

Chapter 7

Visual Interfaces for Semantic Information Retrieval and Browsing

Katy Börner

7.1 Introduction

Today, an unimaginable amount of human knowledge is available in digital libraries, repositories, on the Internet, etc. Computers have become the entry point to a worldwide network of information, services, and other people. However, the Web is extremely unstructured and heterogeneous, making the efficient and effective discovery of relevant data, services, or expertise rather difficult. In sum, we are facing unparalleled challenges and opportunities for universal knowledge collection, access, and management.

Currently, various indexing systems, cataloguing systems, and search engines exist, but their ability to retrieve relevant data from diverse, distributed databases with different formats and multimedia data (text, images, audio, video) is very limited. Among the well-known problems for traditional information systems are the vocabulary inconsistency between user queries and data presentation (Bates, 1998), and the approaches simple keyword-matching usage of exploiting the frequency of co-occurrence of terms.

The vision shared by major information retrieval (IR) experts, as well as the developers of the new “Semantic Web”, is a network of distributed repositories where multilingual, multimedia documents can be searched within and across differently indexed and structured collections (Schatz and Chen, 1996). To make this vision come true, technologies are needed that search repositories despite not only variations in protocols and formats, but also variations in content and meaning. Several approaches to semantic information retrieval and browsing have been proposed and are reviewed in Section 7.2.

Visual interfaces exploit powerful human visual processing to ease information filtering and browsing. By connecting information to space, they also support browsing, i.e. the exploration of information space in order to become familiar with it and to locate information of interest. Browsing requires a working knowledge of both applied knowledge organization (typically alphabetical or hierarchical) and the means to navigate in it (Chen et al., 1998). Section 7.3 reviews related research on visual interfaces for digital libraries and information workspaces.

Section 7.4 explains the inner workings of the LVis Digital Library Visualizer. It details data analysis and data visualization and presents two prototypical

interfaces, the first to citation data from the Science Citation Index and the second to images from the Dido image library. The section concludes with a discussion on how to further adapt visual interfaces to the human cognitive and perceptual system.

The potential of the Semantic Web to improve semantic information retrieval and browsing is discussed in Section 7.5.

7.2 Semantic Information Retrieval and Browsing

Diverse techniques have been proposed to automatically summarize the semantics of a document set. Among them are symbolic machine learning techniques for automatic text classification; graph-based clustering and classification; linguistic approaches; ontologies; and statistics-based multivariate techniques such as multidimensional scaling, latent semantic analysis (LSA), self-organizing maps, and genetic programming. LSA as used in the LVis Digital Library Visualizer, linguistic approaches, and ontologies are reviewed subsequently.

The reader is invited to consult Börner, Chen and Boyack (2003) for an extensive review of diverse data analysis, clustering, and visualization techniques such as factor analysis, multidimensional scaling, latent semantic analysis, pathfinder network scaling, self-organizing maps, etc.

Genetic algorithms are based on principles of evolution and heredity and can be seen as general-purpose search methods that feature stochastic, global search processes. They have been applied for document clustering (Raghavan and Agarwal, 1987), to weight Boolean queries, and to improve recall and precision in information retrieval (Petry et al., 1993). Genetic algorithms have also been combined with neural networks to improve information retrieval (Chen and Kim, 1994).

7.2.1 Linguistic Approaches

Lexical ambiguity is a fundamental problem in IR. Many words are used ambiguously. Statistical approaches to linguistic annotation can help to disambiguate based on existing terminological resources in order to develop more satisfactory indexing techniques for IR:

... Linguistic annotation covers any descriptive or analytic notations applied to raw language data. The basic data may be in the form of time functions – audio, video and/or physiological recordings – or it may be textual. The added notations may include transcriptions of all sorts (from phonetic features to discourse structures), part-of-speech and sense tagging, syntactic analysis, “named entity” identification, co-reference annotation, and so on.¹ (Steven Bird, 2001)

Linguistic annotation can be done via segmentation, morphological analysis, syntactic tagging, and semantic tagging. First, the text is segmented into sentences, words, and other units (e.g. by punctuation). Each word is then analysed and annotated with its lexical description(s), typically the base form and

¹<http://morph ldc.upenn.edu/annotation/>

a morphosyntactic description of the word. Lexical information is then converted to syntactic tags which often depend on the application.

The Interspace project enables concept navigation across repositories using “scalable semantics” (Schatz, 1995). The system prototype semantically indexes real collections of millions of documents. It applies artificial intelligence techniques for phrase extraction and statistically correlates relationships between these concepts within the collection using context co-occurrence.

7.2.2 Ontologies

Ontologies establish a shared conceptualization between members of a community of interest. Examples of ontologies are catalogues, IR indexes, entity-relationship models in database design, dictionaries, thesauri from the computational linguistics community, object-oriented class definitions from the software engineering community, and so on. Ontologies can be represented as logic-based (first-order logic), frame-based (frame logic), or web-based (RDF, KIF).²

An example is the Unified Medical Language System (UMLS) Semantic Network³ under development by the National Library of Medicine as part of the Unified Medical Language System project. UMLS defines 134 semantic types that are linked by 54 important relationships in the biomedical domain and provides a consistent categorization of all concepts represented in the UMLS Metathesaurus.

The growth and dynamic range of digitally available data and the Web make it difficult to maintain manually created thesauri or ontologies that aim to cover a broad spectrum of dramatically changing topics. Automatic techniques promise to alleviate the need for time-consuming human analysis and avoid human limitations and biases. Consequently, the semi-automatic generation, mapping, and evolution of ontologies are hot research topics in Artificial Intelligence research. For example, Ding and Engels (2001) adopted co-occurrence theory to develop a lightweight domain ontology (see also Ding, 2001).

Heflin and Hendler (2000) developed SHOE⁴, a web-based knowledge representation language that supports multiple versions of ontologies. SHOE uses XML-like tags and agent technology to understand specific types of web pages (e.g. university course descriptions), to mark them up with semantic information, and to draw inferences over the collection for interesting search problems, such as student academic planning.

ScholOnto is an ontology-based DL server designed to support scholarly interpretation and discourse. It enables researchers to describe and debate, via a semantic network, the contributions of a document and its relationship to the literature (Shum, Motta and Domingue, 2000).

7.2.3 Latent Semantic Analysis

Latent semantic analysis (LSA), also called latent semantic indexing (LSI), is a vector space retrieval method that overcomes the so-called vocabulary mismatch

²<http://www.ladseb.pd.cnr.it/infor/Ontology/Papers/OntologyPapers.html>

³<http://www.nlm.nih.gov/pubs/factsheets/umlssemn.html>

⁴<http://www.cs.umd.edu/projects/plus/SHOE/>

problem (Deerwester et al., 1990; Landauer, Foltz and Laham, 1998). By considering the context of words, LSA overcomes two fundamental problems faced by traditional lexical matching schemes: synonymy (variability of word choice, e.g. “dad” vs. “papa”) and polysemy (a word with multiple meanings, e.g. “bank”, “crane”).

LSA uses an advanced statistical technique, singular value decomposition (SVD), to extract latent terms. SVD, a form of factor analysis, constructs an n -dimensional, abstract term-document space in which each original term and each original (and any new) document are presented as vectors.⁵ In SVD, a rectangular term-by-document matrix X is decomposed into the product of three other matrices: W , S , and P' with $\{X\} = \{W\} \{S\} \{P'\}$. W is an orthonormal matrix and its rows correspond to the rows of X , but it has m columns corresponding to new, specially derived variables such that there is no correlation between any two columns, i.e. each is linearly independent of the others. P is an orthonormal matrix and has columns corresponding to the original columns but m rows composed of derived singular vectors. The third matrix S is an m -by- m diagonal matrix with non-zero entries (called singular values) along only one central diagonal. A large singular value indicates a large effect of this dimension on the sum-squared error of the approximation. The role of these singular values is to relate the scale of the factors in the other two matrices to each other such that when the three components are matrix multiplied, the original matrix is reconstructed.

A latent term may correspond to a salient concept described by several keywords (e.g. the concept of information retrieval). Each document is represented by a number of those latent terms. The key innovation of LSA is to retain only the n most salient latent terms. All small singular values are set to zero and semantic relationships from the term-document matrix are preserved while term usage variations are suppressed.

The amount of dimensionality reduction (i.e. the choice of n) is an open and critical issue in the factor analytic literature. Ideally, n should be large enough to fit the real structure in the data, but small enough such that noise, sampling errors, or unimportant details are not modelled (Deerwester et al., 1990). Often, good retrieval and browsing performance is used as an operational criterion. The reduced dimensionality solution then generates a vector of n real values to represent each document. The reduced matrix ideally represents the important and reliable patterns underlying the data in X .

Similarity between existing documents and a new document/query can be computed on the basis of the three matrices as follows:

1. Multiply matrix W (its rows correspond to the terms [rows] of X) by the reduced matrix S .
2. Normalize resulting vectors (i.e. divide each vector by its length). This way all vectors are mapped onto a hypersphere.
3. Determine the document–document matrix of inner products that defines the similarities between all documents.

The dot product between two row vectors of the reduced matrix reflects the degree to which two documents have a similar pattern of terms. The matrix XX' is the square symmetric matrix containing all these document–document dot

⁵Readers interested in the existence and uniqueness of the SVD may consult (Deerwester et al., 1990).

products. Since S is diagonal and W is orthonormal, the following holds: $XX' = WS^2W'$. Thus the i, j cell of XX' can be obtained by taking the dot product between the i and j rows of the matrix WS . That is, if one considers the rows of WS as coordinates for documents, dot products between these points facilitate comparison between documents.

The construction of the SVD matrix is computationally intensive and is typically done in a pre-processing step. SVD on an m -by- n matrix is worse than quadratic in complexity on the smallest dimension. Yet, an effective dimensionality reduction helps to reduce noise and automatically organizes documents into a semantic structure more appropriate for information retrieval. This is a prime strength of LSA – once the matrix has been calculated, retrieval based on a user's query is very efficient. Relevant documents are retrieved, even if they did not literally contain the query words. The LSA matrix can also be used to calculate term-by-term or document-by-document similarities for use in other layout routines.

LSA has been successfully applied to information filtering and retrieval in Generalized Similarity Analysis (Chen, 1997, 1999), StarWalker (Chen and Paul, 2001), and the LVis – Digital Library Visualizer (Börner, 2000a; Börner, Dillon and Dolinsky, 2000) among others.

7.3 Visual Interfaces

A considerable body of research aims at processing, analysing, and visualizing large amounts of diverse online data.⁶ Given the typical user's inability to specify information needs, it is widely recognized that information retrieval is an iterative process. Most direct manipulation interfaces provide users with overview, zoom, and filter functionalities, as well as details on demand as proposed by Shneiderman (1996).

Sensemaking theory (Russell et al., 1993) as well as information foraging theory (Pirolli and Card, 1995) identify four main human activities for iterative information gathering: (1) searching for specific answers using keyword or relevance feedback searches; (2) browsing or exploring an information space without a specific question or representation in mind; (3) using overviews to discover areas of high information density or relevance and to ease navigation tasks in browsing and search; and (4) utilization of information workspaces to aggregate and assimilate relevant information. Users seem to shift frequently between these activities in response to the information gathered and the information task. Ideally, information visualization facilitates and supports all four activities.

7.3.1 Visual Interfaces to Digital Libraries

Visual interfaces to digital libraries aim to support search and browsing activities. Diverse data mining and visualization algorithms can be applied to extract salient semantic structures and/or co-citation relationships among

⁶Note that the analysis and visualization of the entire set of documents stored in a digital library is a special case of visualizing the retrieval result – namely, one in which all documents are retrieved.

documents and to layout documents spatially, helping users to visualize, locate, and remember documents more quickly. Several systems provide spatial cues by presenting documents in three-dimensional (3D) space and enabling the user to “walk” through this space.

Some of the first visual interfaces to DLs are: a two-dimensional (2D) map proposed by Orendorf and Kacmar (1996) as a method to structure DLs and their content and to ease document location and access; SemNet (Fairchild, Poltrock and Furnas, 1988); Cone Trees (Robertson, Mackinlay and Card, 1991); a spatial data management system (Herot, 1980), Populated Information Terrains (Benford, Snowdon and Mariani, 1995); VR-VIBE (Benford et al., 1995); and Bead (Chalmers, 1992, 1993), which uses statistical techniques to analyse and group documents based on their semantic similarity and to create visualizations of bibliographies. The section “Three-Dimensional Information Spaces”⁷ in Martin Dodge’s *Atlas of Cyberspaces* (2002). See also Crossley et al.’s (1999) Knowledge Garden provides an environment where people can meet colleagues and share relevant information.

Research by Chen (1999) aims to allow people to explore and navigate intuitively in a semantically organized information space. Latent Semantic Analysis and Pathfinder Network Scaling (Schvaneveldt, 1990) were applied to create a semantically organized information space providing access to 169 long papers from the ACM SIGCHI conference proceedings (1995–1997) and all papers from the ACM Hypertext conference proceedings (1987–1998). The approach was implemented in StarWalker, a system that uses Blaxxun’s community platform to display citation networks as a set union of all possible minimum spanning trees (Chen and Carr, 1999). StarWalker is the very first system that uses a tightly coupled spatial-semantic model as focal point in a multi-user environment. Chen et al. (1999) showed that the proliferation of information visualization (IV) models can play a significant role in extending and enriching the design of inhabited, multi-user virtual environments. In StarWalker, multiple users can simultaneously examine a complex visualization of a static data set. Social navigation is supported by the visibility of other participants and the ability to chat.

All systems mentioned thus far extract the semantic structure or co-citation networks of specific document sets. Because of the computational expense of the applied data analysis algorithms, the resulting visualizations are mostly static.

7.3.2 Information Workspaces

Research on information workspaces, as introduced by Card and his colleagues (1991), is related to visual interfaces to DLs as it refers to an environment, such as the Information Visualizer (Card et al., 1991), in which documents of interest can be compared, manipulated, and stored for future use. Diverse workspaces have been proposed. Among them are Balley’s (1994) Workscape, which enables the 3D spatial layout of documents; WebBook and Web Forager (Card, Robertson and York, 1996), which provides an information workspace for the WWW; the Cat-a-cone interface (Hearst and Karadi, 1997), which uses a 3D+ animated WebBook to visualize a moderate number of documents; PadPrints

⁷http://www.cybergeography.org/atlas/info_spaces.html

(Hightower et al., 1998), which uses the Pad++ zoomable interface⁸ to implement a thumbnail image-based web history mechanism; Data Mountain (Robertson et al., 1998), which is an alternative Web browser system that offers a generic 3D space where users can freely place thumbnail images of web pages for which they wish to keep references; the Task Gallery (Robertson et al., 2000), which extends the desktop metaphor into a 3D workspace represented as a long gallery with paintings symbolizing different tasks; and David Gelernter's Scopeware,⁹ which organizes documents chronologically. These systems differ in the way documents are organized (by users in piles, automatically in a list, partition, or hierarchy, etc.) and how documents are presented (as text, icons, image thumbnail representations, etc.). All the systems mentioned make use of spatial cognition.

However, visual interfaces thus far had only limited success. One reason for this is the fact that they do not provide the user-centred cues that help in remembering locations (Tan et al., 2001); see also the discussion in Section 7.4.4.

7.4 LVis Digital Library Visualizer

The LVis Digital Library Visualizer system was developed in interdisciplinary collaboration by Börner, Dillon and Dolinsky in 1999/2000. It uses LSA (see Section 7.2.3) to extract the semantic similarity of documents; applies clustering techniques and a so-called utility measure to select the partition that best reveals the semantic structure of a document set; and visualizes the result in a two-dimensional desktop screen or in a three-dimensional virtual reality environment. Two prototype interfaces have been implemented that are intended to ease browsing through retrieval results of citation data and image data. Subsequently, data analysis and visualization are explained in detail and examples of the two prototype systems are given. This section concludes with a discussion of how to increase the fit between human information seeking and available technologies.

7.4.1 Data Analysis

Data analysis comprises parsing the textual representations of documents (e.g. document author(s), title and source, addresses, abstract, keywords), analysing the semantic relationships between documents using LSA, and determining the best partition of the data set.

Latent Semantic Analysis: The SVDPACK by M. Berry (1993) was used to determine the semantic similarity between documents. The SVDPACK comprises four numerical (iterative) methods for computing the singular value decomposition of large sparse matrices using double precision. From this package, a sparse SVD via single-vector Lanczos algorithm was selected because of its comparatively low memory and processing requirements by low to moderate accuracies. In the case of digital libraries, terms may stand for words in abstracts

⁸<http://www.cs.umd.edu/hcil/pad++/>

⁹<http://www.scopeware.com/> by Mirror World Technologies.

or titles, author names, or cited references.¹⁰ Only author names, titles, abstracts, and cited references occurring in more than one document are included in the analysis.¹¹

Hierarchical Clustering: Nearest-neighbour-based, agglomerative, hierarchical, unsupervised conceptual clustering is applied to create a hierarchy of clusters grouping documents of similar semantic structure. Clustering starts with a set of singleton clusters, each containing a single document $d_i \in D$, $i = 1, \dots, N$; where D equals the entire set of documents and N equals the number of all documents. The two most similar clusters over the entire set D are merged to form a new cluster that covers both. This process is repeated for each of the remaining $N - 1$ documents. A complete linkage algorithm is applied to determine the overall similarity of document clusters based on the document similarity matrix. Merging of document clusters continues until a single, all-inclusive cluster remains. At termination, a uniform, binary hierarchy of document clusters is produced.

Selection of the Best Partition: The partition showing the highest within-cluster similarity and lowest between-cluster similarity is determined by means of a *utility* measure

$$utility = \frac{wSim}{wSim + bSim}$$

that contrasts the sum of within-cluster similarities $wSim$ with the sum of between-cluster similarities $bSim$.

7.4.2 Data Visualization

Documents are presented visually in a way that reveals their semantic similarities. Rather than being a static visualization of data, the interface is self-organizing and highly interactive. Data are displayed in an initially random configuration, which sorts itself out into a more-or-less acceptable display via a force-directed placement (FDP) algorithm (Battista et al., 1994; Fruchterman and Reingold, 1991). The algorithm works by computing attraction and repulsion forces among nodes. In our applications (see Section 7.4.3), the nodes represent articles or images, which are attracted to other nodes (articles/images) to which they have a (similarity) link and repelled by nodes (articles/images) to which there is no link. Initially, the nodes are placed randomly and are then moved according to the attraction and repulsion forces upon them until a local energy minimum is achieved. If the algorithm does not produce a visually acceptable layout, or if the user wishes to view the results differently, nodes can be grabbed and moved.

FDP algorithms cannot handle larger data sets. Therefore, we recently proposed *Semantic Treemaps* (Feng and Börner, 2002) to visualize the cluster structure of documents as well as the semantic relationships of documents in each cluster.

¹⁰In order to prevent articles with too large a number of citations or too great a length from having differential effects on the model, rows or columns of the term-by-document matrix can be reweighted. Weights can also be used to impose specific preconceptions about the importance of particular terms.

¹¹We do not exclude common words (like “a”, “and”, etc.), stem words, or map variants of words to the same root form.

7.4.3 Prototype Systems

Two prototype systems have been implemented using the data organization and visualization methods described above.

The first system visualizes query results of citation data from the *Science Citation Index Expanded* (SCI-EXPANDED) as published by the Institute for Scientific Information (ISI). This index provides access to current bibliographic information and cited references in more than 5,600 journals. The second system enables users to browse search results from the Dido Image Bank at the Department of History of Art at Indiana University. Dido stores about 9,500 digitized images from the Fine Arts slide library collection of over 320,000 images.¹²

Note that document retrieval can be done with existing search engines or based on LSA. Computationally expensive data analysis is done in a pre-processing step. Both systems are implemented in Java and run locally at Indiana University's School of Library and Information Science (SLIS).

7.4.3.1 LVis Interface to the Science Citation Database

Querying SCI-EXPANDED via the ISI Web of Science interface¹³ results in an often huge number of matching documents (maximally 500 are accessible) organized into lists of ten that can be marked, saved, and downloaded for detailed study.

To demonstrate a visual browser for this type of database, a query result data set named DAIV188 containing 188 articles matching the topic "data AND analysis AND information AND visualization" will be used. It was retrieved from SCI-EXPANDED in June 1999. The documents are represented in the usual Web of Science data output format (including author(s), article title and source, cited references, addresses, abstract, language, publisher information, ISSN, document type, keywords, times cited, etc.), as shown in Figure 7.1.

ISI data provides access to co-citation linkages. However, there are very few co-citation links in each retrieval result set. For example, exactly four articles in the set of 188 articles cite another article in this set. Therefore, LSA was applied over keywords and abstracts of articles to extract their semantic similarity. After clustering, the 167th partition was selected for visualization. It shows the highest utility value partition containing 19 clusters grouping from one to 53 articles.

Figure 7.2 shows all 188 articles. Book articles are represented by a rectangle. Journal articles are denoted by an oval. Each article is labelled by its first author. Lines between articles denote co-citation links.

The 2D layout of articles corresponds to the document-document similarity matrix as well as to the forces applied by the FDP algorithm to generate an acceptable layout. Cluster boundaries are represented by rectangles. Their colour denotes the level of similarity among articles (lighter colours denoting higher similarity). Each cluster is labelled by the keyword used most often in the documents it groups.

Users can click on articles to retrieve detailed information displayed in an additional Web browser window. Note that there are three clusters labelled

¹²<http://www.dlib.indiana.edu/collections/dido/>

¹³<http://webofscience.com/>

```

FN ISI Export Format
PT J
AU Small, H
TI Visualizing science by citation mapping
SO JOURNAL OF THE AMERICAN SOCIETY FOR
INFORMATION SCIENCE
LA English
DT Article
NR 46
SN 0002-8231
PU JOHN WILEY & SONS INC
C1 Inst Sci Informat, ...
ID SCIENTIFIC LITERATURES; COMBINED COCITATION;
WORD ANALYSIS; INFORMATION; RETRIEVAL; MAP
AB Science mapping is discussed in the general
context of information visualization. ...
CR *I SCI INF, 1997, NAT SCI IND DISK 198
AMSLER RA, 1972, UNPUB APPL CITATION
BALDI S, 1995, P 5 BIENN C INT SOC, P43 ...
TC 3
...
PY 1999
...
    
```

Figure 7.1 Partial ISI article representation.

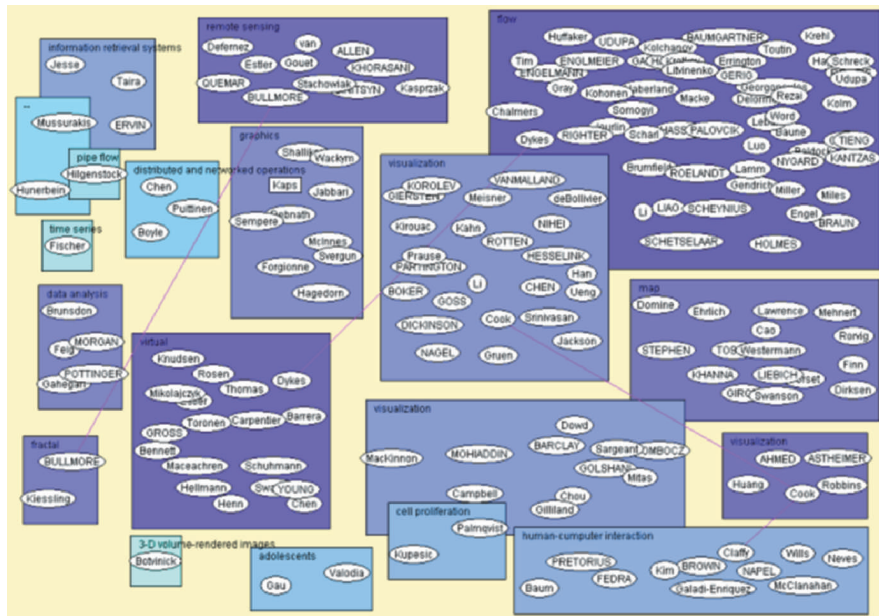


Figure 7.2 Java interface to DAIV188. (© 2000 ACM, Inc. Reprinted by permission.)



Figure 7.3 “Monet.Haystacks in the Snow” image and textual description.

“visualization”. Closer examination reveals that two of these clusters contain articles on visualizations in chemistry and physics. The third cluster groups articles that deal more generally with the analysis and visualization of information.

7.4.3.2 *LVis Interface to the Dido Database*

The Dido database permits convenient access and use of images for teaching and research purposes via a Web interface.

Each image in Dido is stored together with its thumbnail representation as well as a textual description (see Figure 7.3).

Upon querying Dido, a list of textual image representations is returned, based on string matching between search string and textual image description. Thumbnail representations can be clicked on to retrieve an image.

LVis was used to provide a visual interface to the retrieved image data sets. In a first step, SVD as well as the clustering techniques are applied to organize Dido images into clusters of semantically similar images. For demonstration purposes, the set of 21 images matching the keyword “MONET” was retrieved and is displayed in Figure 7.4.

The cluster contains 21 images including two portraits of Claude Monet drawn by Edouard Manet. Thumbnail representations of images fetched from the Dido database show some of Monet’s favourite themes such as haystacks, cathedrals, and water lilies. Each cluster is labelled with the word that occurs most often in the titles of its images – in this example the name of the artist.

In addition, a 3D CAVE interface has been implemented and evaluated (see Figure 7.5). Results are reported in Börner et al. (2000). An alternative 3D desktop interface to image data retrieved from Dido enables multiple participants to examine and discuss search results collaboratively (Börner, 2001a).

7.4.4 Finding the Perfect Match

Ideally, visual interfaces to digital data should match human perception and cognition capabilities as closely as possible. However, research is frequently “technology-driven” instead of “human-centred”.

For example, it is not clear how closely spatial layouts of documents produced by current data analysis and layout algorithms resemble human sorting



Figure 7.4 The MONET cluster. (© 2000 ACM, Inc. Reprinted by permission.)

behaviour. There are very few studies that compare the results of algorithms (in terms of document–document similarity or spatial data layout) with how human subjects organize documents. An inter-algorithm comparison based on the application of diverse algorithms to one data set was presented in Börner et al. (2003).

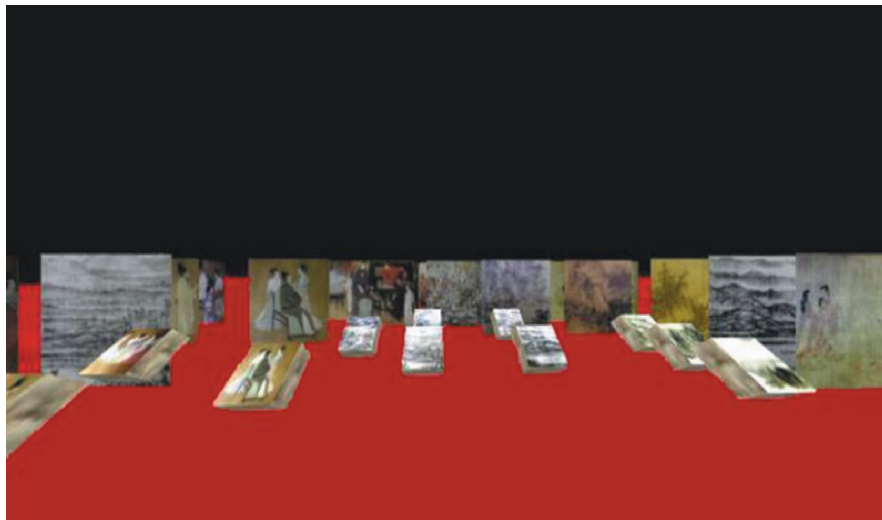


Figure 7.5 Chinese paintings displayed in a 3D virtual CAVE environment.

A comparison of results generated by LSA and by human sorting is reported by Börner (2000b). However, more detailed studies are required to determine and improve those algorithms to support human information-seeking behaviour.

So far concept-based (textual) image retrieval and content-based image retrieval (image recognition) have developed in parallel, with little interaction between the two groups of researchers and their literature (Enser, 1993). For example, LVis is limited to textual analysis – the textual representation of images is used exclusively to determine semantic similarity. LVis supports human browsing of retrieved images based on visual features. Browsing enables users to recognize rather than formulate a precise query in advance. This seems to be of particular relevance for information image retrieval. This combination of textual search and visual browsing was shown to be the most effective method for finding images of interest in a study by Chang, Eleftheriadis and McClintock (1998). Still, it seems to be desirable to combine text analysis and image processing techniques – potentially augmenting both processes by intermediate results to improve retrieval and browsing.

Spatial organizations of digital documents are supposed to ease document access, management, and relocation. However, the problems of spatial navigation in 3D (Darken, Allard and Achille, 1998; Darken and Sibert, 1996), particularly using 2D input devices, often outweigh the advantage of having an additional dimension. Abstract, large-scale information landscapes are especially difficult to navigate. The large scale requires the user to integrate the information provided by successive viewpoints into a coherent mental representation, often called a “cognitive map” (Chase, 1986), which is then used for navigation. Landmarks and their layouts have been identified as critical for navigation (Golledge, 1995; Lynch, 1960; Vinson, 1999); see also Thorndyke and Hayes-Roth’s (1982) work on differences in spatial knowledge acquired from maps and navigation. Darken and Sibert (1996) explored cognitive and design principles as they apply to large virtual spaces. Furnas (1997) explored view traversability and navigability for effective navigation through large data structures.

Experiments on navigation through 3D document visualizations of document databases have also been conducted by Vincow and Wickens (1998). Subjects navigated through a semantically-organized 1000-document information space; and the influence of frame of reference (egocentric, immersed perspective or more exocentric, bird’s-eye perspective) and use of map on document search and spatial judgement tasks were investigated. Performance was best with the exocentric display. Interestingly, in this experiment the map hindered rather than helped. Related research on frames of reference indicates that local information is best made available via an egocentric perspective, whereas global information should be provided by an exocentric perspective (Wickens et al., 1994).

7.5 Information Retrieval and Browsing in the Semantic Web: An Outlook

Currently, information in repositories, digital libraries, and on the Web is consumed and more or less successfully digested by humans. Computers have a hard time making “sense” of digital data; their ability to support humans via agent programs or direct manipulation interfaces, or a combination of the two, is rather limited.

The Semantic Web as envisioned by Berners-Lee and his colleagues (2001) “... is an extension of the current web in which information is given well-defined meaning, better enabling computers and people to work in cooperation.” It will provide access to human readable and machine understandable information. By using ontologies, different data formats can be implemented, translated, and utilized, making semantic search and browsing over different formats for diverse DLs possible.

However, the question of how to provide the necessary (shared or translatable) semantics to make it all work is still open. The most promising approach seems to be the creation of decentralized databases and their interconnection via “semantic translators”. The databases would be built by small communities (e.g. companies, scholarly communities, etc.) using vocabulary and semantics specific to their domain. The different databases could then be interlinked via “semantic translators” provided by specialized companies using semi-automatic approaches. Efforts such as the Stanford Scalable Knowledge Composition (SKC) project¹⁴ and the Bremer Semantic Translation project¹⁵ are already tackling the interoperability between different ontologies.

Exactly one semantic translator would be needed to connect a new database to a universal web of knowledge and services. A query would be translated among diverse databases and processed differently. This way, the Semantic Web would facilitate cross-disciplinary search, communication, and understanding and could “profoundly change the very nature of how scientific knowledge is produced and shared, in ways that we can now barely imagine” (Berners-Lee and Hendler, 2001).

In parallel to the development of the Semantic Web, computers are evolving from mere number crunchers to facilitators of human-human interaction and human-information access. Digital libraries can be viewed as managed and organized information spaces. Meaningful spaces evolve into places in which people can meet and access and manage information collaboratively. Visual interfaces to digital libraries introduce the concept of (semantic) proximity, which DLs and the Web have not provided, but which is successfully used in real-world libraries. Arranging documents in space according to their semantic proximity ensures easy access to related work and to people interested in the same topic (Börner, 2001b; Börner et al., 2002). In the near future, instead of examining long scrolling lists of matching documents, we may visit virtual, semantically organized data landscapes in which matching documents are highlighted for easy access and other users can be queried for suggestions.

7.6 Acknowledgements

We wish to thank the students involved in the LVis project: Andrew J. Clune, Ryan Schnizlein, Ho Sang Cheon, Kevin Kowalew, Sumayya A. Ansari, and Tyler Waite as well as Elin Jacob who gave very detailed and insightful comments on an earlier version of this paper.

We are grateful to Helen Atkins, ISI, and Eileen Fry from Indiana University for their insightful comments on this research as well as their support in using

¹⁴<http://WWW-DB.Stanford.EDU/SKC/>

¹⁵<http://www.semantic-translation.com/>

the Science Citation Index Expanded and the Dido Image Bank. Dave Pape is the software architect for XP, the underlying framework used to implement LVis in the CAVE. The SVDPACK by Michael Berry was used for computing the singular value decomposition. This research was supported by a High Performance Network Applications grant of Indiana University, Bloomington.

7.7 References

- Ballay, J. M. (1994). Designing Workspace: An interdisciplinary experience. *Proceedings of CHI'94 Conference*, Boston, MA, ACM Press, pp. 10–15.
- Bates, M. J. (1998). Indexing and access for digital libraries and the Internet: Human, database, and domain factors. *Journal of the American Society for Information Science*, 49(13), 1185–1205.
- Battista, G., Eades, P., Tamassia, R., and Tollis, I. G. (1994). Algorithms for drawing graphs: An annotated bibliography. *Computational Geometry: Theory and Applications*, 4(5), 235–282.
- Benford, S., Snowdon, D., Greenhalgh, C., Ingram, R., Knox, I., and Brown, C. (1995). VR-VIBE: A virtual environment for co-operative information retrieval. *Computer Graphics Forum*, 14(3), 349–360; <http://www.crg.cs.nott.ac.uk/~dns/vrvibe.html>
- Benford, S., Snowdon, D., and Mariani, J. (1995). Populated information terrains: First steps. In R. A. Earnshaw, J. A. Vince and H. Jones (Eds.), *Virtual Reality Application*, Academic Press, pp. 27–39.
- Berners-Lee, T., and Hendler, J. (2001). Publishing on the Semantic Web. *Nature*, 410, 1023–1024; <http://www.nature.com/nature/debates/e-access/Articles/bernerslee.htm>
- Berners-Lee, T., Hendler, J., and Lassila, O. (2001). The Semantic Web. *Scientific American*, 284(5), 43–43; <http://www.scientificamerican.com/2001/0501issue/0501berners-lee.html>
- Berry, M. (1993). *SVDPACKC (Version 1.0) User's Guide*, University of Tennessee Tech. (CS-93-194).
- Börner, K. (2000a). Extracting and visualizing semantic structures in retrieval results for browsing. *Proceedings of ACM Digital Libraries*, San Antonio, TX, 2–7 June, pp. 234–235, ACM Press.
- Börner, K. (2000b). Searching for the perfect match: A comparison of free sorting results for images by human subjects and by Latent Semantic Analysis. *Proceedings of Information Visualization 2000, Symposium on Digital Libraries*, London, England, 19–21 July, pp. 192–197, IEEE Press.
- Börner, K. (2001a). Adaptation and evaluation of 3-dimensional collaborative information visualizations. *Proceedings of Workshop on Empirical Evaluations of Adaptive Systems, 8th International Conference on User Modelling*, Bavaria, Germany, 13–17 July, pp. 33–40, Lawrence Erlbaum Associates.
- Börner, K. (2001b). iScape: A collaborative memory palace for digital library search results. *Proceedings of International Conference on Human-Computer Interaction*, Smith, M. J., Salvendy, G., Harris, D., and Koubek, R. J. (Eds) Vol. 1: *Usability Evaluation and Interface Design*, New Orleans, LA, 5–10 August, Lawrence Erlbaum Associates, pp. 1160–1164.
- Börner, K., Chen, C., and Boyack, K. (2002). Visualizing knowledge domains. *Annual Review of Information Science and Technology*, 37 (to appear).
- Börner, K., Dillon, A., and Dolinsky, M. (2000). LVis – Digital Library Visualizer. *Proceedings of Information Visualization 2000, Symposium on Digital Libraries*, London, England, pp. 77–81, IEEE Press.
- Börner, K., Feng, Y., and McMahon, T. (2002). Collaborative visual interfaces to digital libraries. *Proceedings of the Second ACH+IEEE Joint Conference on Digital Libraries*, Portland, Oregon, USA, 14–18 July (to appear).
- Card, S. K., Robertson, G. G., and Mackinlay, J. D. (1991). The Information Visualizer: An information workspace. *Proceedings of ACM Conference on Human Factors in Computing Systems (CHI '91)*, ACM Press, pp. 181–188.
- Card, S. K., Robertson, G. G., and York, W. (1996). The WebBook and the Web Forager: An information workspace for the World-Wide Web. *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '96)*, Vancouver, Canada, 13–18 April 1996, ACM Press, pp. 111–117; <http://www.acm.org/sigchi/chi96/proceedings/papers/Card/skc1.txt.html>
- Chalmers, M. (1992). BEAD: Explorations in information visualization. *Proceedings of SIGIR '92*, Copenhagen, Denmark, June 1992, ACM Press, pp. 330–337.
- Chalmers, M. (1993). Visualization of complex information. In L. J. Bass, J. Gornastaev and C. Unger (Eds.), *Third East-West International Conference on Human-Computer Interaction (EWHCI '93)*, Lecture Notes in Computer Science 753. New York: Springer Verlag, pp. 152–162.

- Chang, S.-F., Eleftheriadis, A., and McClintock, R. (1998). Next-generation content representation, creation and searching for new media applications in education. *IEEE Proceedings, Special Issue on Multimedia Processing and Technology*, 86(5), 884-904.
- Chase, W. G. (1986). Visual information processing. In K. R. Boff, L. Kaufman and J. P. Thomas (Eds.), *Handbook of Human Performance, Vol. 2: Cognitive Processes and Performance*. Chichester: Wiley.
- Chen, C. (1997). Tracking latent domain structures: An integration of Pathfinder and Latent Semantic Analysis. *AI and Society*, 11(1-2), 48-62.
- Chen, C. (1999). Visualizing semantic spaces and author co-citation networks in digital libraries. *Information Processing and Management*, 35(2), 401-420.
- Chen, C., and Carr, L. (1999). *Trailblazing the literature of hypertext: Author co-citation analysis (1989-1998)*. Paper presented at the 10th ACM Conference on Hypertext (Hypertext '99), Darmstadt, Germany, ACM Press, pp. 51-60.
- Chen, H., and Kim, J. (1994). GANNET: a machine learning approach to document retrieval. *Journal of Management Information Systems*, 11(3), 7-41.
- Chen, C., and Paul, R. J. (2001). Visualizing a knowledge domain's intellectual structure. *Computer*, 34(3), 65-71.
- Chen, H., Houston, A., Sewell, R., and Schatz, B. (1998). Internet browsing and searching: User evaluations of category map and concept space techniques. *Journal of the American Society for Information Science, Special Issue on "AI Techniques for Emerging Information Systems Applications"*, 49(7), 582-603.
- Chen, C., Thomas, L., Cole, J., and Chennawasin, C. (1999). Representing the semantics of virtual spaces. *IEEE Multimedia*, 6(2).
- Crossley, M., Davies, J., McGrath, A., and Rejman-Greene, M. (1999). The knowledge garden. *BT Technology Journal*, 17(1).
- Darken, R. P., and Sibert, J. L. (1996). Navigating large virtual spaces. *International Journal of Human-Computer Interaction*, 8(1), 49-71.
- Darken, R. P., Allard, T., and Achille, L. B. (1998). Spatial orientation and wayfinding in large-scale virtual spaces: An introduction. *Presence*, 7(2), 101-107.
- Deerwester, S., Dumais, S. T., Landauer, T. K., Furnas, G. W., and Harshman, R. A. (1990). Indexing by latent semantic analysis. *Journal of the American Society for Information Science*, 41(6), 391-407.
- Ding, Y. (2001). IR and AI: The role of ontology. *Proceedings of 4th International Conference of Asian Digital Libraries*, Bangalore, India, 10-12 December.
- Ding, Y., and Engels, R. (2001). IR and AI: Using co-occurrence theory to generate lightweight ontologies. *Proceedings of Workshop on Digital Libraries (Dlib2001), 12th International Conference on Database and Expert Systems Applications (DEXA2001)*, Munich, Germany, 3-7 September.
- Dodge, M. (2002). *An Atlas of Cyberspaces*. Available: <http://www.geog.ucl.ac.uk/casa/martin/atlas/>
- Enser, P. G. B. (1993). Query analysis in a visual information retrieval context. *Journal of Document and Text Management*, 1, 25-52.
- Fairchild, K., Poltrock, S., and Furnas, G. (1988). SemNet: Three-dimensional graphic representations of large knowledge bases. In R. Guidon (Ed.), *Cognitive Science and its Applications for Human-Computer Interaction*, Lawrence Erlbaum Associates, pp. 201-233.
- Feng, Y., and Börner, K. (2002). Using semantic treemaps to categorize and visualize bookmark files. In Erbacher, R. F., Chen, P. C., Grohn, M., Roberts, J. C., Wittenbrink, C. M. (Eds) *Proceedings of SPIE Conference on Visualization and Data Analysis*, San Jose, CA, 20-25 January, SPIE - The International Society for Optical Engineering. Vol. 4665. pp. 218-227.
- Fruchterman, T. M. J., and Reingold, E. M. (1991). Graph drawing by force-directed placement. *Software-Practice and Experience*, 21(11), 1129-1164.
- Furnas, G. W. (1997). Effective view navigation. *Proceedings of CHI '97*, Atlanta, GA, ACM Press; <http://www.acm.org/sigchi/chi97/proceedings/paper/gwf.htm>
- Golledge, R. G. (1995). Primitives of spatial knowledge. In T. L. Nyerges, D. M. Mark, R. Laurini and M. J. Egenhofer (Eds.), *Cognitive Aspects of Human-Computer Interaction for Geographic Information Systems*. Dordrecht: Kluwer, pp. 29-44.
- Hearst, M., and Karadi, C. (1997). Cat-a-Cone: An interactive interface for specifying searches and viewing retrieval results using a large category hierarchy. *Proceedings of 20th Annual International ACM/SIGIR Conference*, Philadelphia, PA, 27-31 July 1997, ACM Press, pp. 246-255; <http://www.sims.berkeley.edu/~hearst/papers/cac-sigir97/sigir97.html>
- Hefflin, J., and Hendler, J. (2000). Dynamic ontologies on the Web. *Proceedings of 17th National Conference on Artificial Intelligence (AAAI-2000)*, Menlo Park, CA, AAAI/MIT Press, pp. 443-449.
- Herot, C. F. (1980). Spatial management of data. *ACM Transactions of Database Systems*, 5(4), 493-513.

- Hightower, R. R., Ring, L. T., Helfman, J. I., Bederson, B. B., and Hollan, J. D. (1998). Graphical multiscale Web histories: A study of PadPrints. *Proceedings of 9th ACM Conference on Hypertext and Hypermedia (Hypertext '98)*, New York, NY, ACM Press, pp. 58–65;
<http://www.cs.umd.edu/hcil/pad++/papers/hypertext-98-padprints/hypertext-98-padprints.pdf>
- Landauer, T. K., Foltz, P. W., and Laham, D. (1998). Introduction to latent semantic analysis. *Discourse Processes*, 25, 259–284; <http://lsa.colorado.edu/>
- Lynch, K. J. (1960). *The Image of the City*. Cambridge, MA: MIT Press.
- Orendorf, J., and Kacmar, C. (1996). A spatial approach to organizing and locating digital libraries and their content. *Proceedings of 1st ACM International Conference on Digital Libraries*, ACM Press, pp. 83–89.
- Petry, F., Buckles, B., Prabhu, D., and Kraft, D. (1993). Fuzzy information retrieval using genetic algorithms and relevance feedback. *Proceedings of ASIS Annual Meeting*, pp. 122–125, Knowledge Industry Publications.
- Pirolli, P., and Card, S. K. (1995). Information foraging in information access environments. *Proceedings of CHI '95*, Denver, CO, ACM Press, pp. 51–58;
http://www.acm.org/sigchi/chi95/proceedings/papers/ppp_bdy.htm
- Raghavan, V. V., and Agarwal, B. (1987). Optimal determination of user-oriented clusters: An application for the reproductive plan. *Proceedings of 2nd International Conference on Genetic Algorithms and Their Applications*, Cambridge, MA, pp. 241–246, Lawrence Erlbaum Associates.
- Robertson, G., Czerwinski, M., Larson, K., Robbins, D., Thiel, D., and van Dantzych, M. (1998). Data mountain: Using spatial memory for document management. *Proceedings of 11th Annual ACM Symposium on User Interface Software and Technology (UIST '98)*, pp. 153–162, ACM Press;
<http://www.research.microsoft.com/users/marycz/home.htm>
- Robertson, G., Dantzych, M. v., Robbins, D., Czerwinski, M., Hinckley, K., Ridsen, K., Thiel, D., and Gorokhovskiy, V. (2000). The Task Gallery: A 3D windows manager. *Proceedings of Conference on Human Factors in Computing Systems CHI2000*, The Hague, 1–6 April 2000, ACM Press, pp. 494–501.
- Robertson, G. G., Mackinlay, J. D., and Card, S. K. (1991). Cone trees: Animated 3D visualizations of hierarchical information. *Proceedings of CHI '91*, New Orleans, LA, 28 April – 2 May 1991, pp. 189–194, ACM Press.
- Russell, D. M., Stefik, M. J., Pirolli, P., and Card, S. K. (1993). The cost structure of sensemaking. *Proceedings of INTERCHI '93 Conference on Human Factors in Computing Systems*, Amsterdam, pp. 269–276, ACM Press.
- Schatz, B. R. (1995). Building the Interspace: The Illinois Digital Library Project. *Communications of the ACM Computing Survey*, 38(4), 62–63.
- Schatz, B. R., and Chen, H. (1996). Building large-scale digital libraries. *Computer theme issue on the US Digital Library Initiative* 29(5), (<http://computer.org/computer/dli/>).
- Schvaneveldt, R. W. (1990). *Pathfinder Associative Networks: Studies in Knowledge Organization*. Norwood, NJ: Ablex.
- Shneiderman, B. (1996). The eyes have it: A task by data type taxonomy for information visualizations. *Proceedings of Symposium on Visual Languages*, Boulder, CO, IEEE, pp. 336–343.
- Shum, S. B., Motta, E., and Domingue, J. (2000). ScholOnto: An ontology-based digital library server for research documents and discourse. *International Journal on Digital Libraries*, 3(3), 237–248.
- Tan, D. S., Stefanucci, J. K., Proffitt, D. R., and Pausch, R. (2001). The Infocockpit: Providing location and place to aid human memory. *Proceedings of Workshop on Perceptive User Interfaces*, Orlando, FL; <http://www.cs.ucsb.edu/PUI/PUIWorkshop/PUI-2001/PUI2001Schedule.html>
- Thorndyke, P., and Hayes-Roth, B. (1982). Differences in spatial knowledge acquired from maps and navigation. *Cognitive Psychology*, 14, 560–589.
- Vincow, M. A., and Wickens, C. D. (1998). Frame of reference and navigation through document visualizations: Flying through information space. *Proceedings of 42nd Annual Meeting of the Human Factors and Ergonomics Society*, Santa Monica, CA, Human Factors and Ergonomics Society, pp. 511–515.
- Vinson, N. G. (1999). Design guidelines for landmarks to support navigation in virtual environments. *Proceedings of CHI 1999*, Pittsburg, PA, ACM Press.
- Wickens, C. D., Liang, C.-C., Prett, T. T., and Olmos, O. (1994). Egocentric and exocentric displays for terminal area navigation. *Proceedings of Human Factors and Ergonomics Society 38th Annual Meeting*, Santa Monica, CA, Human Factors and Ergonomics Society, pp. 16–20.