.[1]

1 Introduction

We are developing *ontologies* of digital library content, services, and licenses, to support the evolution of distributed, intelligent information systems. Ontologies are highly expressive representations of objects, classes of objects, and relations in some domain. The central organizational notion is subsumption—relatively specific concepts are a "kind of" more general concepts—but concept definitions also include any number of many different types of relations to other concepts (see figures 1, 2, and 9 for illustrations). In the University of Michigan Digital Library (UMDL) [Birmingham et. al. 94], ontologies provide language that computational entities, called agents, use to describe their capabilities, which we call *services*. Agents interact in a system architecture that is "open", such that at some point we expect these systems to evolve in a highly decentralized manner, with momentum and ubiquity akin to the Internet itself.

We see our work as early steps in an inevitable transformation of information science. As participants in several recent workshops on digital library metadata, we have sensed a shared feeling verging on despair. The scope and quantity of information is expanding so rapidly that to classify it all within the traditional institutional structure is an overwhelming task. Furthermore, the variety of user needs, in combination with the complexity of the data, means that to agree on a single world-wide set of useful metadata standards is virtually impossible. The reception to ontologies in these workshops is often less than enthusiastic. It is true that ontologies are difficult both to develop and to apply. We argue, however, that when faced with a very difficult problem, the best way to make progress is to redefine the problem, leading to new strategies to achieve the desired ends by different means.

We predict that information science will be transformed in the following ways. Firstly, many more people, including non-professionals such as students, will classify and describe information. Quality assurance will be achieved with a variety of mechanisms, involving the use of intelligent tools, embedded in an institutional framework, that permit the efficient application of professional expertise. Secondly, there will be a proliferation of metadata standards, but these will map directly to extracts from ontologies. Ontologies will also proliferate, but within conventions and constraints that enable useful automatic translation between them. Thirdly, the expressiveness and precision of ontologies will support capabilities that are not currently available. For example, it will be possible to search for content with queries whose meaning is equivalent to a natural language sentence, rather than a combination of keywords.

The title of this paper, *seed ontologies*, is appropriate in several ways. A seed contains a genetic plan that constrains all aspects of a plant's structure and behavior. Similarly, our ontologies will integrate descriptions of all aspects of the creation, maintenance, and distribution of information. Like genes, ontologies will guide the growth of computational systems, which will develop to meet challenges posed by digital library environments. To enable this process, the ontologies are themselves designed to grow: additional terms can be added, and new ontologies can be defined in relation to existing ontologies. Finally, the title is also an imperative.

Compared to description systems developed in the library community, ontologies are more expressive, precise, and powerful. They are expressive because each term is defined in relation to other terms in a complex web (directed graph) of meaning (Gorman [92] makes a wistful plea for such a structure). They are precise in the sense of formal logic: the semantic denotation of each concept is a set of objects, and

[1] This work was supported by the NSF/ARPA/NASA Digital Library Initiative under grant CERA IRI-9411287.

relations denote sets of tuples. They are powerful because their precision supports automated reasoning. Note that ontologies can be used both to define sets of descriptive attributes (such as library systems ranging from the Dublin Core [Dempsey and Weibel 96] to MARC [MARC 94]), and as systems for classifying knowledge (such as Dewey Decimal or Ranganathan's facets).

Our research builds most directly on work in the field of knowledge sharing in artificial intelligence. In recent years, substantial progress has been made towards developing "knowledge interchange" languages (such as KIF [Genesereth and Fikes 92]). Ontologies defined in an interchange language can be automatically translated to numerous "representation languages", which are used by systems that do something with the knowledge [Gruber 93]. Ontologies are now developed as modules, that can be combined with other ontologies for reuse in different problem contexts [Farquhar 96]. An important use for ontologies is to enable communication between agents in distributed systems (e.g. as in the PACT system for circuit design [Olson and Cutkowsky 95]). In the UMDL, we advance this use of ontologies to support systems that have dynamic and evolving populations of agents [Weinstein and Birmingham 97]. By using ontologies to describe agent services, we permit agents to use services from new agents that they have never worked with before.

This paper describes the design and application of ontologies in the UMDL. Section 2 presents ontology concepts for describing digital library content; we have concentrated our definition efforts to date in this area. An interesting application of the ontology, however, is in the area of services. Section 3 describes the implementation of the UMDL as a "society of agents". We show how ontologies for agent services support a system that encourages outside organizations to develop new agents to participate in the system. Section 4 discusses some of the technical challenges that must be solved to realize our vision for digital libraries, and Section 5 concludes.

2 The UMDL Digital Library Ontology

This section reports on our progress towards defining concepts to describe digital libraries, with a primary focus on content. We discuss our design process, the concepts themselves, and the relationship between the major areas of the ontology. We also offer examples of how the content ontology will be used in the digital library, and consider its utility as a basis for relating multiple sets of metadata.

Few of our concept definitions have come easily. We spent hours considering alternative ways to conceptualize familiar concepts such as SONG, or DOCUMENT, and substantially more effort extracting order from slippery areas such as genre (we use upper-case to reference ontology concept definitions, as opposed to their meaning in the world). From this experience we derive two important lessons. First, we need to find ways to make ontology definition easier. We discuss this problem further in Section 4. Second, we believe that any attempt to impose a single conceptualization on all digital libraries is bound to fail. We do not seek one classification system to be used by everybody, but rather, a system that may be useful to many, and an approach that may be useful to many more.

Our concept definitions are structured as frames. When complete, a definition has a name, a natural language description, and a set of relations to other concepts. These relations should constrain the definition to express the meaning of the natural language description. In slightly different contexts, relations may alternatively be called roles, slots, or dimensions. Objects that are instances of the concept inherit its relations. Objects that instantiate associated concepts linked by relations are called "role fillers", or just "fillers". A variety of restrictions can be placed on filler values, such as numerical restrictions, or constraints in relation to other fillers. Also, an instantiation of a concept may be partial, with some roles filled, and others waiting for fillers.

Currently, the UMDL concept definitions are described in stylized natural language and with diagrams. We agree with Uschold and King [95] that these are the best mediums for developing initial definitions. Our natural language definitions can be viewed at
http://www.umich.edu/~peterw/Ontology/ontology.html.

Eventually, we hope to translate (more or less) automatically from our stylized natural language to KIF, and then, via Ontolingua [Gruber 93], to multiple representation languages. Our current practice is to translate manually, directly from stylized natural language to the representation language.

2.1 Work

The backbone of the UMDL ontology is a hierarchy of concepts adapted from a proposal by the International Federation of Library Associations and Institutions (IFLA) [IFLA 96]. We have sharpened their definitions, changed some concept names, and extended their hierarchy to meet the needs of digital libraries. We have also examined and incorporated insight from other metadata efforts, including US MARC, GILS, Dublin Core, and MCF.

The work hierarchy describes the stages of realization of work. The creative process starts with an idea. As long as it is immaterial, we call the work a CONCEPTION. After the work has been specified in some material form, it is an EXPRESSION. When the EXPRESSION is published, it becomes a MANIFESTATION. After electronic encoding, a MANIFESTATION is a DIGITIZATION. A particular copy of a DIGITIZATION is an INSTANCE. “Work” itself is not a defined concept, but refers to the full hierarchy. We believe this formalization is consistent with informal definitions of work by Miksa and Doty [94], and others in the library community.

Each concept in the work hierarchy inherits from its predecessor. Thus, an INSTANCE is also an DIGITIZATION, MANIFESTATION, EXPRESSION, and CONCEPTION. Concepts are in many-to-one relationships with their subsumed concepts: a single CONCEPTION can have many EXPRESSIONs, and so on.

The UMDL descriptions of these concepts are far more terse, but essentially equivalent to IFLA's. We have renamed IFLA’s WORK to CONCEPTION, and ITEM to either DIGITIZATION or MATERIALIZATION, for digital and traditional content, respectively. We have added INSTANCE to accommodate digital distribution. A summary is provided in Table 1.

Figure 1 illustrates the work hierarchy in relation to other concepts that give the work definitions meaning (relations that are not labeled indicate inheritance). One way to define EXPRESSION is as a CONCEPTION with a MODE (the most general concept in a set that describes a work’s genre). A MANIFESTATION is an EXPRESSION with a PUBLISHING-FORMAT, and so on.

Table 1: A hierarchy for the realization of work

Concept	Definition	Example
CONCEPTION	A concept, plan, or design for work.	An idea for a story.
EXPRESSION	Work with specified content.	A manuscript for a novel.
MANIFESTATION	An expression packaged in a publishing format.	A published edition of the novel.
DIGITIZATION	A manifestation encoded in a digital format.	The novel in SGML format.
INSTANCE	A particular copy of a digitization.	A copy of the SGML file.

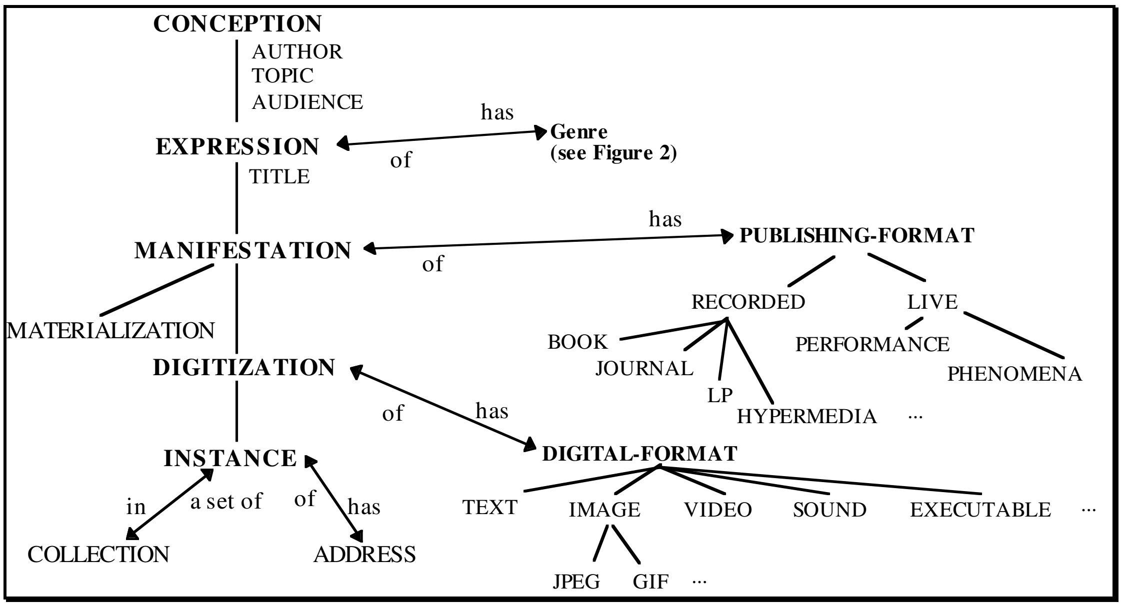

Figure 1: Ontological structure describing the realization of work

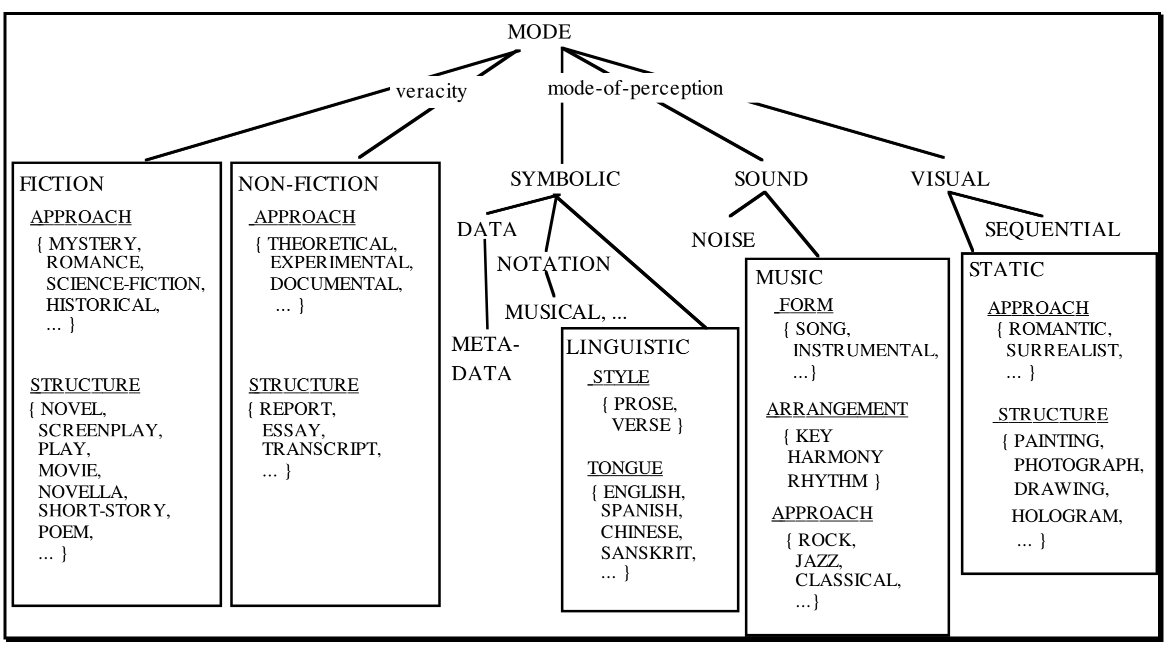

Figure 2: The structure of genre

Genre has been especially difficult to define, and we consider this area of the ontology to be provisional. We use "genre" as the Library of Congress does [MARC 94], with a much more general meaning than in common English. Like "work", no single definition can capture all of the meaning required; rather, "genre" refers to MODE and all of its subsidiary concepts. MODE has multiple dimensions. VERACITY divides FICTION from NON-FICTION, which are mutually exclusive. MODE-OF-PERCEPTION can be SOUND, SYMBOLIC, and/or VISUAL presentation: these are not mutually exclusive. Thus, an instance of an SONG can be FICTION, SOUND, and SYMBOLIC all at the same time.

Figure 2 illustrates the structure of genre. Only inheritance relations are shown. The full structure is not a tree, but a web of connections. For example, the concept SONG is defined as:

> SONG. MUSIC with lyrics, thus a LINGUISTIC-STYLE and a TONGUE.

Stating that SONG is MUSIC indicates a kind-of inheritance relation, while the references to other concepts identify roles.

Accurate description of digital library content requires some complications. Each concept may be instantiated in a composite form. Thus, a MANIFESTATION may consist of a set of MANIFESTATIONs (for example, a boxed set that includes several books).

Furthermore, many works transform or extend other, previous work. We call these connections ONTOGENY relationships, since they describe the origin and development of work. We distinguish between DERIVATION, which transforms the earlier work, and ADDITION, which extends it. Examples of DERIVATION include translations, abridgments, and summaries. Supplementary indices and sequels are examples of ADDITION. These categories are typically, but not necessarily disjoint. For example, a new edition may be both.

Every instance of an ontogenic relationship must be linked from a specific stage of realization of the preceding work, to some specific stage of realization of the subsequent work. Fortunately, the rule for identifying the correct stages of realization is very simple:

- Be as specific as possible when identifying the original work (for example, link from a MANIFESTATION rather than an EXPRESSION).
- Link to the most general stage of the subsequent work that is not shared by the preceding work.

Figure 3 illustrates two situations. In the first, (a), the subsequent work is realized from the same CONCEPTION, but has a different EXPRESSION. An example would be a supplementary index: since page numbers are specific to a PUBLISHING-FORMAT, the link's origin is at the MANIFESTATION level. The supplementary index's EXPRESSION has a MODE that is SYMBOLIC DATA, but its CONCEPTION is shared with the original work. The second situation, (b), shows an ONTOGENY relationship that would be appropriate for a sequel. These works typically deserve a separate CONCEPTION. This sequel refers to events in the preceding work that identify a particular EXPRESSION, but does not refer to a specific MANIFESTATION.

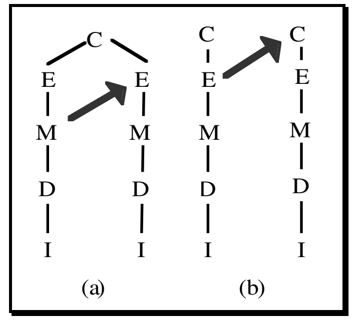

Figure 3: Placing ontogenic relationships in the work hierarchy

There is a long-standing issue in the library community regarding when to consider a new edition a distinct work, thus requiring a separate main entry (see Yee [94] for an in-depth analysis). The UMDL ontology requires judgments that are also subjective, but more detailed and technical. For example, the cataloger may need to decide between declaring a new, derived CONCEPTION, or a new EXPRESSION with a shared CONCEPTION. In either case, the system can reason about the connection between the works; thus, the cataloger's decision will perhaps be less critical.

We have conducted numerous thought experiments to test the work hierarchy and its auxiliary concepts, and have found our ontology to be robust and, for the most part, straightforward to apply. For example, we can describe:

- A digital version of *Through the Looking Glass* by Lewis Carroll, with the addition of special modern illustrations unique to the online version,
- A web page displaying real time data from a satellite run by NASA that is recording solar wind activity,
- A data set of GIS information, in Arc View format from the US Geological Survey,
- A collection of videos of Fiedler and the Boston Pops in concert,
- An interpretive photograph of the Mona Lisa—both the (digitized) original and the derived work have entries.

We are now in the process of creating ontology-based descriptions for a set of 500 US MARC records that have something to do with Beethoven. This knowledge base will support a query interface that will be the first application of the content area of the UMDL ontology to be empirically evaluated.

2.2 Linking Content, Services, and Licenses

Services are conceptualized as actions, which change the state of the world. Services are executed by an actor, for one or more actors, and performed in accordance with some agreement. Licenses are a form of agreement, granting authorization to an actor to provide a set of services.

Each stage of realization in the work hierarchy is an anchor for certain services and licenses. Services and rights associated with a work change as the nature of the work changes. Services associated with distribution or access will be suitable further along in the creative process, such as at the DIGITIZATION and INSTANCE stages. Specific licenses granting an individual a right to distribute will have to be negotiated for those aspects of the work alone. When creating a derivative work, a different type of license will be needed to give the new creator the right to base the new work on the preceding work.

In general, services operate on work. Licenses pertain to services, for work. These dependencies are illustrated in Figure 4.

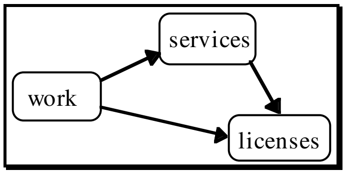

Figure 4: Information dependency between the major areas of the ontology

2.3 Inference

The precision of ontological definition permits extensive computational reasoning. Section 3 describes an important use for reasoning about kind-of relations. Inference rules can also be associated with concepts, leading to chains of deduction triggered when the concepts are instantiated. Table 2 describes some ideas for how this capability can be put to use.

2.4 Ontologies and External Representations

It is useful to think of metadata as external representations of ontologies. It seems one core set of metadata will continue to be too much, too little, or too vague for use across all forms of information goods and services. Alternatively, multiple sets of external representations can be derived, or otherwise related to, a single ontology. Their shared conceptual foundation will enable inference and translation among the metadata sets. If useful translation between ontologies is achieved, the benefits will transfer to metadata associated with the ontologies. Thus multiple representations are not only inevitable, but preferred as a solution to the needs of specificity and compatibility in metadata. We discuss translation between ontologies briefly in Section 4.

Table 2: Uses for inference in digital libraries

Use for inference	Goal	Example
Access	To optimize searching/browsing when only some information is known.	If you know a work is published in article form, the system can know to suggest resources which contain or index magazines, newspapers and journals, because they are collections of articles.
Description	To create more complete descriptions of resources from partial descriptions.	If the system knows the name of an author A of a novel N, and also knows that some manifestation of the novel was published in book form by publisher P, it can infer author A has been published_by publisher P, and that publisher P has_published_author A.
Rights	To identify which rights of a work are available for licensing, based on the stage of realization of the work.	If, as in the above example, novel N of the work W was published in book form by publisher P, then a set group of rights are automatically assigned by inference to the published_form book of novel N. The assignment of rights may perhaps be inferred from the standard license practices of the publisher P or author A.
Services	To suggest classes, or particular tools that are available for operating on work at a given stage of realization.	If the book edition B of novel N is in an image format, and there is no license for user U to print or alter novel N, then the set of tools available to user U to access the novel are view_image_only tools. Other tools allowing transfer or permanent alteration are unavailable.

3 Service Classification

Ontologies make a crucial contribution to a distributed, intelligent system architecture that may become a commonplace paradigm for large-scale systems in the next century. These systems, called *societies of agents*, have been the focus of research in artificial intelligence for over a decade. In this section, we show how we use ontologies to support the growth of societies of agents, and demonstrate that the precision and expressiveness of ontologies are the keys to their usefulness in this context.

Societies of agents are one kind of distributed, intelligent system. A distributed system includes programs that reside on computers that are geographically dispersed. By "intelligent", we don't mean a computer like HAL from 2001 Space Odyssey, but merely that the system demonstrates capabilities that are typically associated with humans: for example, the ability to plan procedures to meet explicit objectives, given a certain understanding of the state of the world. *Agents* are computer programs that communicate with other agents. Potentially, agents have autonomy to manage their resources to pursue their own ends. In other words, an agent is capable of "saying no" (another example of an intelligent capability). A society of agents consists of many agents that are capable of working together. When a user gives an agent a task, that agent finds other agents to work with to solve the problem. When one agent chooses to work with another, we call this selecting a *service*. Services are the capabilities that agents make available to other agents.

We are interested in societies of agents that are large, dynamic, evolving, and terminologically heterogenous. Societies are large if there are many agents participating, of too many kinds for any agent to know about all of them. Societies are dynamic if the availability of agents is in flux. In evolving societies, new agents appear and old agents change their behavior. If the agents use different language to describe equivalent services, or the same language to describe different services, then the society is terminologically heterogenous. The infinite variety of problems and contexts makes terminological heterogeneity inevitable—and indeed, desirable. After all, language is a tool; we need to invent new words and meanings to help us think about new problems.

We use the phrase *organic society of agents* to refer to societies that are large, dynamic, evolving, and terminologically heterogenous, to highlight the decentralized way that we expect these systems to develop. For obvious reasons, accurate selection of services in organic societies of agents is difficult.

We use *runtime classification of services* to support selection of services in societies of agents that are dynamic, and evolving. Figure 5 illustrates this process. Agents describe their services using terms from a set of nested ontologies. The figure shows three ontologies that are specific to digital libraries. The outer ontology conceptualizes content, services, and licenses generic to any digital library. The middle ontology includes terms that are specific to the UMDL. These include descriptions of auctions, goods, and other concepts that capture the UMDL's configuration as a computational economy [Mullen 95]. These ontologies are "static", in the sense that they are modified slowly, over time, by committees of persons. The inner ontology holds concepts that describe available agent services. We call this a *dynamic ontology*, because agents add concepts at runtime.

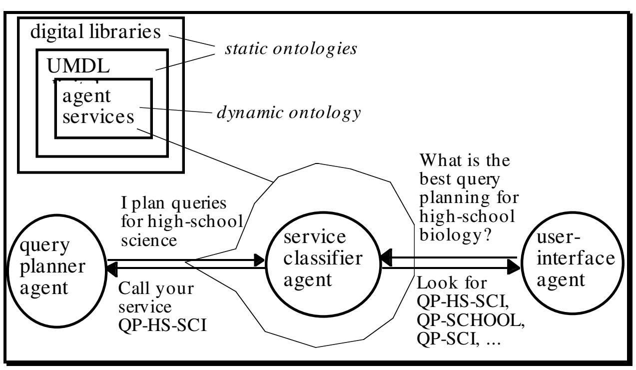

Figure 5: Runtime classification of services

The dynamic ontology is maintained by a Service Classifier Agent (SCA).[2] Agents use the SCA in a manner illustrated by Figure 5. To advertise, for example, a query planner agent tells the SCA about its service: that it recommends collections, and is especially well-suited for high-school science inquiries. The query planner tells the SCA that, if possible, it would like its services to be labeled with a certain tag. The SCA compares the service description with other service descriptions using a process called "automatic classification", which we will describe below. In this example, it finds that it does not already know about a service that is logically equivalent to the query planner's. Therefore, it accepts the query planner's request for a tag, and adds the concept to the dynamic ontology (the thin line around the SCA in Figure 5 has a bulge to indicate this addition to the agent services ontology). Subsequently, an agent working with a user to formulate a query submits a request asking the SCA for the best available service to meet its needs. The query has two parts: a description of the ideal service, and a search strategy. The ideal service in this example is query planning for high-school biology. No service specific to biology is available. Following the given search strategy, the SCA says the best choice is a service for high-school science, followed by other services in increasingly general order.

The SCA maintains a taxonomy of service concepts organized by subsumption. Logically, concepts denote sets of objects. Concept A *subsumes* concept B if every instance of concept B is also an instance of concept A. Subsumed concepts inherit the roles and fillers of their parents. In Figure 6, the Query-Planning service is the most general, subsuming all other services shown. The figure shows only subsumption relations, excluding many other links—a subsumption taxonomy can be thought of as the skeleton of an ontology. Nevertheless, the taxonomy in Figure 6 is not a tree: for example, there are several ways to get from Query-Planning to QP-School-Science.

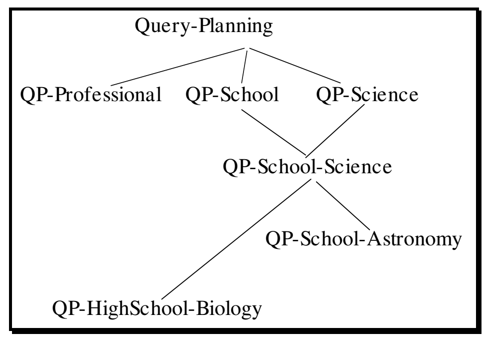

Figure 6: A subsumption-based taxonomy of services

To maintain the taxonomy of services, the SCA uses Loom, a description logic system in the KL-ONE family [MacGregor 91]. Description logics can do *automatic classification*—algorithmically placing new concept definitions into their proper location in the concept taxonomy. The judgment of whether one concept subsumes another is based on the structure and content of the concept definitions. For example, in Figure 7 two services are defined, both a kind of RECOMMEND-DLCOLLECTION. QP-SCHOOL is found to subsume QP-HS-SCI because every role in its definition corresponds to a role in the definition of QP-HS-SCI, with a filler value that subsumes the corresponding filler value. Therefore, automatic classification depends on a precise meaning for every relation in the ontology. See Woods [91] for a thorough discussion of automatic classification. Do not confuse this technology with "automatic classification" based on keywords, or other non-logical approaches (such as the methods described in Woodward [96]).

[2] An interactive demonstration of the SCA is at http://www.umich.edu/~peterw/Ontology/sca.html.

comparisons

```
QP-SCHOOL
(:and recommend-dlcollection
     (:all recommend-dlcollection.has.audience
          school)

QP-HS-SCI
(:and recommend-dlcollection
     (:all recommend-dlcollection.has.audience
          high-school)
     (:all recommend-dlcollection.has.topic
          science)))
```

Figure 7: Intensional check for subsumption[3]

In Figure 8, QP-HighSchool is a new service. Loom first finds one or more "most-specific-subsuming" (mss) concepts: those that are more general than the new concept, but not more general than another concept which also subsumes the new concept. Next, Loom finds one or more most-general-subsumed (mgs) concepts; then inserts the new concept in its proper location.

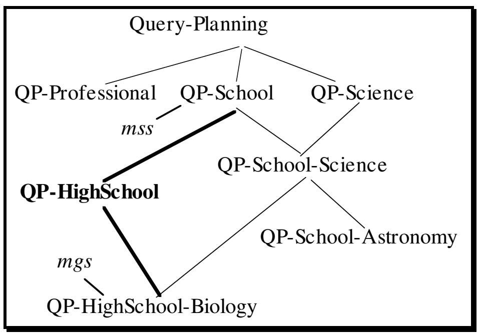

Figure 8: Automatic classification in a subsumption-based taxonomy

While automatic classification depends on the precision of ontologies, useful search depends on their expressiveness. Representing complexity in digital libraries, or any other "real-world" domain, demands expressiveness that is substantially more powerful than that of relational databases. Services have many descriptive dimensions, may be partially described at any level of granularity (with any combination of dimensions), and may be viewed from many perspectives (accessed by different sequences of dimensional values). For example, an agent might seek a service to recommend a collection of articles on volcanoes, for high-school audiences, to be contracted for via auction. Alternatively, the agent might seek an auction that sells a service to recommend collections. (In Figure 9, follow links either from "recommend-dlcollection" at the top, or "auction" at the bottom of the figure). Both perspectives can be represented simultaneously in an ontology, supporting retrieval for queries on any level of granularity, and from either perspective. Declarative formalisms with less expressiveness than ontologies, such as relational databases, force a commitment to particular combinations and orderings of dimensions.

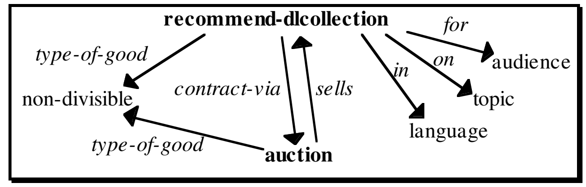

Figure 9: Multi-dimensional, multi-perspective complexity

Currently the SCA provides two search strategies: search by increasing generality (for the most-specific-subsuming service) or by increasing specificity (for the most-general-subsumed service). Relative generality/specificity of concepts is quantified as a function of path distance within the ontology. Search by increasing generality is useful when the service must be appropriate for all possible situations within the defined scope. For example, a science query planner can handle inquiries in both biology and astronomy. Search by increasing specificity is appropriate when the requesting agent doesn't care exactly which service is provided, as long as it is within the defined scope. For example, if you are searching for an image, you might not care if it is in GIF or JPEG format. In general, the classification taxonomy can be used to search for the best available service using any given quantified definition of "best" (i.e. a multi-attribute utility function).

The combination of declarative descriptions of services, with search for the best available service, provides an important benefit for societies of agents that are dynamic and evolving. *New agents can join the society and immediately be utilized for their services, without requiring modification or notification of existing agents*. This occurs if a) an agent requesting services states its needs as specifically as possible using the current ontologies, b) the new agent meets these needs better than existing agents, and c) the agent requesting services periodically repeats its search for the best available agent.

The UMDL is designed to encourage third party organizations to develop new agents to participate in the society. Runtime classification of services, based on ontologies, greatly reduces entry costs for new agents. Ontologies thereby contribute to supporting a style of system growth—decentralized, creative, vigorous, experimental, adaptive, energetic, haphazard but self-correcting—which has characterized the Internet itself in recent years.

4 Discussion

Eventually, we expect automated tools and institutional structure to be developed to provide a far higher level of service, using ontologies, than is provided today or could be provided tomorrow with less powerful approaches. Several difficult technical and organizational challenges, however, will need to be overcome to achieve this result.

Ontologies are precise compared to traditional metadata, and they are used for additional computational purposes. Thus, it is far more difficult to "work around" a definition that doesn't fit the situation at hand. When a definition is wrong, it will be necessary to address the flaw in a straightforward manner, by adding to the ontology, or to another related ontology, or starting a new ontology altogether. We therefore discuss the following relevant issues: the difficulty of creating new definitions, reuse and integration of existing sources of

---

[3] The service descriptions in Figure 7 are different in one important respect from those actually used by agents in the UMDL system. The Loom keyword :filled-by is used instead of :all, to define roles that have instance fillers rather than concept value restrictions. See Weinstein and Birmingham [97] for a discussion of the dichotomy between concepts and instances in description logics, and the problems that result for service classification.

ontological information, and the need to translate between ontologies.

Progress is being made towards developing knowledge acquisition tools for building ontologies. The Accord system [Kremer 95] is a good example. Accord provides a set of graphical user interface tools for developing “concept maps”, a form of semantic network that is less precise than ontologies. Then, Accord helps users transform these maps into increasingly formal representations. The system provides a form of “semi-automatic classification” to assist in the creation of increasingly precise concepts used to restrict fillers, and checks for consistency and completeness of definitions.

Knowledge acquisition tools that make a greater commitment to the structure of ontologies in particular could potentially have much greater impact than generic tools such as Accord. Such a tool would assist developers to make certain decisions that require a lot of time and effort. For example, it is often clear that a certain role is required, but it is not clear at which level of abstraction to attach it. If attached to a concept that is too general, the attribute is inherited where it isn’t appropriate; if attached too specifically, it is lacking where needed. A tool to permit “what-if?” testing of role attachment, highlighting the effects modifications have on associated concepts, would be very useful. Another common dilemma is the choice between creating a new concept, or adding a role with a value that implicitly indicates the concept. Usually, we find it most satisfactory to do both. It seems plausible that it would be possible to develop heuristics for measuring the “elegance” of an ontological structure, to assist decisions such as these.

The goal of the knowledge sharing community is to enable new knowledge bases to be created by composing existing ones, thus greatly diminishing their cost [Patil et. al. 92]. Currently, many professional communities have developed standard vocabulary definitions for their own use. Although there are numerous variations in the syntax used to define terminology, the degree of concordance is also striking. Lehmann [95] argues that it is possible, and desirable, to combine ontologies, thesauri, and sets of industry standards. Even if compatible syntax is assumed, however, such as within the library of ontologies at Stanford [Farquhar 96], it is not clear that extensive reuse of ontologies will be possible.

The problem is that meaning derives from context, rather than entirely from the word itself. It is not difficult to propagate extrinsic attributes using transitive inference rules, starting from the concepts for which they are intrinsic. What is needed, however, is a way to separate intrinsic and extrinsic attributes such that adequate definitions of concepts can be automatically recreated in different contexts.

For societies of agents to scale, we need to be able to translate between ontologies, so that local systems can develop their own terminology as needed for local problems. We believe that approaches that do "rough mapping" have the most promise. All of these methods start with some sort of overlap between the source and target ontologies. This overlap can be in the form of shared "typical" instances [Lehmann and Cohn 94], shared concepts, or shared parents from which terms in the source and target ontologies inherit [Weinstein 95]. Many of the rough mapping algorithms work by identifying correspondences between subgraphs in the source and target ontologies. There are important connections between these techniques and artificial intelligence research in analogy (see Owen [90] for a lucid overview of this work). We are currently investigating ways to use machine learning to induce rules for successful mapping in particular problem contexts.

Identifying appropriate constraints to structure the relationships between ontologies will be essential for the development of technology to translate between them. Inter-ontology organization will also be necessary for effective system behavior. For example, in large societies of agents, it is impossible for any agent, even a service classifier, to know about all agents in the society. To support society-wide searches for the best available service, ontologies of ontologies will prove useful, with meta-service-classifier agents to help select service-classification service.

5 Summary

We have portrayed the central role that ontologies can and should play in digital libraries implemented as societies of computational agents. Ontologies combine expressiveness and precision to a far greater degree than traditional systems for description and classification, and thus support computational reasoning that can be used in interesting and important ways.

The UMDL ontologies describe content, services, and licenses, with most progress to date in the area of content. A hierarchy of concepts modeling the realization of work forms the backbone of the UMDL conceptualization of content, providing anchors for concepts for services and licenses. This structure can be used to deduce characteristics of content being searched, and to identify operations that are appropriate and available to access content or manipulate it in other ways.

We have used the UMDL ontologies to organize declarative descriptions of services provided by agents in a distributed system. To advertise, agents define their services at runtime. Agents requesting services describe their needs as specifically as possible, then use the available agent that best fits that description. Therefore, new agents can be utilized immediately upon introduction to the system, without requiring any modification of agents requesting services. This capability encourages third party development of new agents, thereby contributing to the growth of the society.

Compared to traditional approaches, ontologies may offer a high performance alternative for describing and classifying information and services in digital libraries. They are difficult to create and apply, but we expect that future research will successfully overcome these obstacles.

Acknowledgments

Kevin Butterfield, Amy Warner, and especially Judy Ahronheim contributed substantially to the development of the UMDL ontologies. Bill Birmingham set up the problem and encouraged us to pursue it.

References

[Birmingham et. al. 94] Birmingham, W.P., K.M. Drabenstott, C.O. Frost, A.J. Warner, K. Willis. 1994. The University of Michigan Digital Library: This is Not Your Father's Library. Digital Library '94 Proceedings, June 1994, 53-60.

[Dempsey and Weibel 96] Dempsey, L. and Weibel, S. 1996. The Warwick Metadata Workshop: A Framework for the Deployment of Resource Description. *D-Lib Magazine*, July/August 1996.

[Farquhar et. al. 96] Farquhar, A., Fikes, R., Rice, J. 1996. The Ontolingua Server: a Tool for Collaborative Ontology Construction. KAW96. Also available as KSL-TR-96-26.

[Genesereth and Fikes 92] Genesereth, M. R., Fikes, R. 1992. Knowledge Interchange Format Version 3.0 Reference Manual, Report Logic-92-1, Computer Science Dept., Stanford University.

[Gorman 92] Gorman, M. 1992. After AACR2R: The Future of the Anglo-American Cataloguing Rules. Pp. 89-94 in Smiraglia, R. P. (ed.) *Origins, Content, and Future of AACR2 Revised*, American Library Association, Chicago.

[Gruber 93] Gruber, T. R. 1993. A Translation Approach to Portable Ontology Specifications. http://www-ksl.stanford.edu/yearindex.html.

[IFLA 96]. International Federation of Library Associations (IFLA). 1996. Study on the Functional Requirements for Bibliographic Records. Deutsche Bibliothek, Frankfurt am Main, Germany. http://www.nlc-bnc.ca/ifla/VII/s13/frbr/frbr.htm.

[Kremer 95] Kremer, R. 1995. The Design of a Concept Mapping Environment for Knowledge Acquisition and Knowledge Representation. Proceedings of the 9th Banff Knowledge Acquisition for Knowledge-Based Systems Workshop, pp. 23.1-19.

[Lehmann 95] Lehmann, F. 1995. Combining Ontologies, Thesauri, and Standards. IJCAI95 Workshop on Basic Ontological Issues in Knowledge Sharing, August 19-20, 1995, Montreal, http://www.csi.uottawa.ca/~doug/ontolWshop/ontolworkshop.html.

[Lehmann and Cohn 94] Lehmann, F.,Cohn, A. G. 1994. The EGG/YOLK Reliability Hierarchy: Semantic Data Integration Using Sorts with Prototypes, Proc. CIKM-94, Third International ACM Conference on Information and Knowledge Management, ACM Press, New York.

[MacGregor 91] MacGregor, R. M. 1991. The Evolving Technology of Classification-based Knowledge Representation Systems. pp. 385-400 in Sowa, John. F. (ed.), *Principles of Semantic Networks*, Morgan Kaufmann Publishers, San Mateo, California.

[MARC 94] Network Development and MARC Standards Office. 1994-. USMARC format for bibliographic data: including guidelines for content designation. Library of Congress, Cataloging Distribution Service, Washington D.C., USA.

[Miksa and Doty 94] Miksa, F. L., Doty, P. 1994. Intellectual Realities and the Digital Library. Digital Library '94 Proceedings, June 1994. http://abgen.cvm.tamu.edu/DL94/paper/miksa.html.

[Mullen and Wellman 95] Mullen, T., Wellman, M. P. 1995. A Simple Computational Market for Network Information Services. First International Conference on Multiagent Systems, June 1995. http://ai.eecs.umich.edu/people/mullen/papers.html.

[Olsen et. al. 95] Olsen, G. R., Cutkosky, M. R., Tenenbaum, J. M., Gruber, T. R. 1995. Collaborative Engineering Based on Knowledge Sharing Agreements. *Concurrent Engineering* V3 N2 p. 145 - 159.

[Owen 90] Owen, S. 1990. *Analogy for Automated Reasoning*. Academic Press, San Diego.

[Patil et. al. 92] Patil, R. S., Fikes, R. E., Patel-Schneider, P. F., McKay, D., Finin, T., Gruber, T., Neches, R. 1992. The DARPA Knowledge Sharing Effort: Progress Report. http://www-ksl.stanford.edu/yearindex.html.

[Uschold and King 95] Uschold, M., King, M. 1995. Towards a Methodology for Building Ontologies. Artificial Intelligence Applications Institute technical report AIAI-TR-183, July 1995.

[Weinstein 95] Weinstein, P. 1995. Matching Expressions with Differentiated Vocabulary: Technology for Knowledge Sharing in Partially Focused Communities. http://www.umich.edu/~peterw/abstract.vocab.html.

[Weinstein and Birmingham 97] Weinstein, P., Birmingham, W. P. 1997. Runtime Classification of Agent Services. Proceedings of the AAAI-97 Spring Symposium on Ontological Engineering. http://www.umich.edu/~peterw/papers.html.

[Woods 91] Woods, W. A. 1991. Understanding Subsumption and Taxonomy: A Framework for Progress. pp. 45-94 in Sowa, John. F. (ed.), *Principles of Semantic Networks*, Morgan Kaufmann Publishers, San Mateo, California.

[Woodward 96] Woodward, J. 1996. Cataloging and Classifying Information Resources on the Internet. *Annual Review of Information Science and Technology*, V31 pp. 189-220.

[Yee 94] Yee, M. M. 1994. What is a Work? Part 1: The User and the Objects of the Catalog. *Cataloging and Classification Quarterly*, V19 N1 pp. 9-28.