# Query-based Navigation in Semantically Indexed Hypermedia

*Daniel Cunliffe*
Hypermedia Research Unit,
Department of Computer Studies,
University of Glamorgan,
Pontypridd, CF37 1DL, UK.
Tel: +44 1443 482718
E-mail: djcunlif@glamorgan.ac.uk

*Carl Taylor*
DK Multimedia,
9 Henrietta Street,
Covent Garden,
London, WC2E 8PS, UK.
E-mail: cdtaylor@dkmm.co.uk

*Douglas Tudhope*
Hypermedia Research Unit,
Department of Computer Studies,
University of Glamorgan,
Pontypridd, CF37 1DL, UK.
Tel: +44 1443 482271
E-mail: dstudhope@glamorgan.ac.uk

**ABSTRACT**
This paper discusses an approach to navigation based on queries made possible by a semantic hypermedia architecture. Navigation via query offers an augmented browsing capacity based on measures of semantic closeness between terms in an index space that models the classification of artefacts within a museum collection management system. The paper discusses some of the possibilities that automatic traversal of relationships in the index space holds for hybrid query/navigation tools, such as navigation via similarity and query generalisation. The example scenario suggests that, although these tools are implemented by complex queries, they fit into a browsing, rather than an analytical style of access. Such hybrid navigation tools are capable of overcoming some of the limitations of manual browsing and contributing to a smooth transition between browsing and query. A prototype implementation of the architecture is described, along with details of a social history application with three dimensions of classification schema in the index space. The paper discusses how queries can be used as the basis for navigation, and argues that this is integral to current efforts to integrate hypermedia and information retrieval.

**KEYWORDS:** Hypermedia, semantic index space, semantic closeness, query-based navigation, cultural heritage, museums

## 1 INTRODUCTION
The conventional boundaries separating hypermedia from related disciplines are dissolving to yield a confluence of hypermedia, information retrieval, and multimedia database technologies [21, 27]. Traditionally, query-based retrieval methods have been considered a very different kind of interaction to the browsing commonly associated with hypertext or hypermedia. Queries require prior planning and a more analytical approach with greater cognitive overheads

for the user. Recent research in information retrieval (IR) has endeavoured to break down this barrier and build hypertext interfaces to IR systems. The separation of media items or information content from a link base is now widely accepted in hypermedia systems [22, 24]. In IR-oriented applications, an important distinction is that between the document space and some form of index space [18, 30, 2]. Terms in the index space are descriptors for media items in the document space—text passages, images, etc. Many hypertext-IR retrieval systems have been based on statistical analysis of terms automatically extracted from free text, [17, 4]; the index space is automatically created and 'vector space' similarity coefficients measure degrees of match between queries and media items, or between two media items. A different approach to IR is based on (usually) manually created index spaces, where semantic relationships exist between index terms [35, 25], and this is the approach we follow. In the cultural heritage domain, traditional practice includes the use of thesauri or classification systems as controlled vocabularies to index the media items in the collection. It is also possible to combine the two IR approaches [2, 47] with a thesaurus used to expand query terms.

One possibility that a semantic index space provides is an organised set of browsable nodes and links as a navigation aid to the associated layer of media items [1, 33]. The user can browse the index space, 'beam down' [12] to view media items of interest, and conversely 'beam up' to the index space from media items. The inclusion of semantic information in the index space, however, provides the opportunity for knowledge-based hypermedia systems that can provide intelligent navigation support and retrieval, with the system taking a more active role in the navigation process rather than relying purely on manual browsing. The navigation via query described here offers an augmented browsing capacity based on measures of semantic closeness between terms in an index space that models the classification of artefacts within a museum collection management system. This paper contributes to the integration of hypermedia and information retrieval from a hypermedia perspective that focuses on the notion of

navigation and hybrid query-navigation tools. All forms of navigation in our system are ultimately based on queries. We have been particularly interested in reasoning over the relationships in the index space (see also [15, 26]). Traversal of transitive semantic relationships makes possible imprecise matching between query and media item, or between two media items. This paper gives examples of such traversal algorithms in different types of tools and navigation techniques in both index and document spaces.

## 1.1 What is navigation anyway?

In the most extensive discussion of navigation to date, Marchionini [27] first distinguishes analytical from browsing strategies, and then identifies navigation as one of four browsing techniques. Traditional IR queries are considered an analytical strategy, requiring planning, greater cognitive overhead, more goal-driven than browsing techniques and more batch-oriented. Browsing strategies, on the other hand, are more heuristic, interactive, data-driven, and opportunistic (discovery as opposed to search). Marchionini recommends that retrieval systems should support both strategies. Much of the HCI and hypermedia literature equates hypermedia navigation with browsing. Marchionini however, distinguishes the four browsing techniques: Scan, Observe, Monitor, and Navigate with different degrees of interactivity, cognitive effort, goal-orientation and information structure. Navigation is characterised by high interactivity in a structured environment with the destination seldom pre-determined. Navigation is often a compromise between user and system responsibility; an incremental process with the user making choices from directions and feedback provided by the system. This aspect of discovery emphasised by Marchionini finds echoes in recent calls by museum professionals for a different approach to information retrieval in future public-access systems: "pathways to discovery instead of answers to questions ... Users seek information that is unknown, often not knowing what they want until it is seen" [38], a guide to "the serendipity of information" [39].

Another notion intrinsic to navigation is the idea of *movement* through an information space. The concept of a current position, or context, is crucial. The next navigational move is based on the results of previous navigations; hypermedia navigation is an example of an inter-referential mode of interaction [16]. This typically occurs when output (results of a query, destination of a link) forms part of the next input to the system. However in some cases, a previous input to the system is involved (navigation history, relevance feedback). Waterworth and Chignell [46] describe a model for information exploration with dimensions that distinguish target orientation (the degree of goal directedness in the user's cognitive attitude) from responsibility (the proportion of the responsibility for controlling the search shared by user versus system). Our query-based navigation corresponds to 'mediated browsing' in their model. They observe that rather than being dichotomies, these dimensions are continuum,

and that systems should offer a smooth transition between querying and browsing.

Hypermedia has tended towards the manual navigation of explicitly authored links, while IR has traditionally been oriented to the construction of analytic queries returning a ranked set of documents. Recent work in both disciplines has attempted to integrate these two approaches, although they still tend to appear as distinct phases of a user session. IR has gone a long way to incorporating browsing into essentially analytical strategies via the navigational interpretation of query results. If we consider hypermedia to have a different perspective, the issue is the incorporation of query into essentially browsing strategies. Modern hyperbase systems, where computed links can be implemented as queries on underlying linkbases, offer opportunities for rich and hybrid forms of query/navigation. For example, Garzotto et al [20] and Amman et al [5] have investigated how navigation can be based on database queries. Microcosm [24] supports IR free text queries (which do not require exact match) in its Computed Links, and various systems have made use of IR techniques to generate guided tours or trails [9, 23, 7]. Several applications now include content-based image matching as a component of navigation.

A query-based interpretation of navigation should allow augmented navigation tools to dynamically return the results of complex queries as link destinations, but retain the essential character of hypermedia navigation. This can be characterised by interaction that is relatively simple, does not require high cognitive effort, but involves user choice of navigational moves. A session will have a broad direction but not be so goal-oriented as to leave no room for discovery. Next moves will be based on the current position and context, but the query component can overcome an over-reliance on the immediate neighbourhood and context that some see as a disadvantage of navigational browsing [46].

## 2 CULTURAL HERITAGE INFORMATION SYSTEMS

Within the cultural heritage domain the advantages of a controlled vocabulary with semantic relationships between terms are well recognised [3]. Controlled vocabularies, such as ICONCLASS [11], and the Art & Architecture Thesaurus [32], already exist for large scale multimedia information systems. Due to the heavy investment of both time and expertise in such vocabularies and their role in professional practice, they will almost certainly have a part to play in future museum based multimedia information systems. For example, the ongoing European Aquarelle project makes use of a Semantic Index System, based on an object-oriented semantic model [14]. Such semantic schema provide a controlled vocabulary used by a human indexer to impose a classification on a set of objects, rather than automatically deriving classifications directly from the features of the objects themselves. Although manual classification is time

intensive, the skilled indexer can identify an object and place it within its wider context in a way that automatic content-based indexing alone cannot. In practice the classification of a complex media item, such as a photograph, is an expert task, and efforts must be taken to produce frameworks that encourage consistency in index term selection, although some variation can be compensated by intelligence in retrieval tools [29].

In general the vocabulary is organised into some form of semantic network, for example a hierarchy of broader/narrower terms and sometimes equivalent and related terms. This schema provides an appropriate frame of reference for both indexing and search. Museum information systems have typically concerned themselves with the management of collections, but there is a growing interest in public access systems which allow the non-expert museum visitor to obtain information about the museum, its collections, exhibitions, and so forth. The non-expert user may, however, be unfamiliar with the classification space, term vocabulary and organisation of the material and as a result of this, may have difficulty navigating through multi-faceted index spaces. Unlike tightly focused applications, cultural heritage information is often multi-purpose, with an element of discovery in its use. There is therefore a need to provide advanced navigation tools to assist exploration of such complex information systems. This paper is intended as a step in this direction, describing prototypes of techniques that can form part of such future tools.

## 3 THE SEMANTIC HYPERMEDIA ARCHITECTURE

The SHA [41] can be described in terms of three main layers: a media-base, a semantic index space, and an application layer. The application layer supports the user's interaction with the system, giving access to the set of navigation tools provided in a particular application and providing presentation capabilities for media items. The media-base contains the actual media items, such as photographs, sound recordings and so on, together with authored links between media items. To date within the project we have not been particularly concerned with the implementation of the media base, we simply assume that media items can be retrieved using a reference (to the media item as a whole—we do not currently support sub-component addressing).

Together with the media-base is a link-base of explicit (authored) node-to-node links. (We are currently investigating how to integrate the system with existing hyperbase technology.) However, the focus of our research has been the semantic index space. The index space holds the descriptors (index terms) of the information items in the media-base and the semantic relationships that give the terms meaning.

Both the index space and the explicit link-base are implemented on an underlying database, a form of Binary



**Figure 1: Semantic Hypermedia Architecture overview**

Relational Store with a restricted set of primitive relationships [10]. Thus, the index space consists of a set of binary relationships of the form:

$$<object\ relationship\ subject>$$

where *object* and *subject* are entities, such as media item IDs or index terms, and *relationship* is one of the semantic relationships permitted within the model, such as *A-Kind-Of*. A restricted set of core relationships have been implemented, including notions of hierarchy, aggregation, and instantiation. The database supports three basic operations, insertion, deletion and query, all at the level of a binary relationship triple. Insertion simply adds a triple to the store. Deletion and query are based on Standard Associative Forms [19]. These SAFs are essentially pattern matching queries on the database, using wildcards. Thus a SAF of the form:

$$<?\ A-Kind-Of\ Community\ Life>$$

where ? is a wildcard, will match all triples with the relationship *A-Kind-Of* which also have object *Community Life*. A SAF query with no wildcards is a test for the existence of a specific triple within the database. SAFs can be used for conventional query interactions with the database. They can also form the basis of the query-supported navigation described in Section 5.1. A constraint manager ensures the integrity of the relationships within the database.

The results described in this paper have been generated by a research prototype Lisp/HyperCard implementation. The database is a (memory-resident) collection of unordered triples, manipulated by LISP routines. A recent re-implementation uses a relational database with the constraint manager being implemented as stored procedures, and an expanded set of core relationships. An investigation into the appropriate table and index structures is underway.

### 3.1 Social history museum application

We chose the domain of social history to investigate a particular application of the SHA, a prototype interactive social history exhibit of the town of Pontypridd, mainly consisting of photographs from the archives of the Pontypridd Historical and Cultural Centre, with some text and oral history material.

For a subject index, we adopted SHIC—the Social History and Industrial Classification [37]. Clearly, other classification schemes could have been modelled in the SHA, however SHIC seemed appropriate for the subject area. SHIC is a museum standard which classifies objects according to their context of use, or the sphere of human activity with which they are primarily associated, as opposed to the more common type of object name classifications. It is used by many UK museums (and is currently being considered by the Aquarelle project) for cataloguing and collection management of social history material—we adapted it for public presentation. A small scale in-gallery evaluation of an early version that did not involve advanced navigation tools was conducted earlier in the project [43]. The four first level classification divisions, Community Life, Domestic and Family Life, Personal Life, and Working Life, form separate hierarchies (facets). Each of these terms is further subdivided to form a hierarchy of more specific subterms five levels deep. Four of these are modelled in the SHA, where SHIC terms are connected by semantic relationships to form a hierarchical subject index space (see Figure 2). Since photographs are composite objects, and an object can have several uses, multiple SHIC terms (sometimes from the same facet) were used to index each item, and an earlier version of the system allowed Boolean queries of SHIC terms [43].

Frequently a media item will be indexed by several terms, or combinations of terms used in a query. One interesting approach to the index space is to conceive it as a 'hyperindex' [12, 13]. If connectors between index terms have been defined, the term list becomes an 'index expression'. From this a 'power index expression' can be formed, a lattice-like structure which is the set of all sub-expressions. This lattice can be navigated ('query by navigation'); the user interactively broadens or narrows the search by specifying the current focus within the corresponding multi-term query. Arents and Bogaerts [6, 7] have extended this approach to semantically coupled thesauri with 'semantic hyperindexing' and computed links. Here a priori knowledge of the different facets (dimensions) in their subject thesaurus is used to create rules for meaningful (or interesting) combinations of terms from different facets; some combinations of terms are disallowed. These rules can be used to dynamically generate navigation possibilities, in particular valid links from the current node, or suggested trails. We do not currently have connectors between terms or make use of semantics associated with the combination of facets, however the SHIC traversal algorithm

yields degrees of similarity between sets of terms [45] and this is used in the navigation via similarity tool.

Empirical studies of user behaviour in the humanities and cultural heritage areas show that time and space are important access routes to information [8, 28]. Accordingly, we decided on an index space with three 'dimensions', representing the three primary categories we chose to model social history: subject (SHIC), time and space. Media items are indexed along the three dimensions. The temporal schema models time as a linear sequence of discrete, non-divisible elements—the smallest element of time being the year. These elements are combined to form higher level intervals each with a start year and an end year. Two classes of interval are used within the schema, calendar intervals and cultural intervals. Calendar intervals define a hierarchy of constituent subintervals, e.g. century, decade, year, and so on. Cultural intervals reflect a particular perspective on significant subdivisions of the timeline, e.g. the Victorian Era, or World War I. Temporal query operators, such as *pre*, *post* and *during* are available to the user and can be applied to both types of interval.

The spatial schema is based on the notion of a geographical area connected to other areas via semantic links; unlike Geographical Information Systems there is no geometric data model—in future work we wish to include co-ordinate information. The schema is made up of four components, compositional relationships, a geographic area type hierarchy, proximity relationships, and temporal relationships. Compositional *Part-Of* relationships are used to construct a regional hierarchy, where each area is composed of a set of non-overlapping sub-areas. Within the prototype the fundamental unit at the lowest level of the hierarchy is the street, although prominent landmarks can be defined as a composite part of a street. The geographic area type hierarchy provides a framework for the conceptual classification of geographical areas, e.g. towns, streets, bridges and so on. Although the compositional relationship allows the modelling of topological structure, it does not provide for the modelling of proximity relationships between geographic terms. To represent this, a *Next-To* relationship was added, allowing simple qualitative relations of proximity to be expressed. Temporal relationships are included within the geographical schema to model change in the definition or function of geographical areas over time [40]—a geographical term exists at a specific time.

### 4 SEMANTIC CLOSENESS MEASURES

Semantic closeness measures in the SHA are based on traversals of the relationships in the classification schema. Time and space are first class categories—each of the three dimensions of the classification schema has its own closeness algorithm. A brief description is given here, for a more extended treatment and references to related work, see Tudhope and Taylor [45]. The derivation of measures of closeness over semantic index spaces is a complex issue; a

variety of indications of closeness are possible. Consideration must also be given to the validity of the measures for the intended application and target users, the efficiency of the algorithms, parameterisation and tailorability of the measures, and the visualisation of the effect of a particular choice of parameters. For the purposes of the navigation tools described in this paper, the semantic closeness measures can be viewed as replaceable modules. The types of navigation described below do not depend on a particular implementation of semantic closeness.

Semantic distance between two terms in the conceptual schema is essentially based on the minimum number of semantic relationships that must be traversed in order to connect the terms [34]. Each traversal diminishes the semantic closeness measure by a given factor and the level of diminishment determines the generality of the terms that will be considered semantically close. Generally, the lower the value the more terms will be considered close. In addition to determining whether or not two given terms are semantically close, the semantic distance can also be used to retrieve the set of terms that are semantically close to a given term. Refinements of the semantic closeness measure have included different diminishment values for each type of semantic relationship and a weighting according to the level of specialisation of the traversal. This weighting is based on an intuitive belief that terms separated by a traversal at a high level of specialisation are likely to be closer than terms separated by a traversal at a lower level, effectively considering the granularity of classification at different levels within the classification schema. Therefore the diminishment value is determined by both the type of semantic relationship and the level of specialisation at which the relationship occurs. Measures of closeness between sets of classification terms have also been implemented.

A similar approach to semantic closeness is adopted for the geographical schema, traversals are over the Next-To relationships and again a diminishment factor is used. The temporal closeness function evaluates similarity between both years and temporal intervals. The measure is based on a weighted combination of the distance between the two temporal periods, the overlap between the periods, and the degree of separation between the periods. To measure similarity over the three dimensions of the schema, the semantic closeness measures are applied for each individual dimension and the results are intersected. A more satisfactory approach would be some form of unified similarity coefficient over the three dimensions [44].

## 5 NAVIGATION BY QUERY
A generic model for hypermedia navigation involves three elements: a source context, a destination context, and a transformation function [20]. Navigation can be viewed as a change in focus from the current context to another context. Conventionally, the contexts are media items and the transformation is achieved by the selection of an explicit link. In systems supporting computed links the media item itself is often not the context that is being transformed, rather links are defined via higher level abstractions, such as feature vectors or index terms. The resulting destination context may be a media item, but it could also be a new set of index terms. Queries can be said to be supporting navigation when the result of the query is mediated by the current location in which the request is framed.

Dynamic query based navigation is more flexible than fixed link navigation in that the results of navigation do not have to be pre-processed or fixed by the author of the hyper-media. Changes in the structure of the link space can be automatically reflected in the links presented to the user without the complex task of re-engineering the links. Query as navigation also frees the user from the need to follow navigational trails engineered by the hypermedia author, instead they are able to dynamically generate their own trails.

### 5.1 Query-supported navigation
In our system, all navigation at implementation level is based entirely on queries to the database, by means of the Standard Associative Forms (SAFs) defined in Section 3. The navigation tool in the application layer queries the database via the SAFs, and the resulting triples are passed back to the application layer. The triples can then be post-processed for expression in a particular navigation tool.

The queries to the database can be simple or complex. More complex queries result in the 'augmented navigation' described below. Conventional hypermedia navigation techniques, including both local and global browsers, guided tours, and Boolean queries can be implemented by relatively simple underlying queries [10]. For example, a hierarchical local browser, conventionally implemented by authored links, can be implemented using dynamic queries as the user progresses through the application. In order to implement guided tours, an additional relationship, Next, was defined to impose an ordering on the tour components. Although Next is essentially an explicit link between media items, the links are not embedded in media items but held in the link-base. Four of Trigg's [42] guided tour navigation functions were implemented by queries on the store: Start, Next, Previous, and Jump.

The user can browse through the SHIC classification (again, based on SAF queries) and collect SHIC terms, which can be combined with geographical and temporal terms into a combined query. A SHIC query can return all items indexed directly by a term and those indexed by specialisations of that term. The user can request the number of postings for a SHIC term (number of media items indexed by the term). All results, or destinations of navigational moves, are returned in the media viewer, through which the user can browse. An item in the Media Viewer can be used as the starting point for subsequent navigations. We have also

# DOCKET ALARM

# Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

## Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

## Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

## Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

## API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

### LAW FIRMS
Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

### FINANCIAL INSTITUTIONS
Litigation and bankruptcy checks for companies and debtors.

### E-DISCOVERY AND LEGAL VENDORS
Sync your system to PACER to automate legal marketing.