Web-based Discovery and Dissemination of Multidimensional Geographic Information

UNIVERSITY OF
NEWCASTLE UPON TYNE

Gobe Ernesto Hobona BSc. (Hons.)

Thesis Submitted for the Degree of
Doctor of Philosophy

School of Civil Engineering and Geosciences
University of Newcastle upon Tyne

September 2006

Abstract

A spatial data clearinghouse is an electronic facility for searching, viewing, transferring, ordering, advertising, and disseminating spatial data from numerous sources via the Internet. Governments and other institutions have been implementing spatial data clearinghouses to minimise data duplication and thus reduce the cost of spatial data acquisition. Underlying these clearinghouses are geoportals and databases of geospatial metadata. A geoportal is an access point of a spatial data clearinghouse and metadata is data that describes data. The success of a clearinghouse's spatial data discovery system is dependent on its ability to communicate the contents of geospatial metadata by providing both visual and analytical assistance to a user. The model currently adopted by the geographic information community was inherited from generic information systems and thus to an extent ignores spatial characteristics of geographic data. Consequently, research in Geographic Information Retrieval (GIR) has focussed on spatial aspects of web-based data discovery and acquisition.

This thesis considers how the process of GIR from geoportals can be enhanced through multidimensional visualisation served by web-based geographic data sources. An approach is proposed for the presentation of search results in ontology-assisted GIR. Also proposed is an approach for the visualisation of multidimensional geographic data from web-based data sources. These approaches are implemented in two prototypes, the Geospatial Database Online Visualisation Environment (GeoDOVE) and the Spatio-Temporal Ontological Relevance Model (STORM). A discussion of their design, implementation and evaluation is presented. The results suggest that ontology-assisted visualisation can improve a user's ability to identify the most relevant multidimensional geographic datasets from a set of search results. Additional results suggest that it is possible to offer the proposed visualisation approaches on existing geoportal frameworks. The implication of the results is that multidimensional visualisation should be considered by the wider geographic information community as an alternative to historic approaches for presenting search results on geoportals, such as the textual ranked list and two-dimensional maps.

Table of Contents

List of Figures

List of Tables

List of Abbreviations and Acronyms

API	Application Programmers' Interface
BLOB	Binary Large Object
CAT	Catalogue Services Specification
CCS	Geocentric Cartesian Coordinate System
CLEF	Cross Language Evaluation Forum
CNIDR	Clearinghouse for Networked Information Discovery and Retrieval
CSDGM	Content Standard for Digital Geospatial Metadata
CSG	Constructive Solid geometry
CSW	Catalogue Service for the Web
DAML+OIL	DARPA Agent Markup Language and Ontology Inference Layer
DBMS	Database management System
DEM	Digital Elevation Model
DTM	Digital Terrain Model
EAI	External Authoring Interface
EPSG	European Petroleum Survey Group
FGDC	Federal Geospatial Data Committee
GCS	Geographic Coordinate System
GE	Google Earth
GeoDOVE	Geospatial Database Online Visualisation Environment
GeoVRML	Geographic Virtual Reality Modelling Language
GIR	Geographic Information Retrieval
GIS	Geographic Information System
GML	Geography Markup Language
GOS	Geospatial One Stop
GUI	Graphical User Interface
HTTP	Hypertext Transfer Protocol
IR	Information Retrieval
ISO	International Organisation for Standardisation
JDBC	Java Database Connectivity
JRE	Java Runtime Environment
JSP	Java Server Pages
JVM	Java Virtual Machine

KML	Keyhole Markup Language
LR	Logistic Regression
MBR	Minimum Bounding Polygon
MBR	Minimum Bounding Rectangle
NEEDS	North East Environmental Data Server
ODBC	Open Database Connectivity
OGC	Open Geospatial Consortium
OWL	Web Ontology Language
PCS	Projected Coordinate System
RPN	Reverse Polish Notation
SDI	Spatial Data Infrastructure
SFS	Simple Features for SQL
SOA	Service-Oriented Architecture
SOAP	Simple Object Access Protocol
SPIRIT	Spatially-Aware Information Retrieval on the Internet
SQL	Structured Query Language
STORM	Spatio-Temporal Ontological Relevance Model
TREC	Text Retrieval Conference
UDDI	Universal Description Discovery and Integration
UML	Unified Modelling Language
URI	Uniform resource identifier
VRML	Virtual Reality Modelling Language
W3C	World Wide Web Consortium
WCS	Web Coverage Service
WFS	Web Feature Service
WMS	Web Map Service
X3D	Extensible 3D
XML	Extensible Markup Language

Acknowledgements

During the process of undertaking this research, several people have been supportive and inspirational to me. I would like to take this opportunity to thank all of them as their help and support has been invaluable during my studies.

I wish to thank my supervisors David Fairbairn and Phil James for the years of academic support and guidance. Special thanks to Professors David Parker and Phil Moore for believing in my research potential and enabling me to undertake my studies at the School of Civil Engineering and Geosciences. Thanks to Nosa Erharuyi and Simon Abele for the mentoring; the community of postgraduates and staff at *geomatics@newcastle* for creating a friendly and welcoming atmosphere. Thanks to everyone who participated in the experiments in this study; your patience and assistance were highly instrumental to the completion of this thesis.

Thanks to my fellow yudansha at Kensho karate club for providing recreation away from the geomatics lab and making Newcastle Gateshead a fun place to be.

My family has been highly inspirational during the period of this research; Thank you Ndingililo, Phildah, Beulah, Lerato, David, Bertha, Lucas, Pungie, my nephews and nieces. My parents for the love and support throughout my studies, nda boka, bazwadzi bangu. Thanks to Pinah, Tuffy and Gemma for all the encouragement as I was writing this thesis.

This thesis is produced from research jointly funded by the University of Newcastle upon Tyne and Universities UK through the Overseas Research Students Awards Scheme ORS Ref: 2004029022.

Publications from this Research

The following publications have been produced from the research presented in this thesis:

Hobona, G., James, P. and Fairbairn, D. (2006a) 'Multidimensional Visualisation of Degrees of Relevance of Geographic Data', International Journal of Geographic Information Science, 20, (5), pp. 469 - 490.

Hobona, G., James, P. and Fairbairn, D. (2006b) 'Web-Based Visualisation Of Three-Dimensional Geospatial Data Using Java3D', IEEE Computer Graphics and Applications, 26, (4), pp. 28- 33.

Chapter 1 Introduction

1.1 Context

1.1.1 Geographic Data Infrastructure

Geographic information through maps has been collected for several centuries; however, it is clear that there is more geographic information being collected since the turn of the millennium than at any other time in history. The significant increase in the amount of geographic information is due to the technological advances that have been made in the past few decades in space and airborne geographic data collection. Longley et al.(2001) define geographic data as linking a place, a point in time and some descriptive property. Technologies for geographic data collection have advanced against a backdrop of other improvements in digital data sharing and dissemination.

In response to the growing cost of geographic data collection, initiatives were started by world governments to create Spatial Data Infrastructures (SDI) — a collection of technologies, policies, and frameworks for facilitating the availability and accessibility geographic data. SDIs are therefore meant to provide mechanisms for the discovery, retrieval, evaluation and application of geographic data. In addition to reduced data collection costs, an SDI is essential for providing affordable, timely and effective public services. The first recognisable National SDI (NSDI) was initiated by an Executive Order 12906 by former United States President Bill Clinton in 1994(FDGC, 2005). This was as a result of the realisation that the US government was spending approximately four billion dollars on geographic data annually. In the United Kingdom, the cost of geographic data acquisition is estimated to be about 400 million euros. Consequently, the Office of The Deputy Prime Minister (ODPM) is an active supporter of the SDI initiative (AGI, 2006).

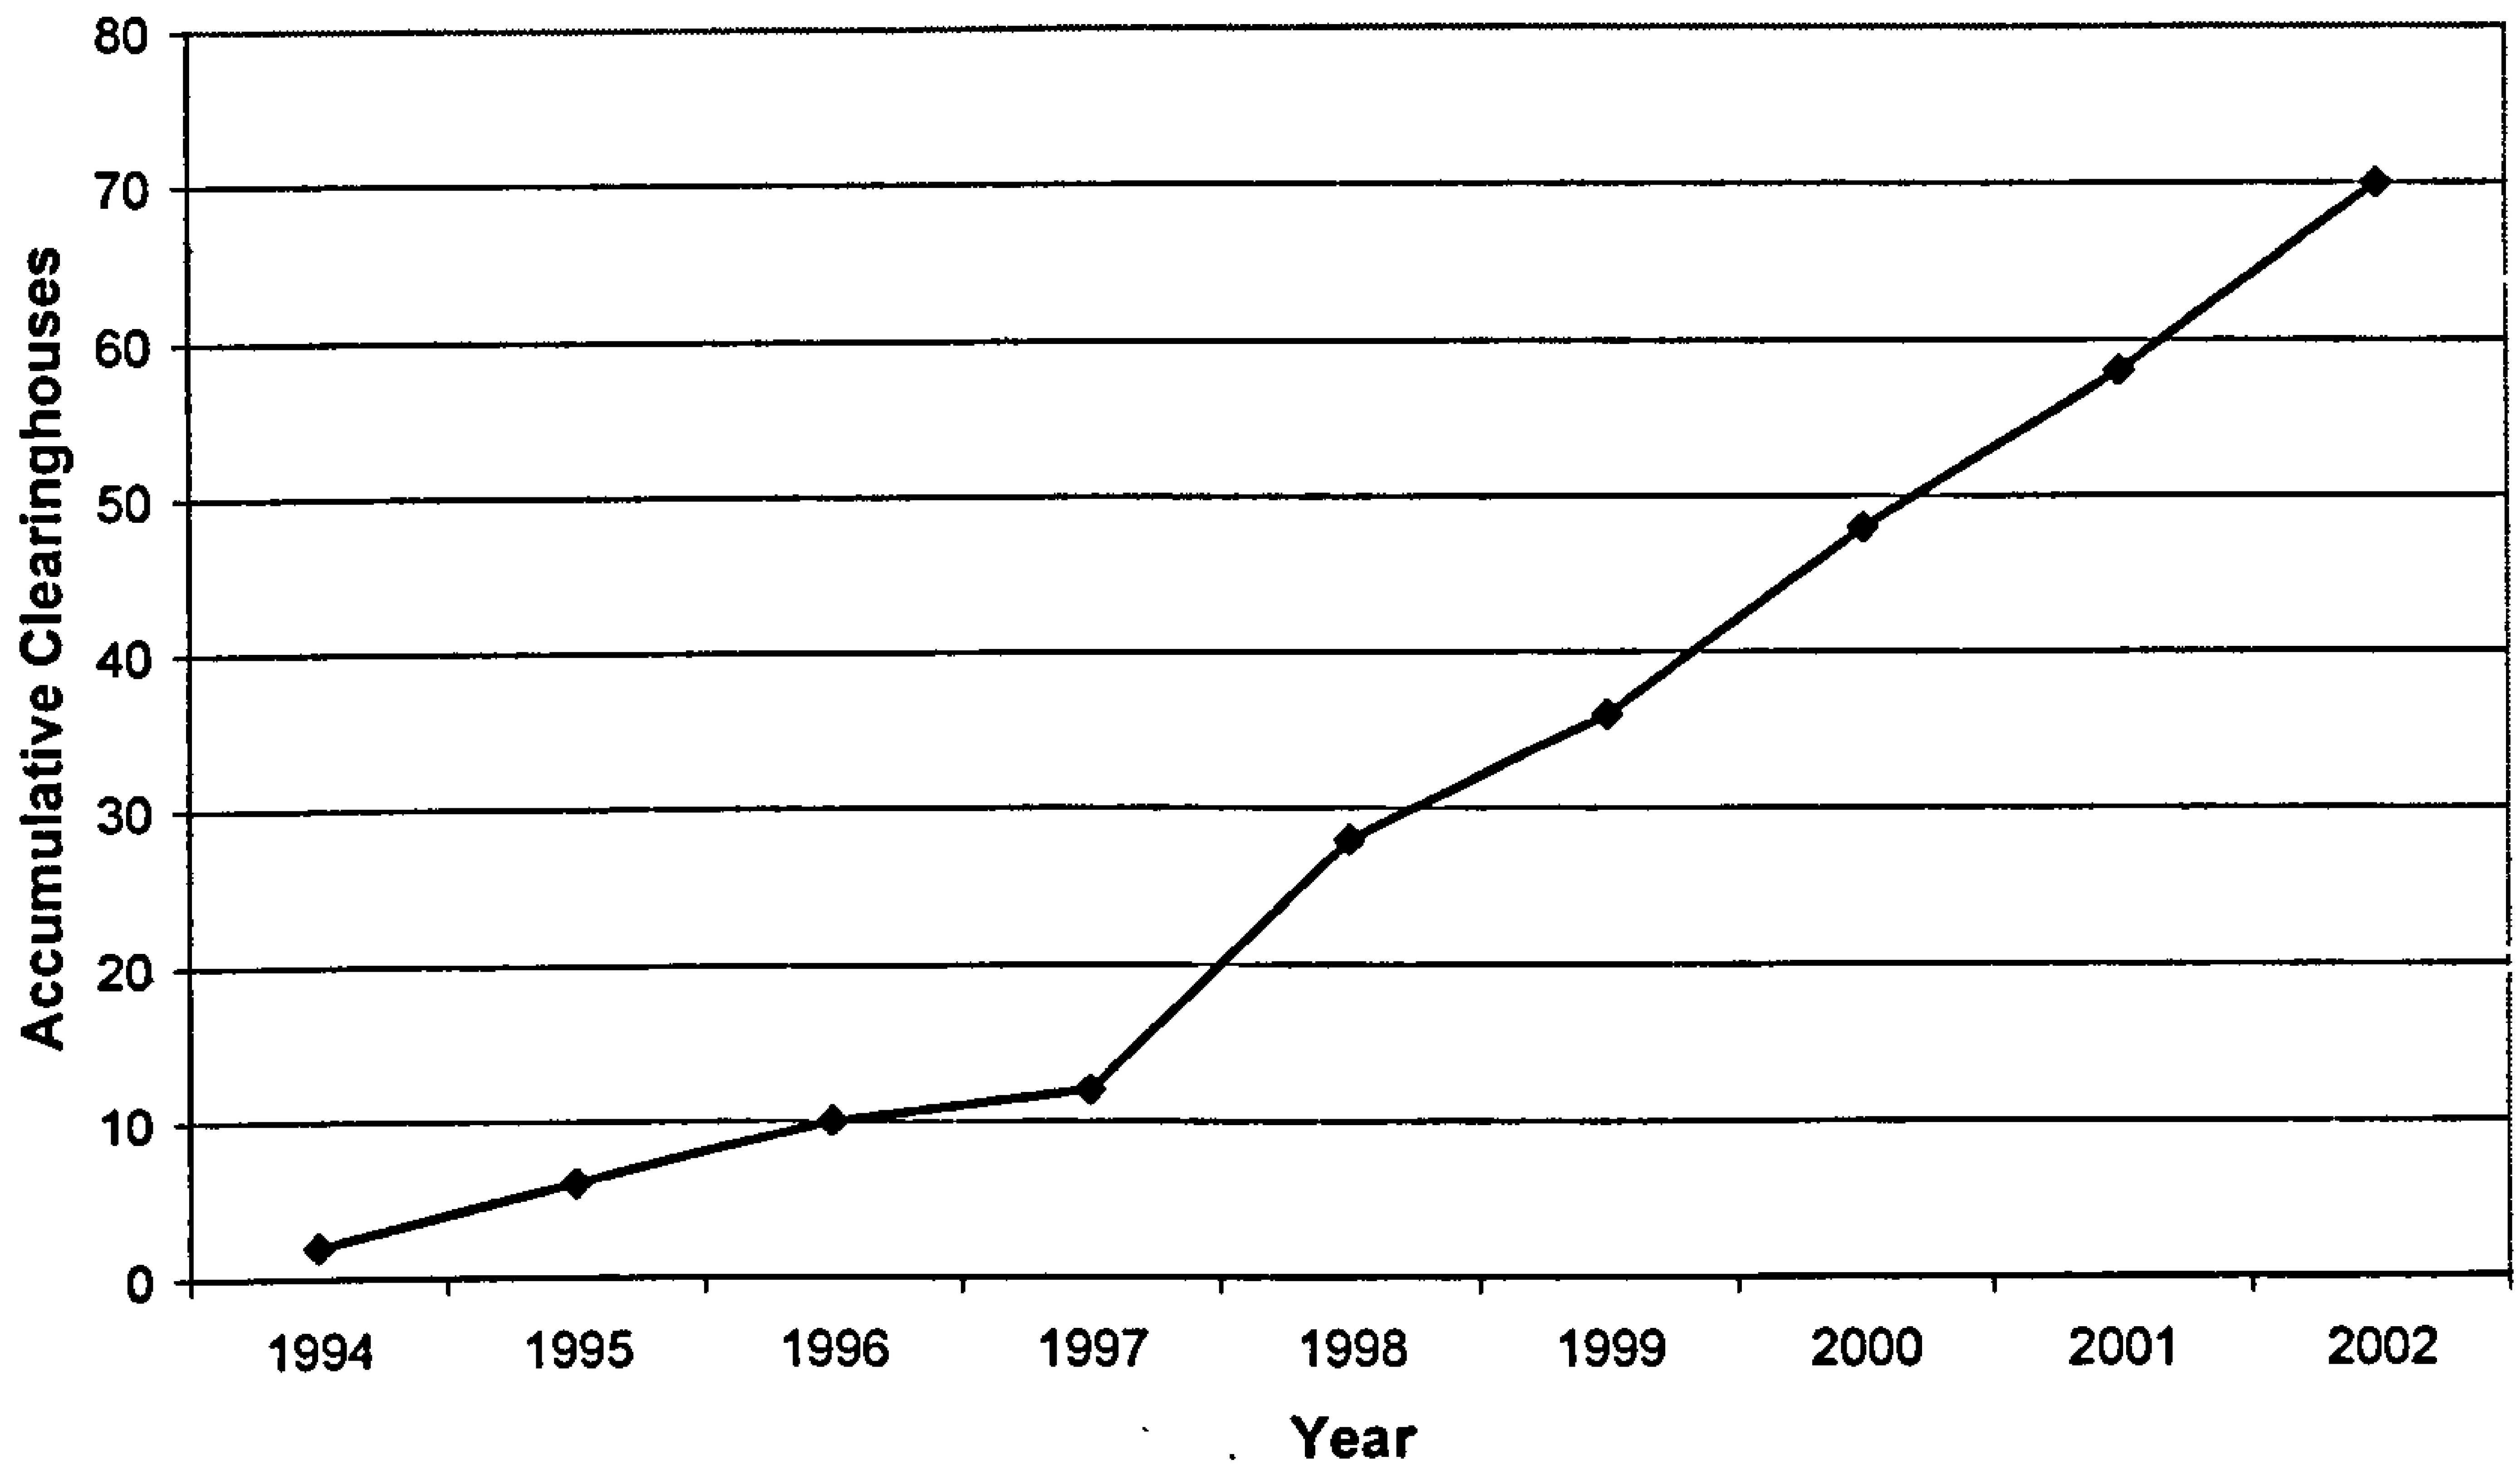

Figure 1.1 Total number of spatial data clearinghouses globally from 1994 to 2002 (adapted from Crompvoets et al., 2004)

"A spatial data clearinghouse is defined as an electronic facility for searching, viewing, transferring, ordering, advertising and/or disseminating spatial data from numerous sources via the Internet and, as appropriate, providing complementary services" (Crompvoets et al., 2004: pp. 665). A clearinghouse adopts a distributed information systems architecture, where the data is not held centrally but is retrieved on-demand. Data producers and software vendors make known what products they have available and how the data can be accessed. The data is described in terms of its quality, spatial and thematic specifications allowing potential customers to evaluate its potential for addressing a particular information need. This form of descriptive information about data is called *metadata* and is discussed later in this chapter. A recent innovation in the development of clearinghouses has been the incorporation of listings of 'wanted' datasets. These listings allow geographic information users to advertise a particular need for information; allowing vendors to respond with possible solutions for their needs.

Since the inception of the first SDI, the number of national spatial data clearinghouses has been steadily increasing with approximately 67 countries having implemented a web-accessible clearinghouse and 13 countries in the process of implementation (Crompvoets et al., 2004). Figure 1.1 shows the accumulation of national spatial data clearinghouse from a global total of one in the year 1994 to 67 by the year 2002, however it excludes the intra-organisational clearinghouses that have been implemented and managed within each country. Intra-organisational clearinghouses are created by local authorities, companies and other organisations within a country. They are therefore not national clearinghouses. For example, the United States has over 180 clearinghouse nodes referenced by the national spatial data clearinghouse, Canada has over 80 clearinghouse nodes within its territory and they are also referenced by the national clearinghouse, and the United Kingdom has approximately 10 published clearinghouse nodes (FDGC, 2005).

Currently the global network of national clearinghouses makes available over 250 000 semi-structured descriptions of geographic information resources (Crompvoets et al., 2004). Unstructured content is also available from websites with geographic references, for example, the *Newcastle* City Council website or the *Newcastle* United Football Club. Other contributions to the global collection of web-based geographic information resources have been the publication of mapping application programming interfaces (API) by popular media companies such as Microsoft, Yahoo! and Google. The geographic information community has adopted their APIs for offering maps on different kinds of issues, for example, transportation, travel, recreation and leisure. It can thus be expected that these factors will contribute to the increase in the number of geographic information resources available on the Internet.

1.1.2 Geoportals

A user accesses a spatial data clearinghouse through a geoportal — a website that acts as a gateway to geographic content on the Internet (Tait, 2005). Therefore, a geoportal is effectively the interface of a spatial data clearinghouse, as illustrated in Figure 1.2. The illustration shows that a GIS service provider *publishes* his products on a geoportal. A GIS user then *searches* for the services on a geoportal. If any relevant services are available, these are *discovered* by the GIS user, after which he then *consumes* the services by using them in a GIS application. Whereas historically, geoportals focussed on publishing geographic data and associated metadata, recent advances in web-based technology have allowed geoportals to offer web-based geospatial processing as well. Although discovered through a geoportal, such web-based processing is made accessible through both a web browser and a desktop GIS.

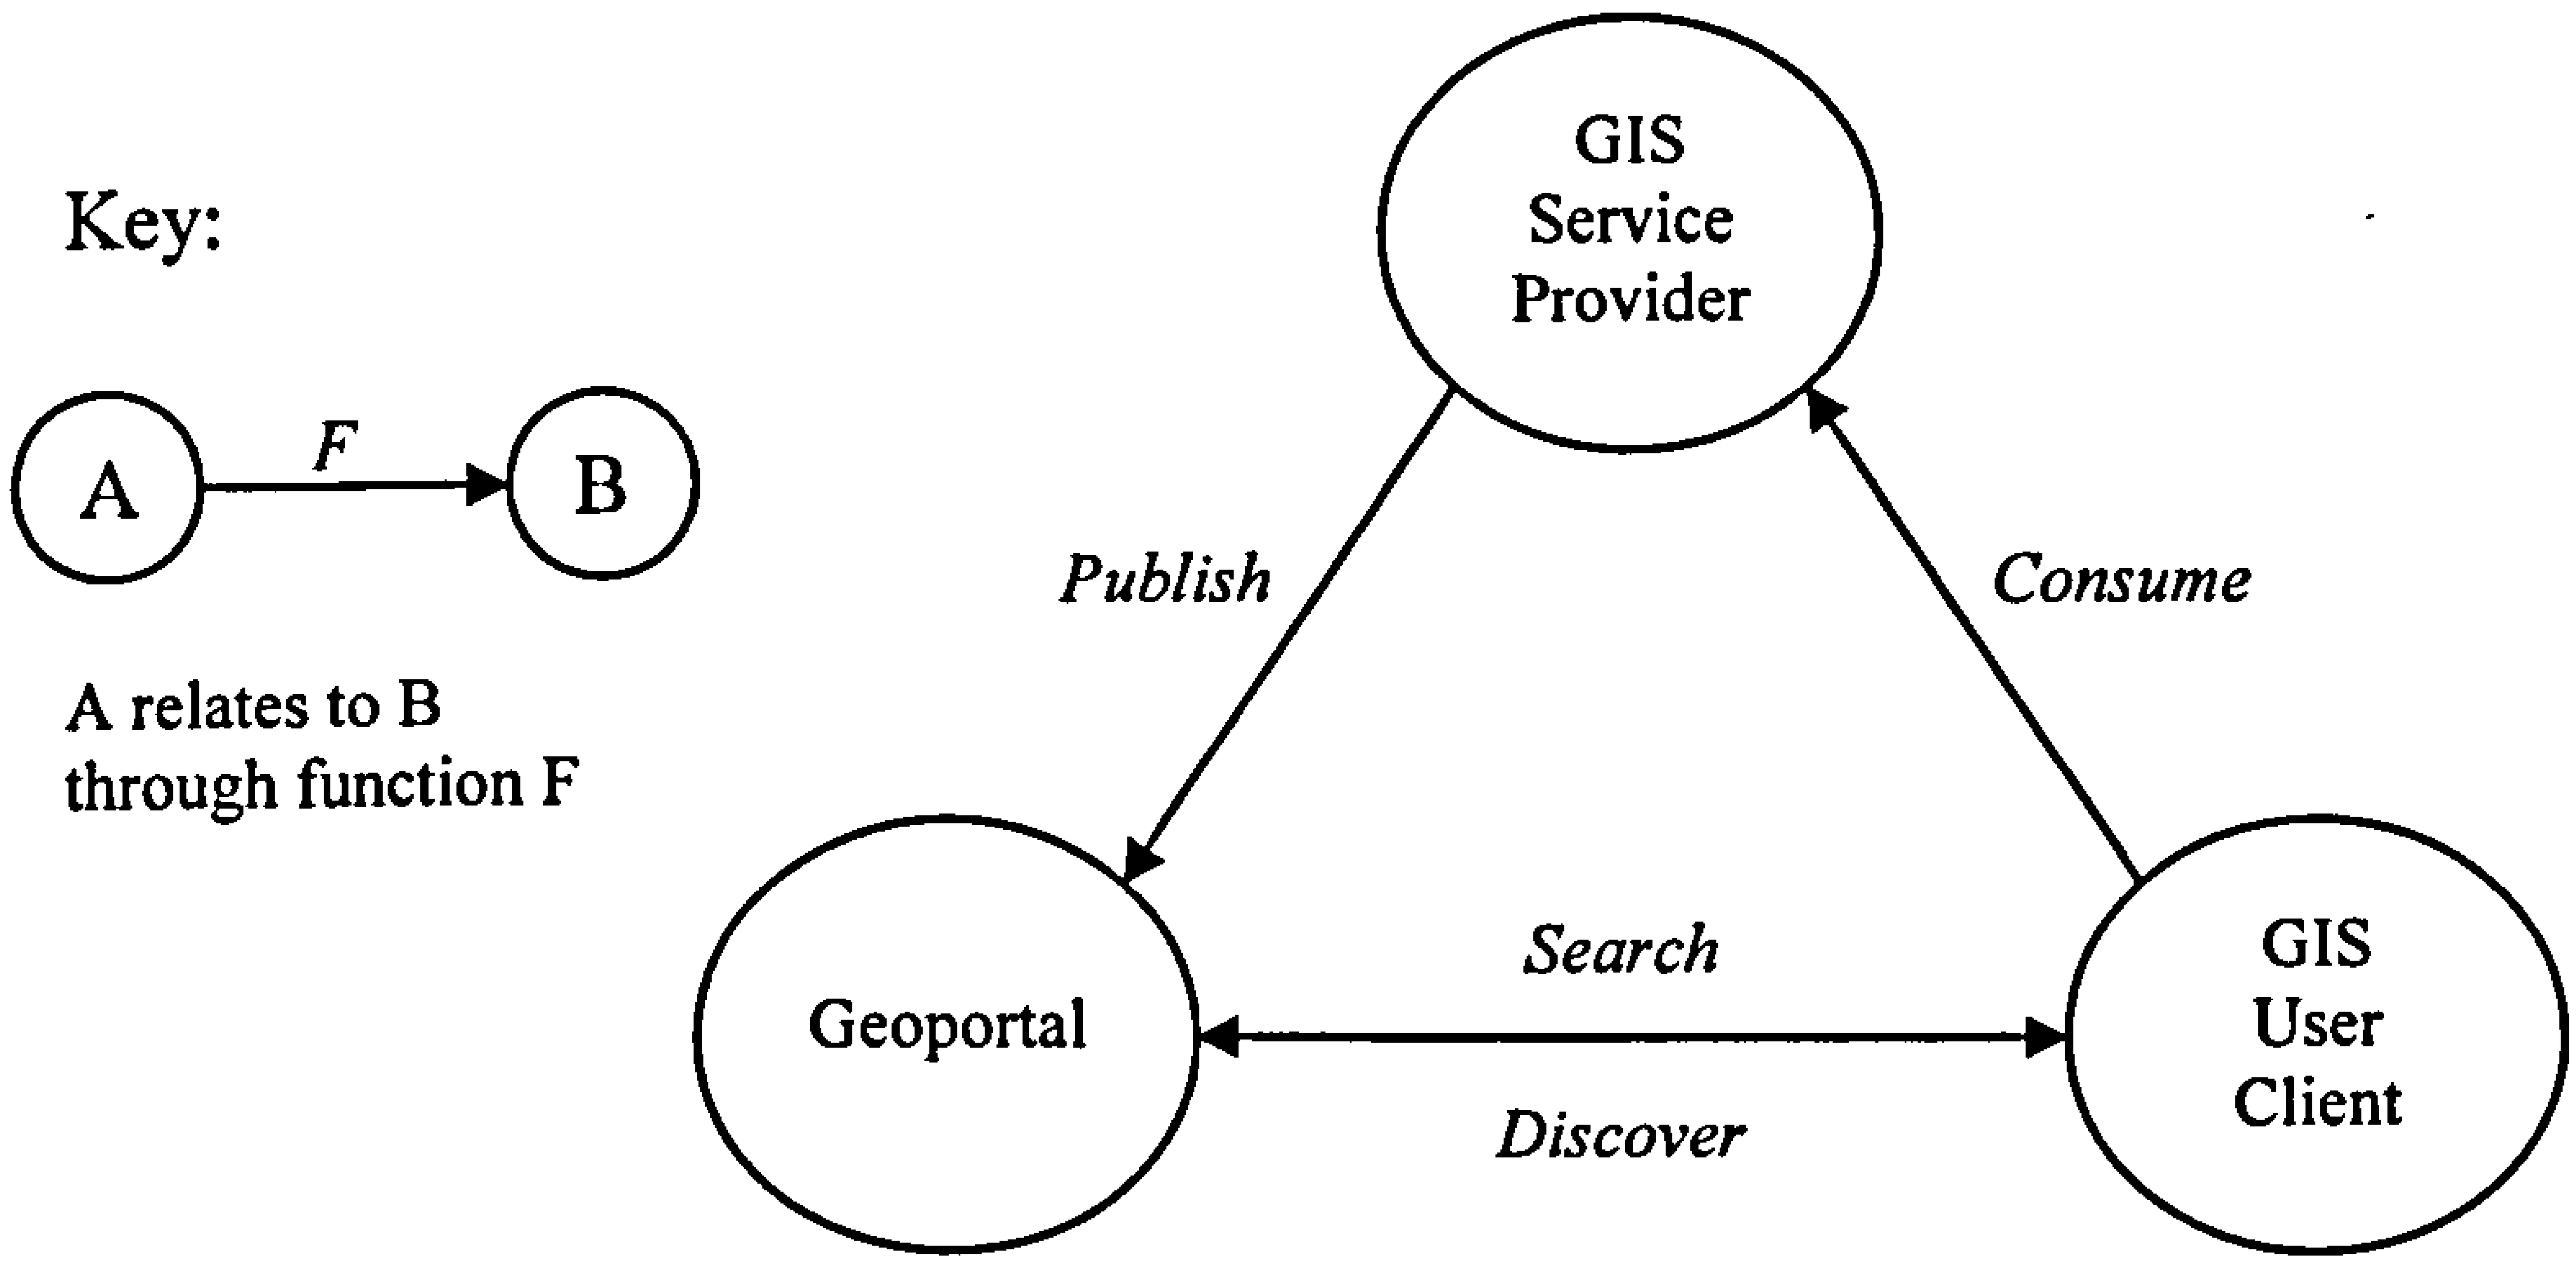

Figure 1.2 The role of a geoportal in an SDI (adapted from Maguire and Longley, 2005)

Maguire and Longley (2005) classify geoportals into two main groups, namely catalogue and application geoportals. Catalogue portals create and maintain indexes of descriptive information about available geographic information resources. They offer a web-based interface that allows users to query the indexes remotely. They are therefore primarily intended for data discovery and are generally targeted towards wide audiences. Application geoportals combine geographic information services to provide advanced GIS functions for addressing specific tasks. Tang and Selwood (2005) suggest an additional type of geoportal, called an enterprise geoportal, which incorporates web-mapping and location services into an existing business web site. This thesis is mainly concerned with catalogue geoportals and the methods through which they enable data to be discovered.

1.1.2.1 Architecture

A geoportal is typically made up of several components. To ensure interoperability between the different components and associated software, international standards have been developed to provide a standardised interface for communication between components. The functionality offered by each component is referred to as a service and an architecture that adopts this form of structure is illustrated in Figure 1.3. The architecture has the following benefits:

- It allows for the integrated use of components from different vendors, communicating through established protocols.
- Components operate independently of one another and consequently, errors are isolated within each component.
- As data is disseminated through a web service, the publisher or originator can impose their own access constraints on the data, for example, maximum permissible resolution on previewed imagery.
- Data is delivered in standardised file formats, making it interoperable with several GIS.

Considering the three-tier architecture illustrated in Figure 1.3, at the lowest level of a geoportal model is a data management tier. This tier includes geographic and non-geographic database management systems (DBMS). By definition, geographic databases include data referenced to locations on the earth; for example, census, flood, weather and geological datasets. The data is stored as vector, raster or tabular content. Querying of the DBMS and meta-database systems is carried out using the Structured Query Language (SQL) or a similar language such as the Reverse Polish Notation (RPN). In addition to these databases, the data management tier also includes repositories of geographic metadata, described in more detail later in this chapter.

The middle tier of the architecture includes a series of web services. These offer functionality for accessing, querying and retrieving information from the data repositories and presenting it to the user via the top-level portal services. The middle-tier services of a geoportal include the data, portrayal and catalogue services. Technologies used for communication between components in this tier and others, include the Geography Markup Language (GML), Simple Object Access Protocol (SOAP), Web Service Description Language (WSDL), Web Map Service (WMS) and Web Feature Services (WFS) specifications. These technologies ensure interoperability between components by providing a standardised interface between web services and clients. A more detailed discussion of web services is offered in Chapter 4.

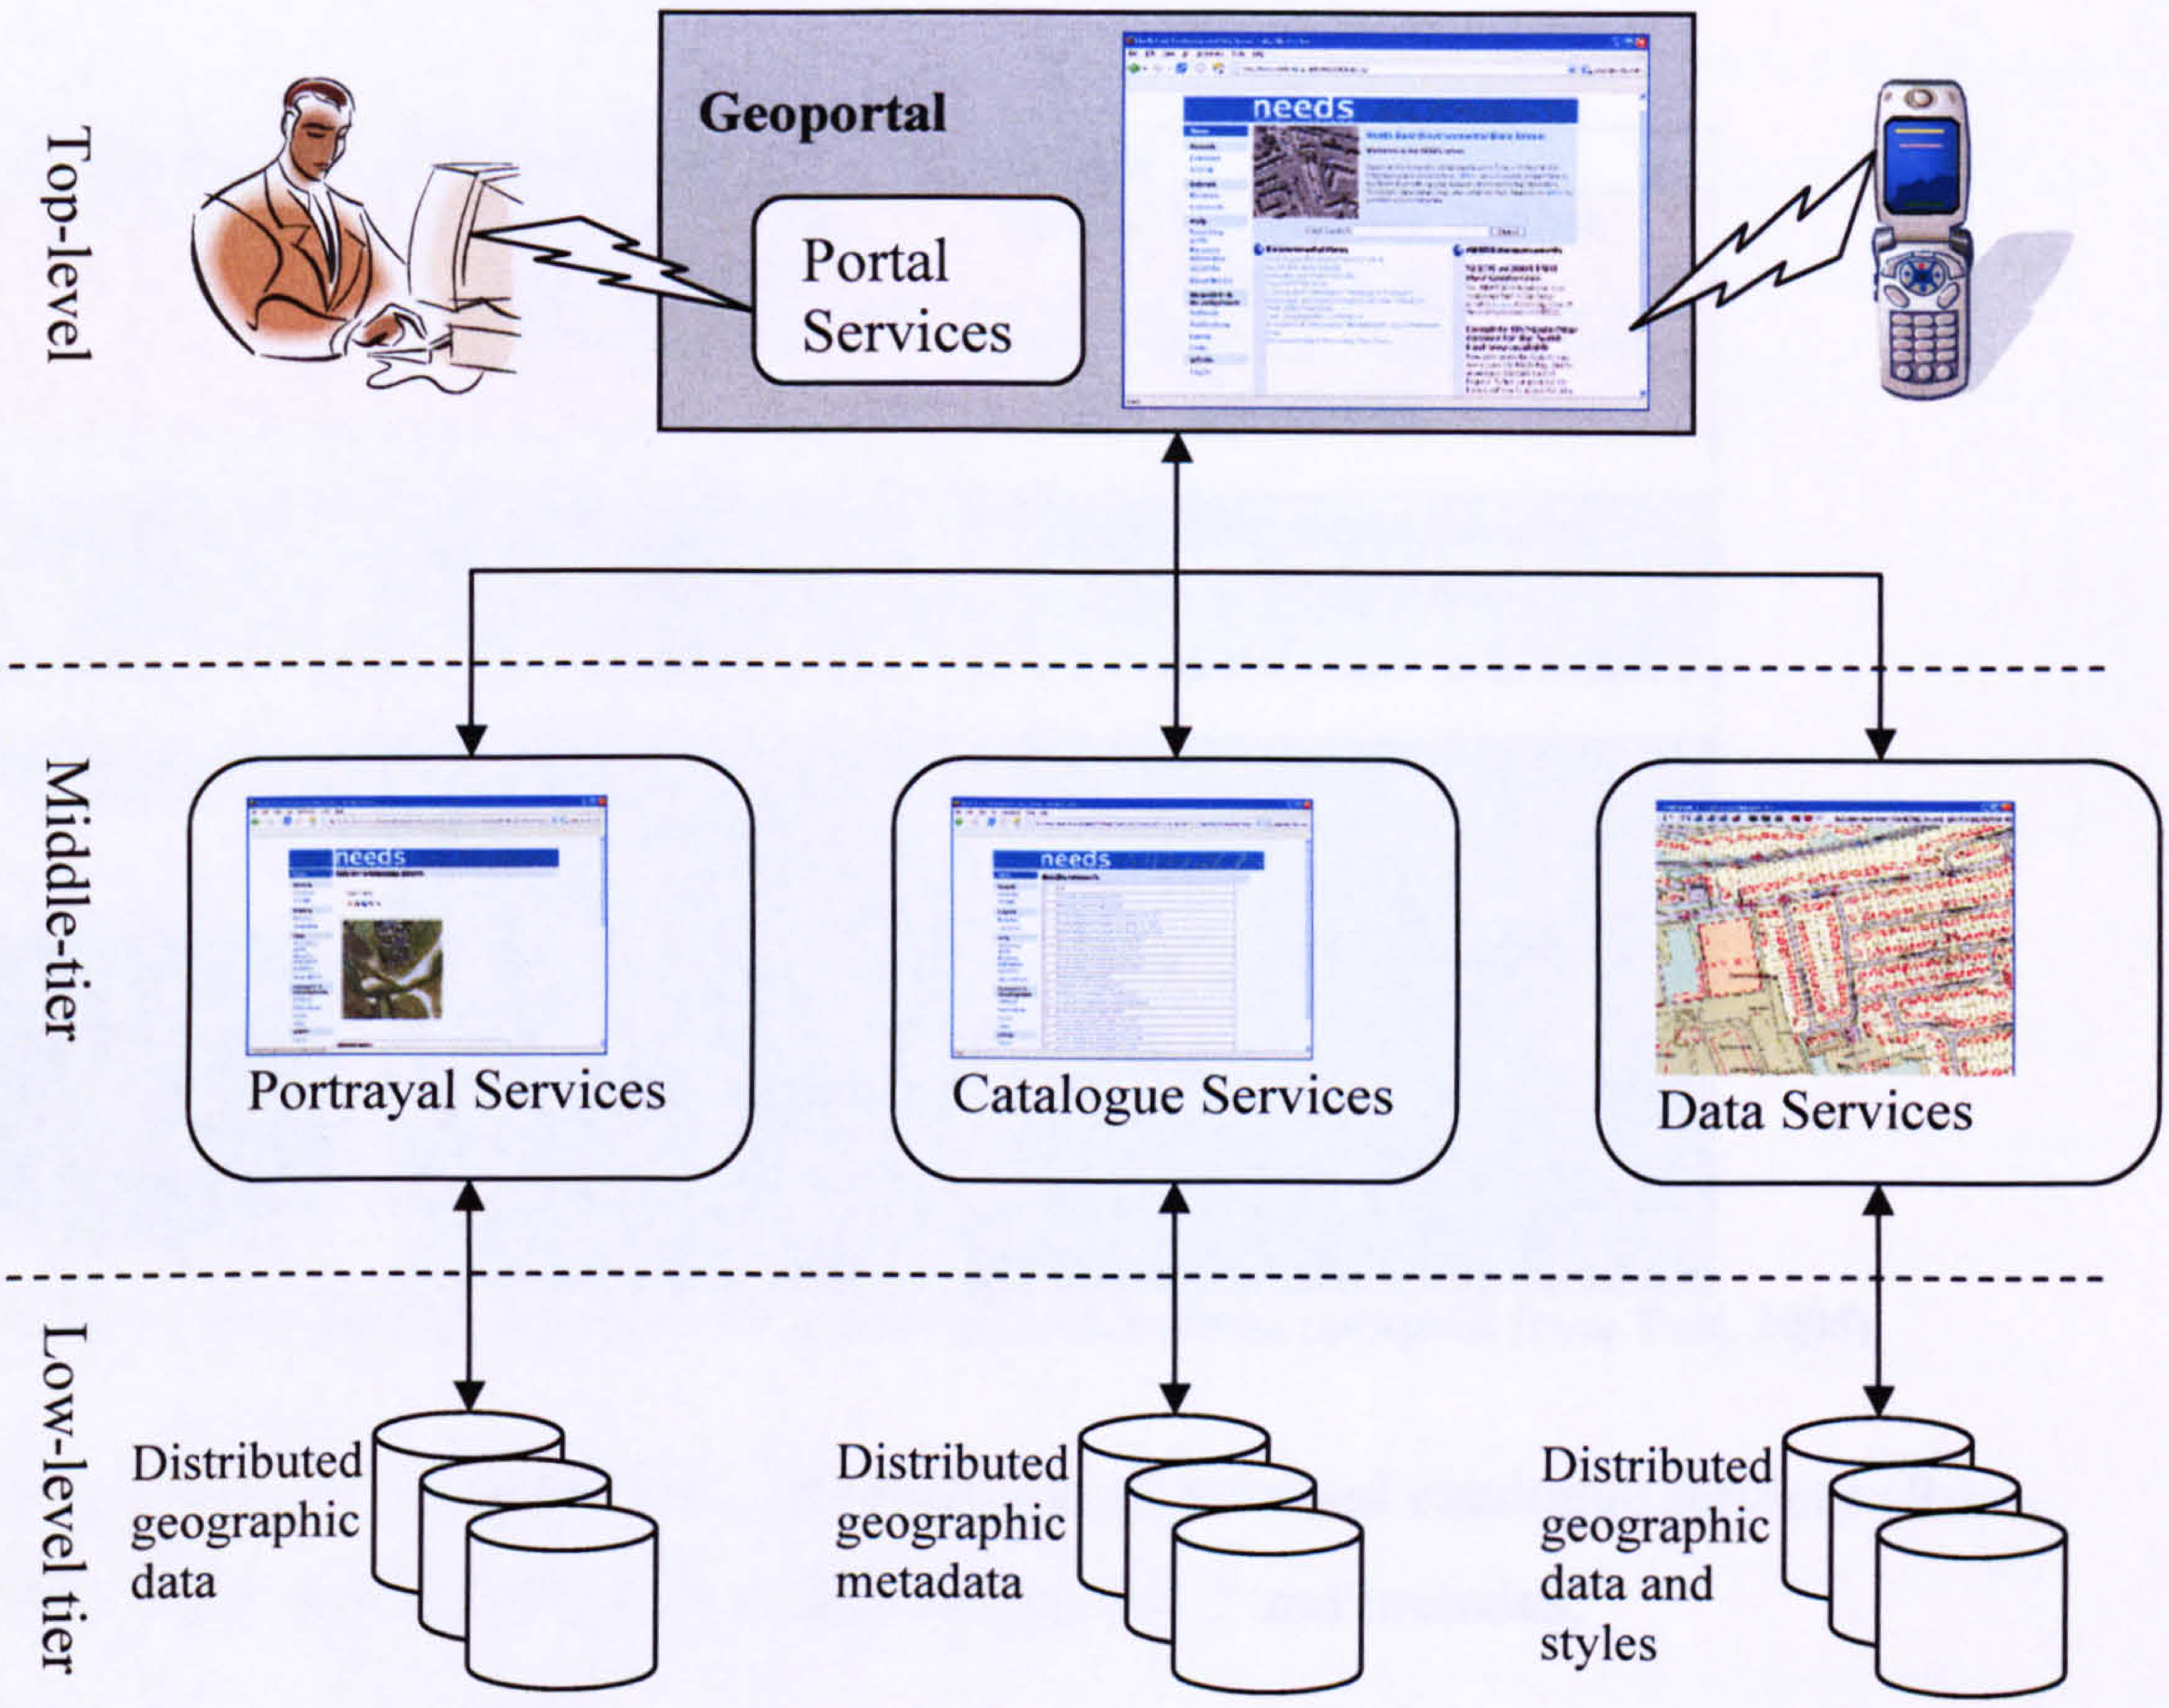

Figure 1.3 Service-based framework for geoportals (adapted from Rose, 2004)

At the highest-level within the geoportal architecture are portal services. These are primarily the user-interface components of the geoportal. There is an increasing demand for these services to be graphical to ensure adequate portrayal of maps and other geographic information. Initiatives such as the European Union INSPIRE framework have included 'view services' as requirements for conformant geoportals (Commission of the European Communities, 2004). Within the INSPIRE framework, 'view services' refer to web mapping and other visualisation components. Other national geoportals have also implemented clients for viewing data and other content, for example the United States Geospatial One Stop (GOS) and the India SDI (Sivakuma et al., 2004).

Components	Environments	Functions
Web Portal	HTML, HTTP, XSL, XML, JSP, ASP	Search, Map Viewer, Publish, Administrate
	Java Beans, .NET	Query, Gazetteer, Mapping, Edit, Geocoding
Geographic Web Services	XML, SOAP, WSDL, WMS, WFS, GML	Query, Map render/feature, Transaction, Geocode
DBMS	SQL	Raster, Vector, Tabular
Geographic & Tabular Data		

Figure 1.4 A techno-centric view of a geoportal architecture (adapted from Tait, 2005)

Services may be classified into portrayal, portal, data and catalogue services (Rose, 2004). The classification is illustrated in Figure 1.3 and includes,

Portal Services: Effectively the user-interfaces for administration, searching, mapping, and publishing functionality on a geoportal. Components providing these services include map viewers, discovery clients, access and management clients.

Portrayal Services: Provide operations for processing and rendering geographic content. These include maps (rendered geographic data) and styling services (providing cartographic information such as symbols and colours).

Data Services: Primarily used for the delivery of geographic content and services. These include gazetteer, feature and coverage services.

Catalogue Services: Used to advertise and locate geographic information resources. These include data discovery, service registry, service discovery, metadata entry and update components.

This classification identifies almost all components, including viewers, as services. The classification is based on the Open Geospatial Consortium (OGC) *Service*

Architecture Abstract Specification. The OGC is an international body of academic, governmental and commercial organisations that aims to improve interoperability between different GIS systems. An alternative architecture, illustrated in Figure 1.4, places web services as the interface layer between the data management components and the portal components (Tait, 2005). Although the two models differ in the names given to the component groups, it is evident that the same functionalities are identified as being key to geoportal composition. The latter figure is basically a more techno-centric illustration of the architecture.

1.1.2.2 Geospatial-One-Stop : A Case Study

One of the most recognisable geoportals is Geospatial-One-Stop (GOS), the national geoportal of the United States. GOS hosts over 100 000 geographic information resources, 1400 of which are live image services, published by over 500 information providers. Tang and Selwood (2005) report that GOS was handling over 600 visitors per day in 2004. Each visitor could define an information need by selecting an area of interest or entering a keyword, then sending a request for any relevant resources through the geoportal. The response from a search is a list of records describing available resources, in most cases accompanied by a link to a live map that can be viewed using a simple web map client. The web map client on GOS implements standard geospatial specifications and can therefore support several geographic data servers. It is also possible to save the composition of the live map, such that it can be shared between different users of GOS.

Currently in its second major version, GOS offers data, news and information categorised into several topics (channels) including for example administrative boundaries, agriculture, atmosphere, biology, business, cadastral, demographic, elevation, environment, facilities, geology, health, inland water, locations, oceans, transportation, utilities imagery and basemaps. A team of experts, known as channel stewards, manages and groups resources into each of these topics. This increases the likelihood of discovering relevant content. The channel stewards are also tasked with seeking contributions from people in their fields of interest. In response to special events, the stewards are tasked with creating temporary channels targeted towards those events. For example, a special channel called 'GIS for the Gulf' was set up to assist emergency management teams responding to hurricanes Katrina and Rita (Tang and Selwood, 2005). The special channels comprised information on topography, cadastre, emergency services, weather conditions and flood levels in the rivers and tidal stations.

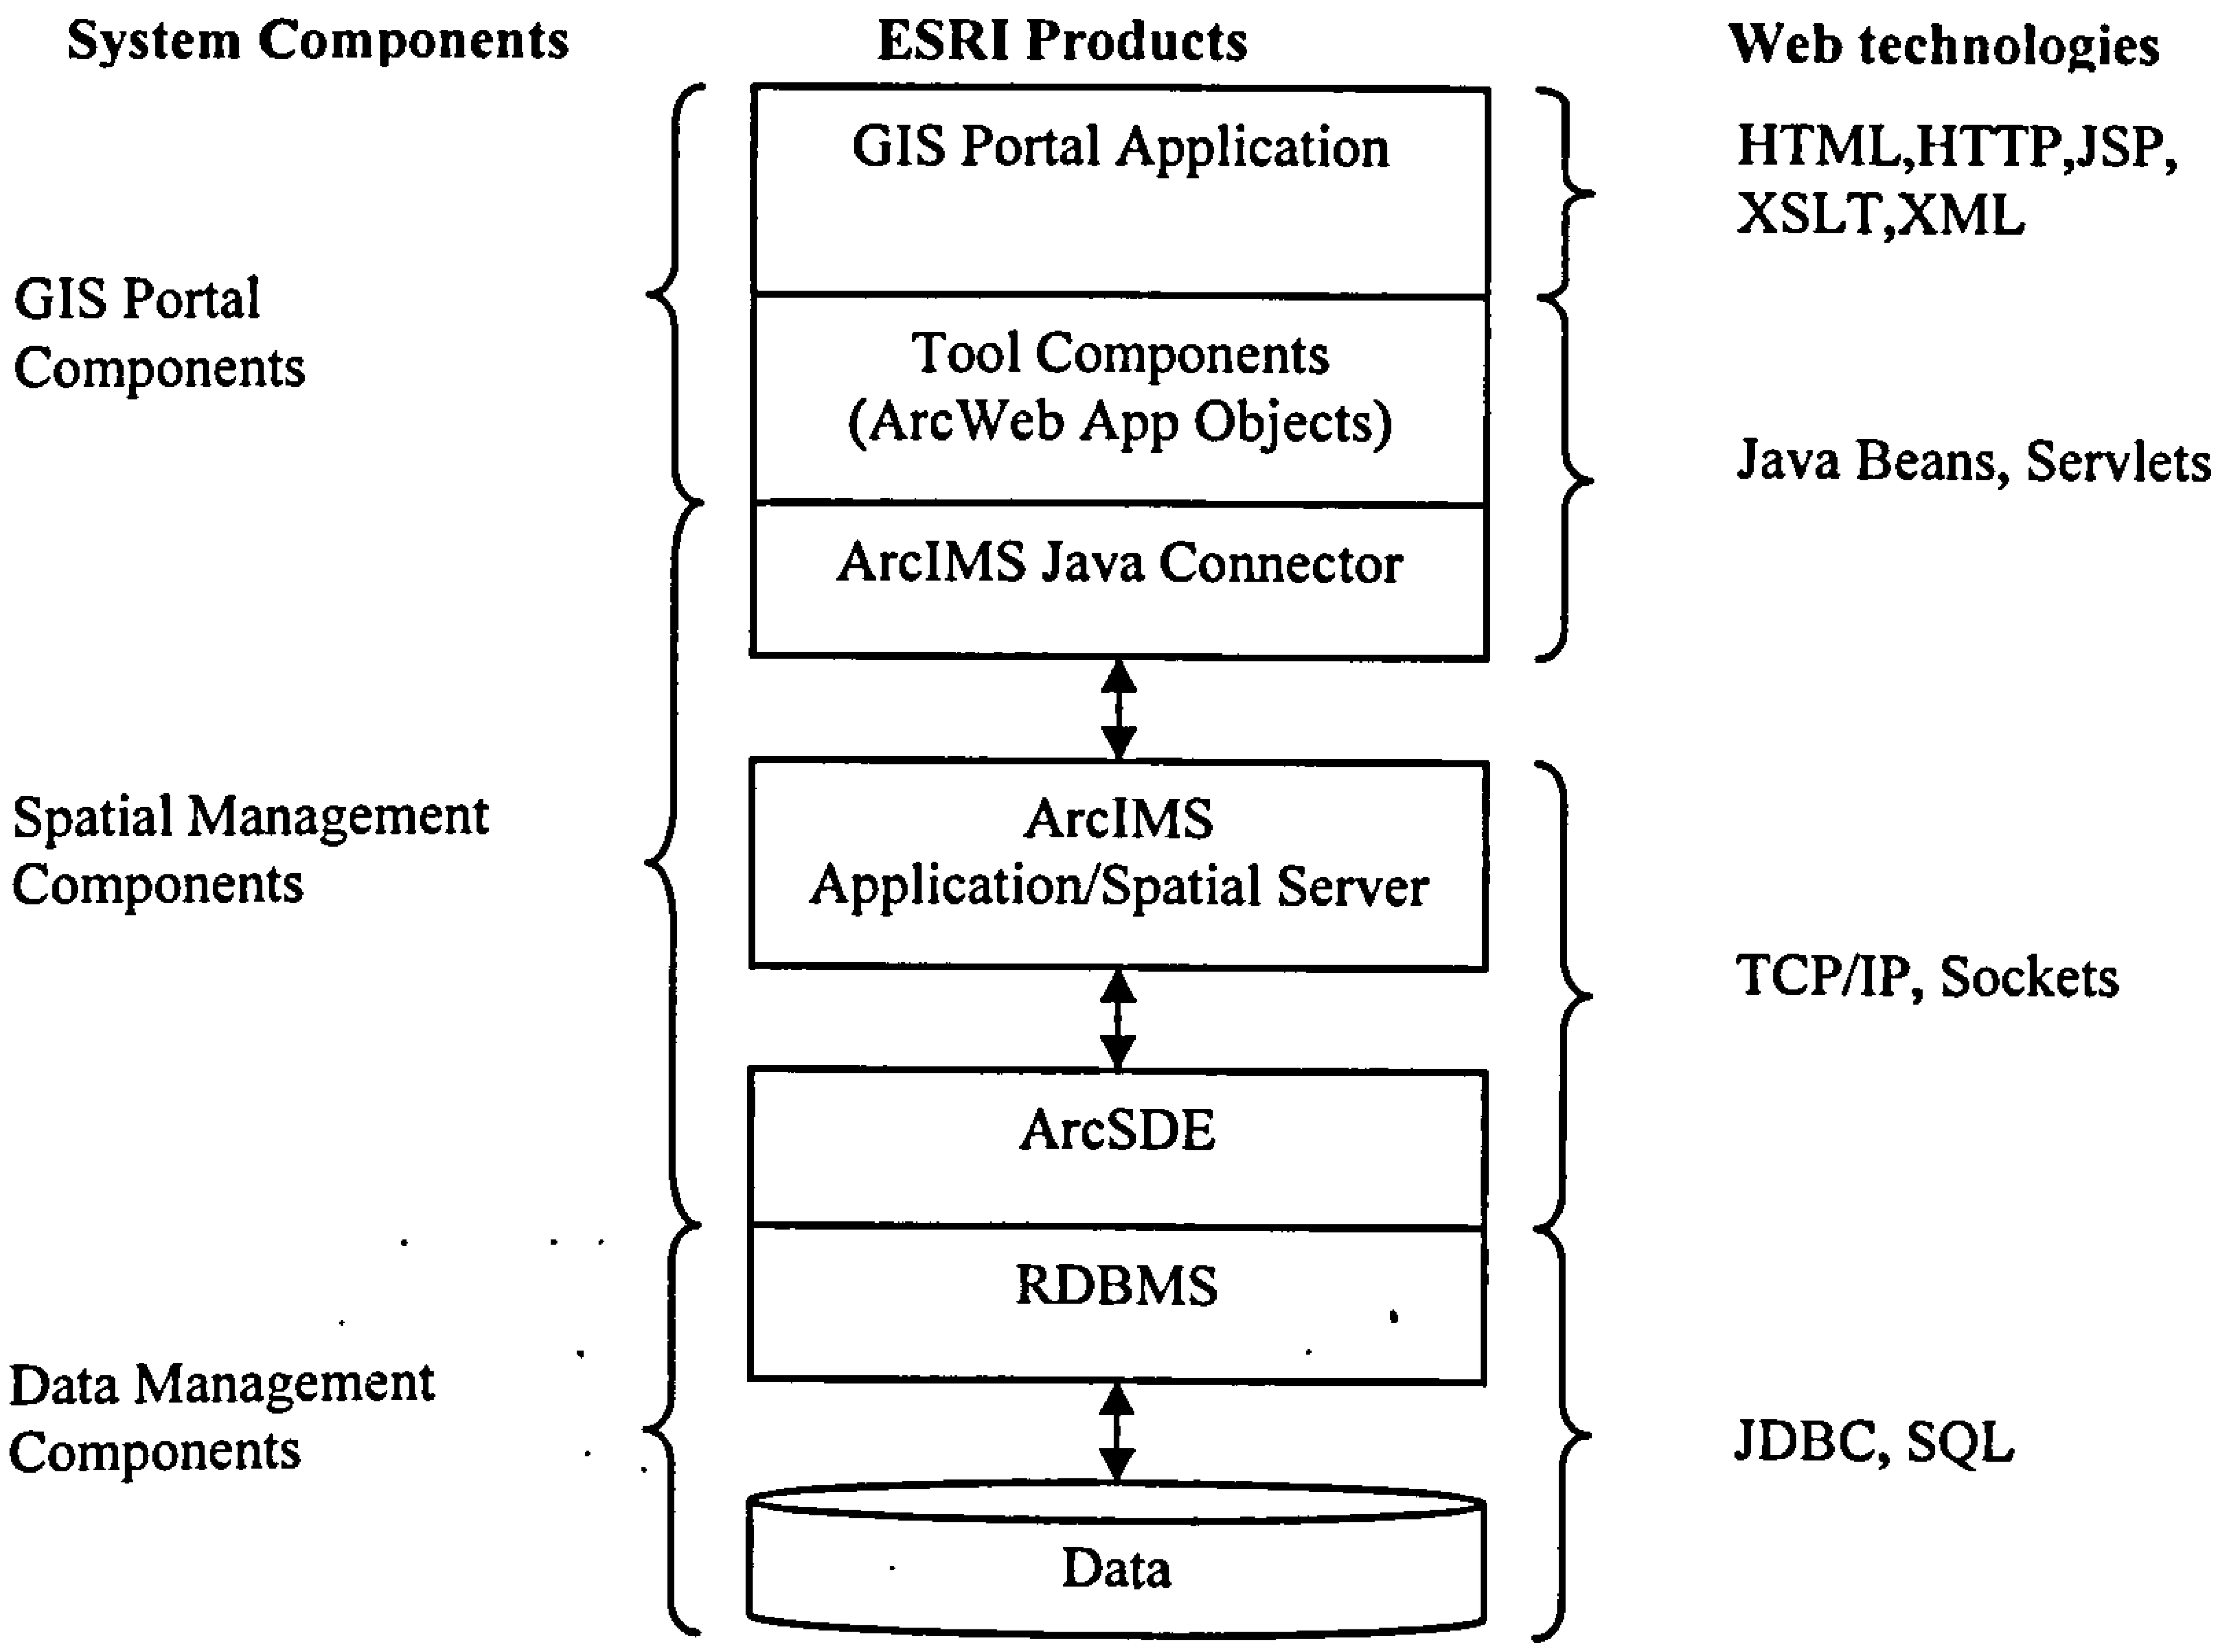

Figure 1.5 Architecture of The GIS Portal Toolkit by ESRI (adapted from ESRI, 2004)

GOS is implemented using ESRI ArcSDE (a spatial data engine that offers spatial capabilities to conventional relational database engines) and ArcIMS (an internet map server for viewing web maps and publishing geographic metadata) (ESRI, 2004). Since the development of the first version of GOS, ESRI has exported the technology to other SDI projects including the European Union INSPIRE and the Indian NSDI projects. The composite product has been named the ESRI GIS Portal toolkit and its architecture is illustrated in Figure 1.5. The illustration identifies the web technologies employed in implementing the aforementioned ESRI products.

1.1.3 Geographic Metadata

A fundamental component of a data discovery system is the repository of metadata. Metadata is defined as descriptive information about data; that can be queried to search for data and resources using space, time and thematic attributes (Maguire and Longley, 2004). The data is organised in fields that have been formalised thus making it possible for information communities, including the geographic information community, to share metadata records. Each metadata record describes any aggregate of information: for example, it can be used to describe a single dataset, a collection of datasets or even a component of a dataset. It allows users to find out what data is available, where to find the data, how to access the data and whether it meets their specific requirements. This section is concerned with geographic metadata and how it is disseminated.

1.1.3.1 Standards

The most widely-adopted geographic metadata standard is the Content Standard for Digital Geospatial Metadata (CSDGM), published by the US Federal Geospatial Data Committee (FGDC). Although this standard is US-specific, most of its elements are still applicable in other parts of the world. Special interest communities have implemented their own profiles of the CSDGM. The profiles make it possible for communities to search for data using properties that are specific to their domain whilst also allowing users from other domains to search the metadata collections using those fields that are common amongst all user groups. Example profiles of the CSDGM include the Biological and Shoreline (Bathymetric) profiles (FDGC, 2005).

An alternative standard for geographic metadata, ISO19115, was published by the technical committee for geographic information of the International Organisation for Standardisation (ISO). The ISO19115 can also be extended through the creation of profiles to allow countries to implement their own specific metadata elements. Unsurprisingly, the ISO standard adopts several fields from the CSDGM, for example the title, abstract, spatial referencing system and others (ISO, 2003). However, some of the fields have been renamed for example, the spatial domain (CSDGM) was renamed as the geographic extent (ISO19115).

Both the ISO19115 and the CSDGM define hundreds of attributes and attribute dimensions(Podolak and Demšar, 2004). We identify the title, abstract, topic category, geographic extent, publication and creation dates as being searchable attributes as they collectively describe the spatial, temporal and thematic scope of a dataset. We therefore discuss them further. First, the *date* of publication or creation offer temporal descriptions of the dataset. Temporal descriptions are particularly important as the historic role of spatial data clearinghouses is to archive data. Unfortunately, both the ISO19115 and the CSDGM allow free-text entry of dates. This means an entry such as "before 2001" or "after 2004" is considered valid. This results in significant variability in temporal metadata, increasing the possibilities of incorrect temporal computations.

The *title* provides a single-line description of the contents of the dataset. The *abstract* provides a brief narrative of the contents of the dataset. It expands on the information provided by the title by i) specifying whether the dataset is part of a larger collection, ii) providing an explanation of the scientific concept behind the dataset, iii) naming which other datasets complement this dataset if it is part of a larger collection, iv) providing a textual description of the location and spatial coverage of the dataset. Both the title and the abstract are specified in 'free text', that is, there are no restrictions to the values of the fields. In addition to the title, metadata includes information on the topic of a dataset. The *topic category* describes the contents of the dataset using a set of predefined themes. This provides a high-level classification scheme that enables developers to implement unsupervised data discovery as all the topics or themes are standardised.

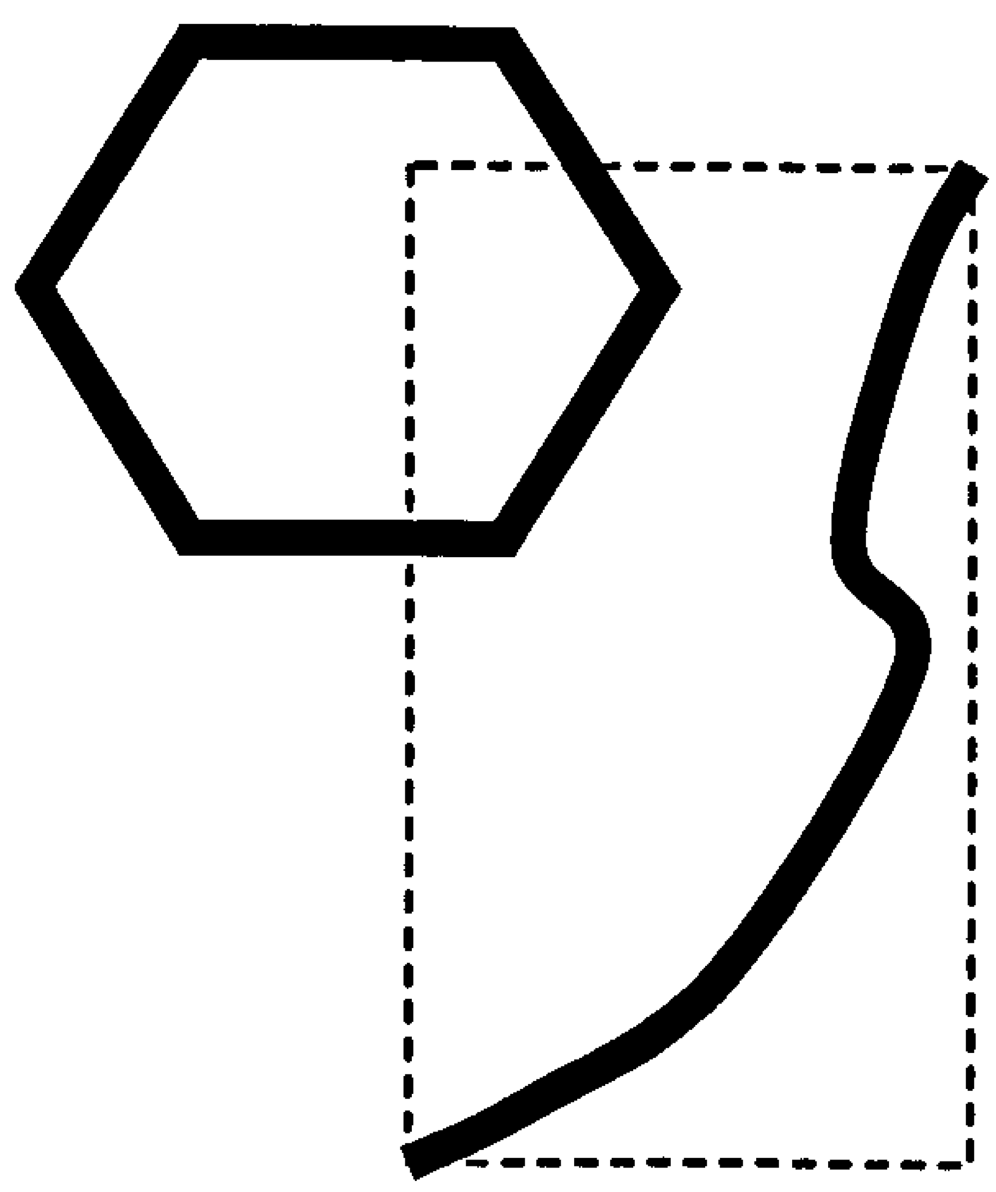

Figure 1.6 Inaccuracy of the spatial operations involving bounding rectangles.

The *geographic extent*, referred to as the spatial domain in the CSDGM, provides an absolute positioning of the spatial footprint of the dataset. This field is usually represented by a Minimum Bounding Rectangle (MBR) of the contents of the dataset, aligned to a geographic coordinate system. However, topological operations using MBRs offer only approximations of spatial relationships (Papadias and Theodoridis, 1997). For example, although the linear feature in Figure 1.6 does not intersect the hexagonal parcel, its MBR does. Spatial operations offered by traditional geographic data discovery systems were based on MBRs as they provided very fast computations of intersections. Further, an MBR reduces the effort needed in creating the actual metadata as only the bounding coordinates East, West, South and North are required during metadata entry. Advances in the development of spatially-enabled Relational DBMS (RDBMS) that can handle Minimum Bounding Polygons (MBP) offer new possibilities for geographic data discovery systems. An MBP is a convex hull of all coordinates in a geographic dataset.

This subsection has presented two geographic metadata standards and identified some of their main metadata fields. The fields provide a geographic reference, a textual and temporal description of a dataset. Other studies discussing geographic metadata include Podolak and Demšar (2004), Nogueras-Iso et al., (2004) and Hodge (2001). The next subsection describes the approaches for geographic metadata dissemination.

1.1.3.2 Metadata Dissemination

Metadata fields of both standards can be grouped according to generic packages such as *identification* or *data quality* information. These packages are listed in Table 1.1. The packages help organise information and form a basis for storing metadata in a DBMS. Two types of database servers have been adopted within the geographic information community for metadata dissemination; relational database engines and Z39.50 servers (FDGC, 2005). Relational database engines such as Microsoft SQL Server, MySQL, PostgreSQL and others allow clients to remotely invoke SQL queries on a database server. Queries are invoked subject to permissions assigned to users by the Database Administrator (DBA). For retrieval-only purposes, a DBA need only allow remote users to invoke SQL *select* statements. Further, relational databases allow for strict data typing and entry of values (i.e. *not null*). Strict data typing ensures that dates and fields are input in a specified format, for example 'YYYY/MM/DD'. The *not null* SQL constraint can be used to ensure that mandatory metadata fields are filled.

CSDGM	ISO19115
• Identification Information • Data Quality Information • Spatial Data Organization Information • Spatial Reference Information • Entity and Attribute Information • Distribution Information • Metadata Reference Information	• Identification Information • Data Quality • Spatial Representation information • Units of Measurement • Extent Information • Reference System Information • Metadata Extension Information • Distribution Information • Citation and Responsible Party Information • Content information • Constraint Information • Portrayal Catalogue Information • Maintenance Information • Application Schema Information

Table 1.1 Packages of the CSDGM and the ISO19115

The second database server used for metadata dissemination involves use of the Z39.50 protocol for Information Retrieval, also known as the ISO23950 standard. This standard is widely used by the geographic information community (FDGC, 2005). Z39.50 servers receive queries from clients and return records that satisfy the query. The records are returned as text-only documents, in contrast to result sets from RDBMS which may contain different types of objects encoded as byte streams. Many current metadata schemes implemented in Z39.50 systems use XML or its superset SGML (Standard Generalized Mark-up Language) (Hodge, 2001). To obtain statistics on the types of servers implemented, we examined over 400 nodes on the FGDC Clearinghouse Registry. The registry lists each Z39.50 node with its connection parameters, spatial footprint covered and several other properties. As shown in Figure 1.7, the results suggested that the CNIDR Z39.50 server is the most popular for spatial data clearinghouses. A possible reason is its inclusion in the freely-downloadable Isite software available from the FGDC website (FDGC, 2005). Unsurprisingly, the ArcIMS metadata server is also widely-implemented. ArcIMS is a web mapping product from ESRI, the leading global GIS vendor.

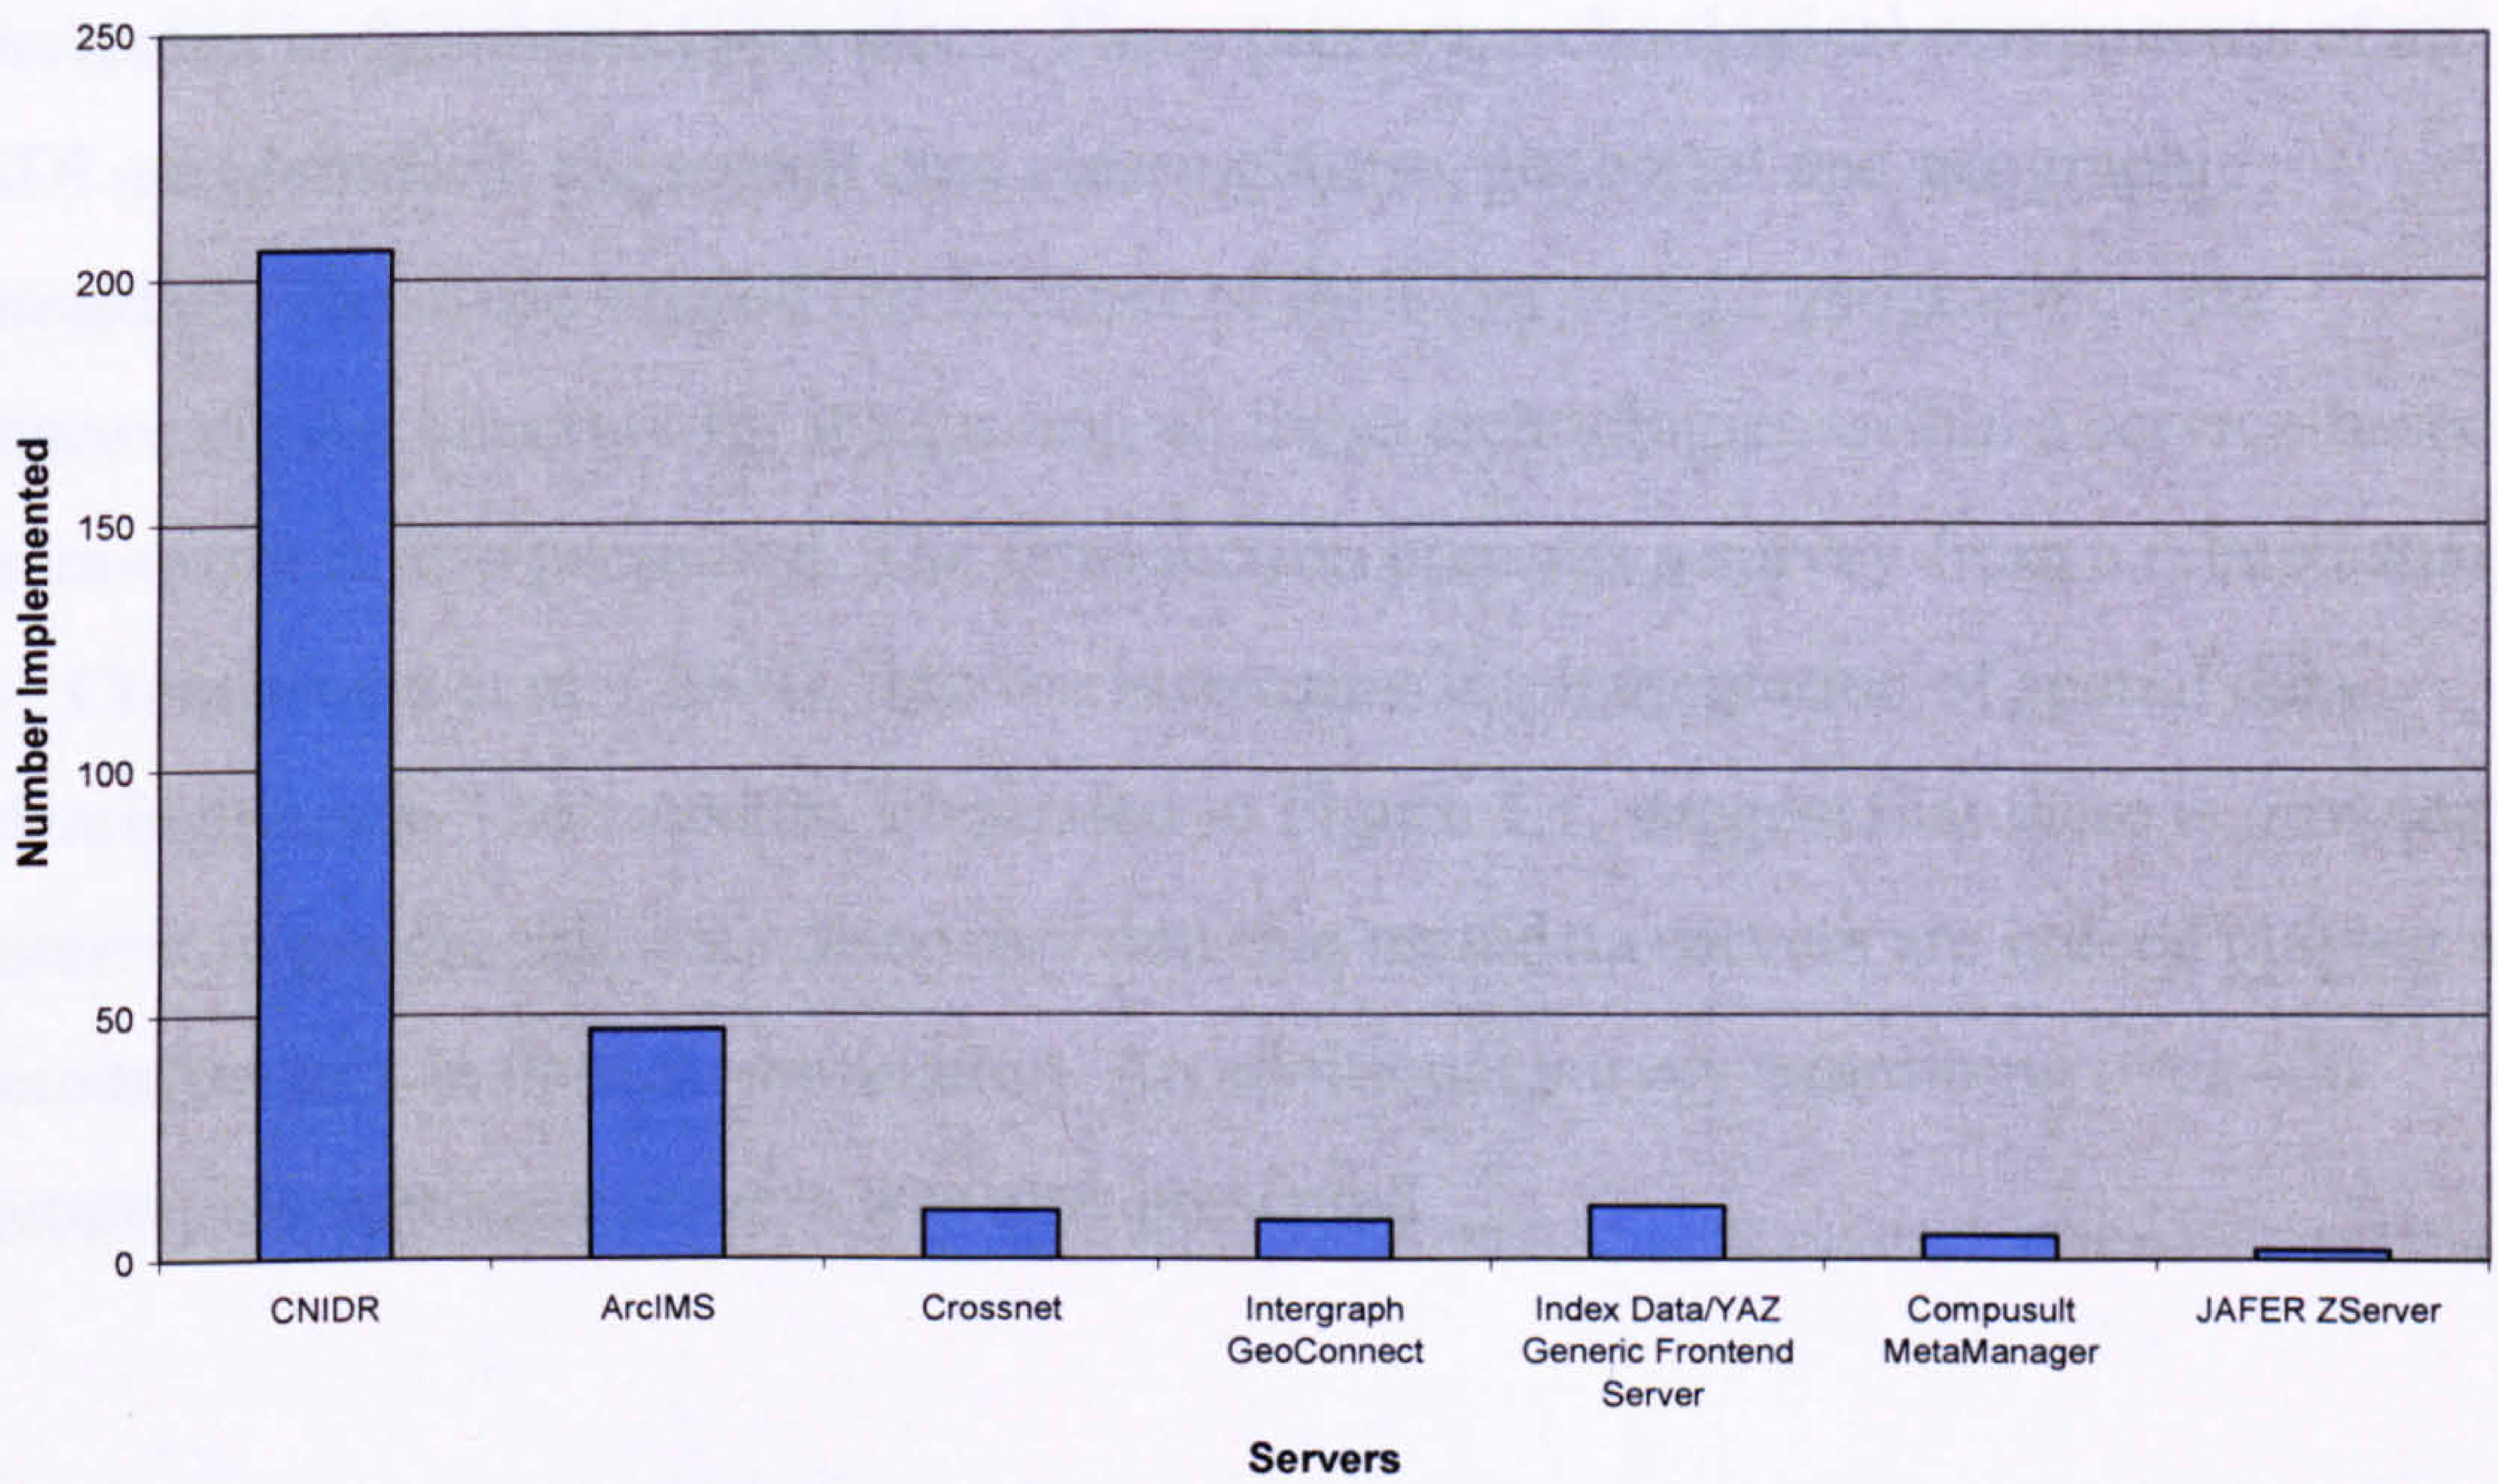

Figure 1.7 Popularity of Z39.50 metadata servers listed on the FGDC registry

Two main architectures for metadata dissemination incorporating the aforementioned database servers include *cross-system* search and *metadata harvesting* (Hodge, 2001). Cross-system discovery architectures allow various clients to access a distributed network of servers to search for records. Metadata harvesting is an automated scheduled process for collecting new and updated metadata from a wide variety of geographic metadata collections (ESRI, 2004). The former approach has the advantage of decentralisation, allowing the client to search up-to-date records. However, the approach relies on referenced servers being available at the time of a user's search. This is in contrast to the latter approach, which stores a copy of each metadata record on a centralised server, thereby ensuring that the metadata is available at the time of search. The latter approach has the advantage that the metadata is well structured as it is harvested and processed prior to searching. The disadvantage is that because metadata harvesting is a scheduled process, the metadata records are accurate only to the last time the harvesting was performed, typically overnight.

1.1.4 Summary

The introduction has examined the technological backdrop within which the research presented in this thesis takes place. Three primary technological components of an SDI are identified: the spatial data clearinghouse, geoportal and geographic metadata. These are singled out because of their key role in geographic data discovery. Architecture for integrating all these technologies within a service-based framework is also presented. The introduction presents a survey from a related study, by Crompvoets et al. (2004), into the increasing implementation of spatial data clearinghouses. Their results, illustrated in Figure 1.1, suggest that there is growing interest in geographic data discovery and that metadata servers are indeed playing an important role in data dissemination. An additional survey examining over 400 geographic metadata servers was also presented.

1.2 Problem Statement

As geographic information resources on the Internet increase, a major challenge facing the Geographic Information community is how to enable a user to efficiently find and retrieve the most relevant resources for geographic tasks. Traditional approaches for presenting the results of geographic search do not adequately illustrate the degrees of relevance of geographic datasets (Beard and Sharma, 1997, Göbel et al., 2002). Consequently, important resources are not discovered during searches, leading to higher data acquisition costs. A new framework that effectively illustrates the relevance of multidimensional geographic datasets, and allows a user to examine each dataset in greater detail, is needed. To increase the possibility of acceptance within the Geographic Information community, the new framework should incorporate some of the existing approaches.

1.3 Research Questions

The hypothesis of this thesis is:

Geoportals offering ontology-assisted multidimensional visualisation services can allow users to discover and retrieve more relevant geographic information resources than those that do not offer these services.

This hypothesis can be divided into the following research questions:

- What are the limitations of existing approaches in the visualisation of candidate datasets during geographic data discovery?
- How can ontology be used to support visualisations of geographic search results?
- From search results, how can a multidimensional geographic dataset be visualised in greater detail to determine its relevance?
- Could the suggested approaches be incorporated into a conventional geoportal?
- What effect would the suggested approaches have on the performance of a geoportal?

This study therefore aims to provide a new framework for the development of geoportals that uses ontology-assisted multidimensional visualisation approaches to improve geographic data discovery and offer interoperable delivery of heterogeneous geographic datasets. It is envisioned the framework will enable users to, more effectively and efficiently, find and retrieve geographic datasets that are relevant to their information needs. To ensure applicability in an international setting, the proposed model will be developed from open-source technologies as proprietary technologies may prove expensive for potential user communities. Further, open-source development would offer the possibility of incorporating any relevant projects from the tens of thousands of other open-source projects available online (Gayle and Rainer, 2005).

1.4 Methodology

This research comprises of the following three main parts:

Analysis: An analysis of the 'state of the art' in geographic data discovery and dissemination sets this study in context. This includes a consideration of the concepts, standards, technologies and other issues involved in the discovery and delivery of geographic data. This aspect of the study addresses the research question "what are the limitations of existing approaches in the visualisation of candidate datasets during geographic data discovery?".

Implementation: This aspect of the study proposes two visualisation approaches that address the research questions "how can ontology be used to support visualisations of geographic search results?" and "how can a multidimensional geographic dataset be visualised in greater detail to determine its relevance?". The proposed visualisation approaches are then implemented in prototypes that include i) a web-based ontology-assisted visualisation tool for searching a geospatial data clearinghouse ii) a web-based tool for multidimensional geo-referenced visualisation and data delivery iii) a geoportal that integrates all other tools developed in this study into a single system.

Evaluation: A variety of evaluation studies are carried out to assess the effect of the proposed visualisation approaches as applied to GIR. The effect of the proposed approaches is reflected by the retrieval performance and usability of the implemented prototypes. The evaluation studies, therefore, address the research questions "could the suggested approaches be incorporated into a conventional geoportal?" and "what effect would the suggested approaches have on the performance of a geoportal?".

1.5 Scope of the Thesis

This thesis centres on the development of mechanisms for the web-based discovery and delivery of multidimensional geographic information. Since geographic information can be disseminated on the internet in a variety of forms, this thesis limits its handling of geographic information resources to geographic datasets and associated metadata conforming to the International Standards ISO19125 and ISO19115 respectively; the former standard being the Simple Features model (ISO, 2004) and the latter, the geographic metadata standard (ISO, 2003). As the tools developed within this study for handling these geographic information resources are intended for different purposes, they are evaluated separately and independently of one another.

A major consideration in this thesis is the transferability of proposed models to existing SDI. As a result, the experimental work uses 'real-world' geographic data collections collected from national spatial data clearinghouses and mapping agencies. Different data publishers have created some of these collections manually; consequently they may contain errors in accuracy of content. For example, some metadata records have erroneous geographic footprints. Errors in metadata may cause some datasets to be missed during searches, or irrelevant datasets to be included in specific search results. Although, these errors in metadata content are undesirable, they are an inescapable element of geographic (meta)data created manually.

1.6 Organisation of the Thesis

Subsequent chapters are organised as follows:

Chapter 2 describes the state-of-the-art in Geographic Data Discovery and Geographic Information Retrieval (GIR). In particular, the chapter discusses concepts within GIR including ranking, relevance and ontology-assisted queries.

Chapter 3 introduces the concepts of information and data visualisation, with reference to geographic data. An approach for presenting the results of a geographic search is proposed. Also, an approach for web-based visualisation of geographic datasets is proposed. A discussion of real-time computer graphics is also offered.

Chapter 4 introduces concepts in the web-based delivery of multidimensional geographic data. Two delivery mechanisms are examined, in particular, through connection to a spatially-enabled Relational Database Management System and through connection to service-based data delivery system.

Chapter 5 describes the design and implementation of the two visualisation prototypes presented in this thesis. The chapter also describes how they are incorporated into an existing geoportal framework.

Chapter 6 provides an evaluation of the two prototypes presented in this study. The chapter also presents observations from usability studies of the completed geoportal framework.

Chapter 7 highlights the contributions of the study, presents the conclusions and suggests some possible areas for further research.

Chapter 2 Geographic Information Discovery and Retrieval

The aspect of Geographic Information Science (GISc) that studies spatial data discovery is referred to as Geographic Information Retrieval (GIR). Bucher *et al.* (2005:pp 1) define GIR as "the retrieval of geographically and thematically relevant documents in response to a query of the form <theme,location>, where the spatial relationship may either implicitly imply containment or explicitly be selected from a set of possible topological, proximity and directional options, and where documents searched are those available on the internet". In short, GIR is "concerned with providing access to geo-referenced information sources"(Larson, 1996). GIR is sometimes considered to be an integration of traditional Information Retrieval (IR) with GISc. This chapter offers a review of the literature on GIR with respect to the discovery of geographic data.

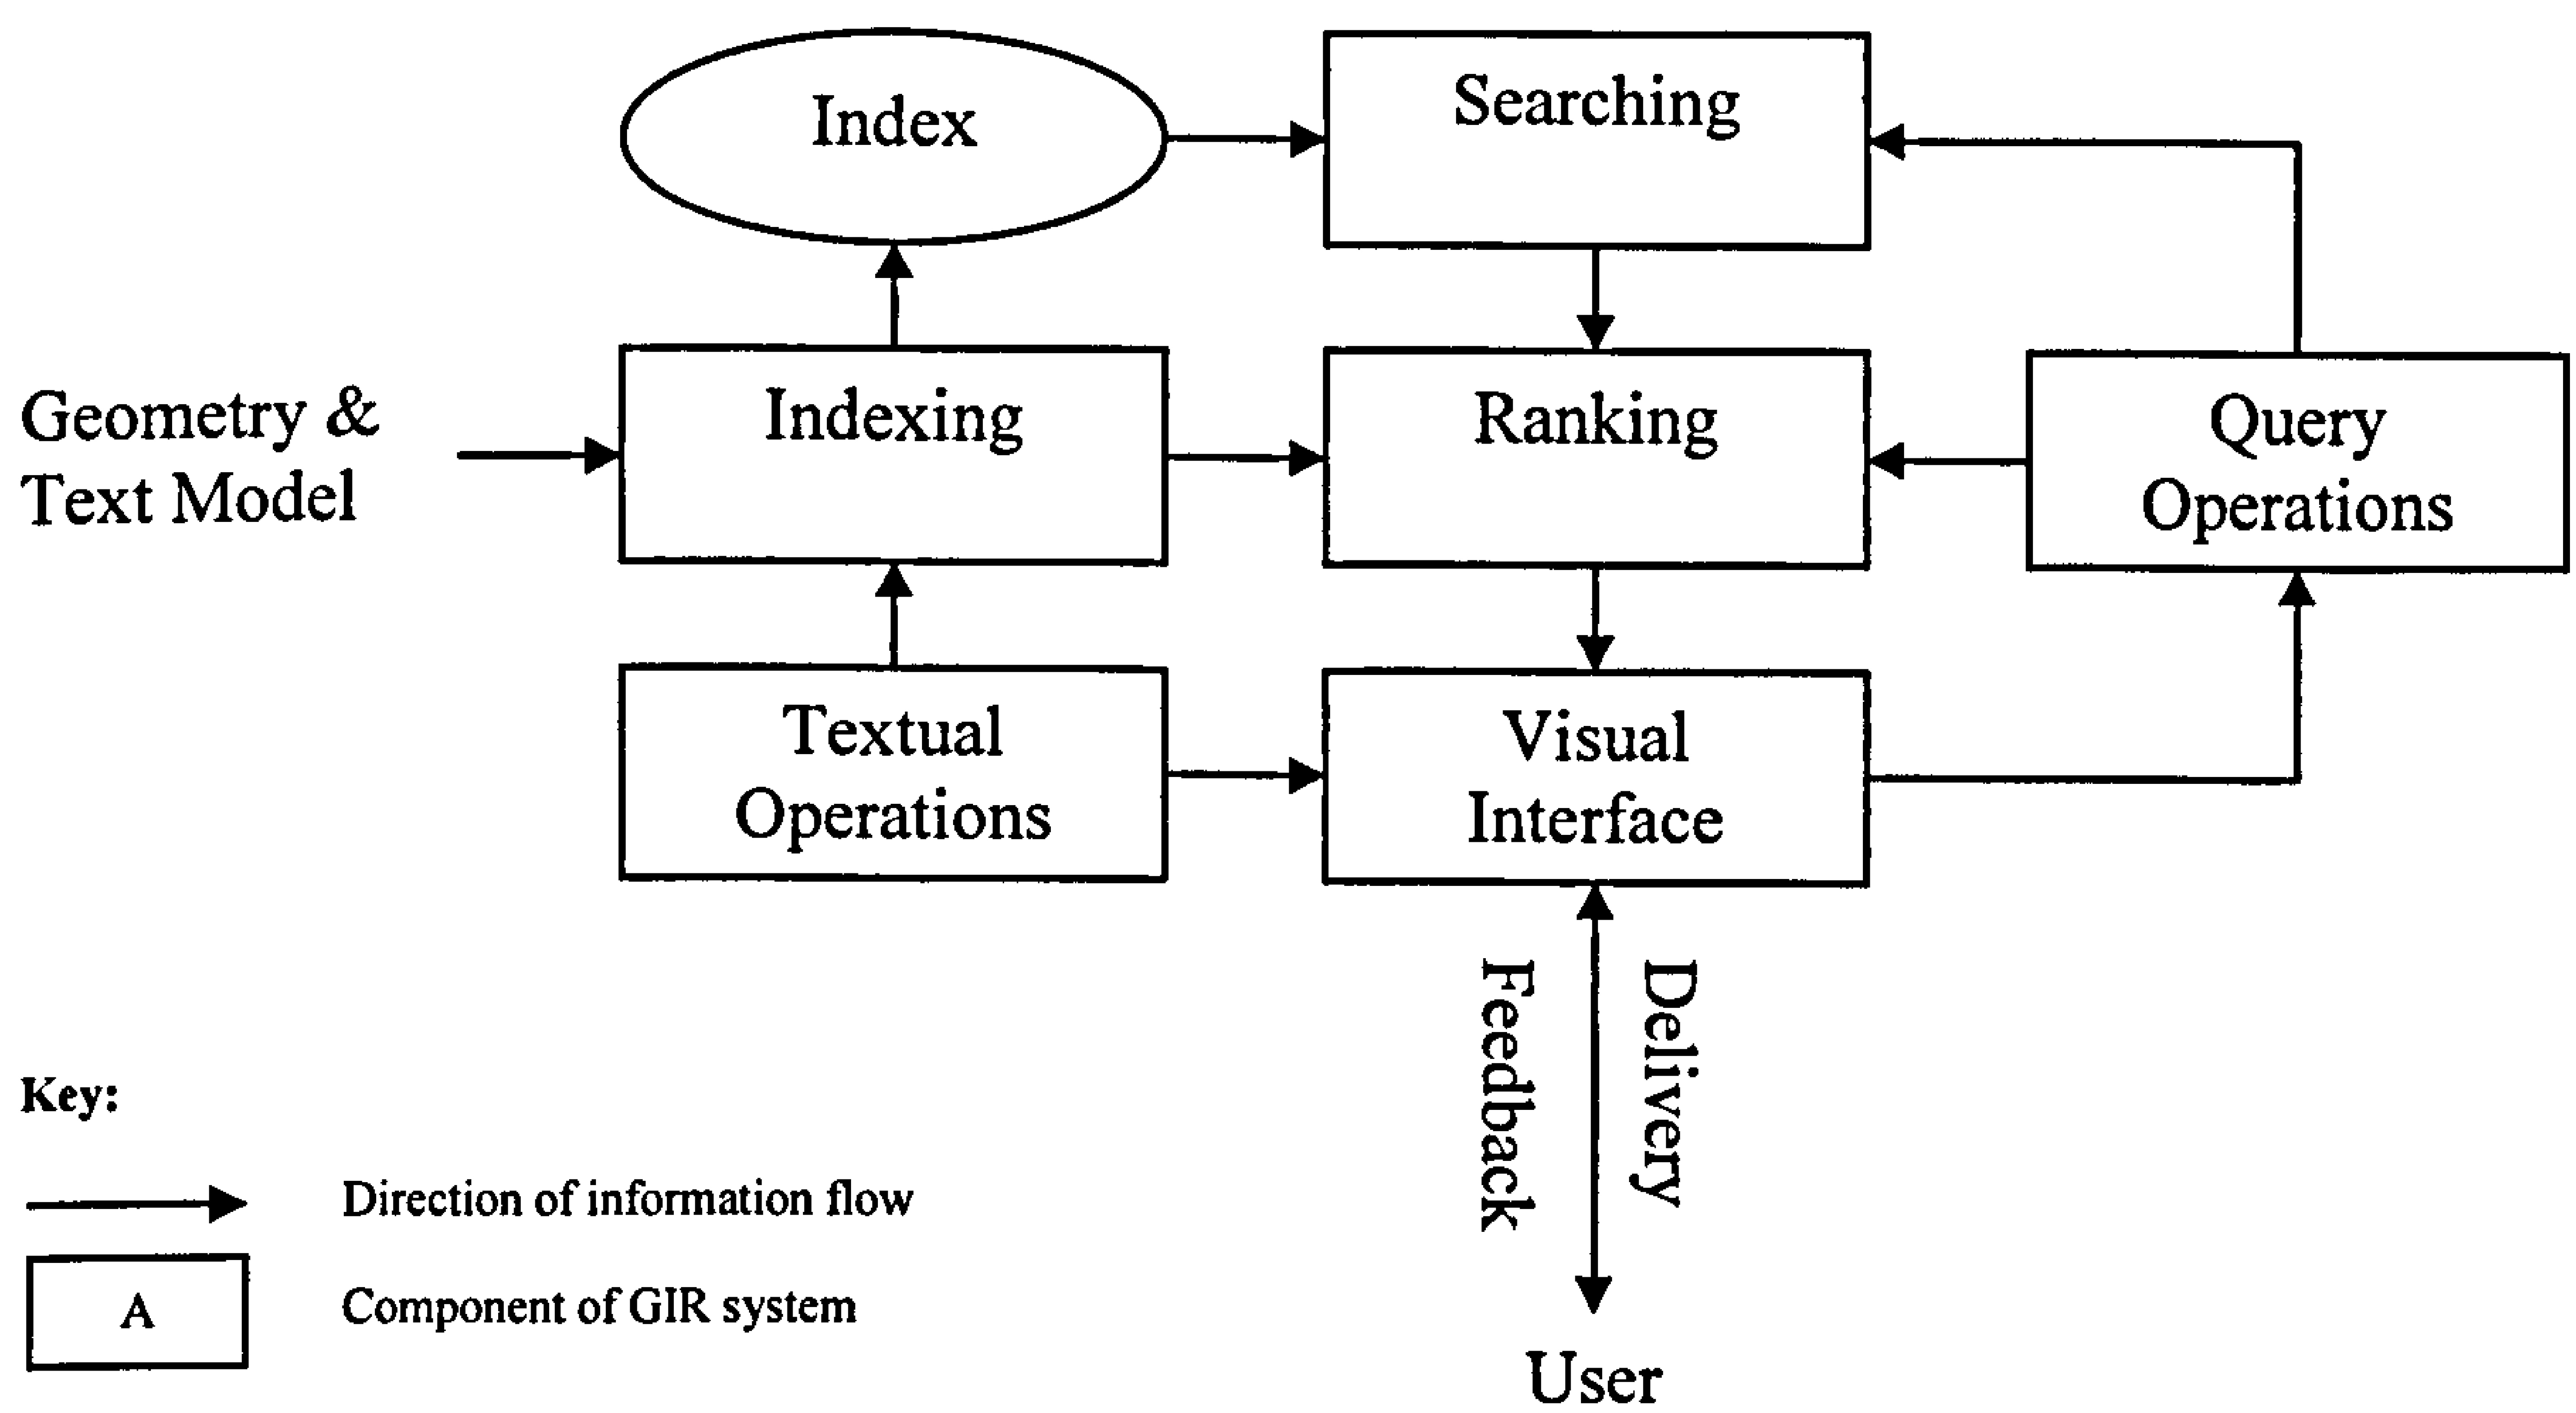

Figure 2.1 The process of retrieving information. (adapted from Baeza-Yates and Ribeiro-Neto, 1999)

The cyclic nature of the process of information retrieval is illustrated in Figure 2.1. The illustration acknowledges the part played by the user in finding a suitable dataset, as the user provides feedback to the system in the form of refined keywords or altered query constraints. The *query operations* invoke the *searching* mechanisms which, in turn, retrieve information from the *index*. The retrieved information, usually metadata, is then arranged according to increasing or decreasing relevance by the *ranking* facility before ultimately being presented to the user on the *visual interface*. This importance of the visual interface is shown by its role of being the bridge between the user and the rest of the IR system. The rest of this chapter discusses the different stages of IR illustrated with Figure 2.1.

2.1 Ranking through Relevance

2.1.1 The Concept of Relevance

A fundamental concept in IR is 'relevance'— the degree to which an item is deemed to satisfy a user's query. Previous research has shown that there are variations in people's perceptions of relevance (Harter, 1996). "An overwhelming body of evidence has accumulated toward the conclusion that the human beings involved in the information retrieval process (indexers, searchers and users) and the products they produce (indexing records, retrieved documents, queries, and relevance judgments) vary enormously from one another" (Harter, 1996: pp. 48). His summary of factors affecting a user's assessment of relevance lists over fifty possible parameters in a user's judgement. These parameters include the title of the document, the author's credibility, the searcher's intelligence and several other factors. A title describes the content of a document; therefore, directly influences a decision on the document's relevance to a task. However, a searcher's ability to understand the description of the document is also a factor in a decision on relevance. The credibility of the author (for example, is the author an expert on a specific subject) influences the searchers judgement of relevance as well.

Similarly, in her work Borlund (2003a) examines the different definitions of relevance by other researchers. She then classifies them into two main groups, system-based relevance and human-based relevance. System-based relevance refers to a relation between a query and a set of information objects as computed through a specified algorithm. Whereas, human-based relevance refers to the judgement made by the searcher through his own individual mental experience (Borlund, 2003a). As the reliability of system-based experiments is generally based on the repeatability of the experiments, system-based evaluation of relevance tends to be static and objective. In contrast, human-based relevance tends to be dynamic and subjective, as it is a function of the searcher's current cognitive state, situational constraints and the type of information problem.

A recent discussion of geographic relevance is provided by Bucher et al., (2005). They observed six scenarios influencing the determination of geographic relevance. Three of the situations are with respect to thematic relevance and the other three are with respect to spatial relevance. Although their discussion does not include temporal relevance, it supports the view that geographic information relevance is multidimensional in nature. In their study they compare relevance judgements by 11 users, based on 5 queries, each query returning 10 documents. A scheme based on the six scenarios was used to evaluate the relevance judgements made by the users. They observed that although users found the scheme easy to understand, they were less confident in judging spatial relevance particularly when users were not familiar with a location. The scheme they used is as follows:

Thematic relevance:

1. A document which contains relevant information about the concept queried and on its own allows you to form a judgement about the document (ie. Requires no external knowledge).
2. A document is relevant, since it points to a resource mentioning the concept but you must consult further pages referenced by the document to perform a judgement.
3. A document does not provide information about the concept provided.

Spatial relevance:

1. A document refers to a location that is in/near the query location and you think that the location in the document has sufficient detail for you to find on a local map of the area

2. A document refers to a location that is in/near the query location but you think that there is insufficient information for you to find that location on a local map of the area.
3. A document does not fall within the query location.

This subsection has examined the views on relevance with respect to geographic data. Key observations include the suggestion that there are variations in how different individuals perceive relevance, suggested by Harter (1996). A classification of relevance into system and human-based assessments is offered by Borlund (2003a). Other studies referenced in this subsection highlight the multidimensionality of the relevance of geographic data. Based on the works discussed in this subsection, we take the view that the relevance of geographic data is indeed a multidimensional property, however it can be categorised into three components namely spatial, temporal and thematic properties (Beard and Sharma, 1997), as per the definition of geographic data suggested by Longely et al. (2001). The next subsection discusses how this relevance is inferred in IR and GIR.

2.1.2 Ranking Geographic Data

The process of computing and ordering the retrieved items in a search is commonly referred to as ranking. The context of ranking within the IR process is illustrated in Figure 2.1. Several algorithms and approaches to relevance ranking have been developed over the years. Examples of ranking algorithms in existence include those for ranking simple text, hypertext, images and spatial documents (van Kreveld et al., 2005, Lynch et al., 2004, Brin and Page, 1998). Several other ranking algorithms exist, however a discussion of all of them would be beyond the scope of this thesis.

An example of ranking algorithms includes the PageRank algorithm for ranking hypertext documents, developed by Brin and Page (1998). This patented algorithm forms the basis of the Google search engine which has become increasingly popular over the past decade. The algorithm computes relevance of a hypertext document based on how many other documents reference that page. It treats each link to a hypothetical page B from page A to be a vote for page B by page A. Additional textual mining is carried out on the page to determine the relevance of content to the prescribed query. Issues taken into consideration include the size of text, for example words in bold or presented in bigger sized fonts are given a higher weighting than normal text. Other issues the algorithm takes into consideration include the possibility of the user 'to get bored' during the search. This is represented in the algorithm by the damping factor. The damping factor therefore adds an element of personalisation to the ranking approach. The PageRank for a page labelled A, is therefore

$$PR(A) = (1-d) + d\,(PR(T_1)/C(T_1) + \ldots + PR\,(T_n)/C(T_n))$$

where :

n is a positive number

$T_1 \ldots T_n$ are pages referencing Page A

d is a damping factor between 0 and 1

$C(A)$ is the number of pages referenced by Page A

Even after the PageRank has been calculated, further evaluation of the returned documents is carried out. This involves counting the number of times the searched terms occurred in each of several types (the title, abstract, large plain text, small plain text and so on). Each type is assigned a type-weight. The number of times the search term occurs on each of the types contributes to the count-weight. The type-weights and the count-weights make up two separate vectors. The dot product of the vectors results in the IR score for the document which is combined with PageRank to give the final rank for the hypertext document. Multi-worded queries however result in multiple result sets. The proximity of the result sets is based on a ten-unit scale from "not even close" to "direct match". Proximity is included in the type-weight to give a type-prox-weight. The IR score is then computed from the dot product of the count-weights and the type-prox-weights. The PageRank algorithm is implemented in the Google search engine, the most widely used search engine, but does not take into consideration the spatial properties of documents. However, geographic data includes both attributive and spatial properties.

An example of a ranking algorithm that considers spatial properties is the Multidimensional Scattered Ranking approach (MSR)(van Kreveld et al., 2005). Taking the vector of spatial and thematic relevance scores as points in multidimensional space, as illustrated in Figure 2.2, scattered ranking lets the proximity between points be the Euclidean distance between the points in multidimensional space. In its basic form scattered ranking favours a document that is furthest away from the already ranked documents. Naturally, the seeding document is the one with the smallest Euclidean distance from the query. As in other IR models, this document is considered to be the most relevant to the query. Given two collections R and U, containing ranked and unranked documents respectively, the next document to be ranked after the seeding document is the one that is furthest from the seeding document. Once this document has been discovered and added to the set of ranked documents R, it is then deleted from the set of unranked documents U. All consecutive rankings are relative to the last ranked document.

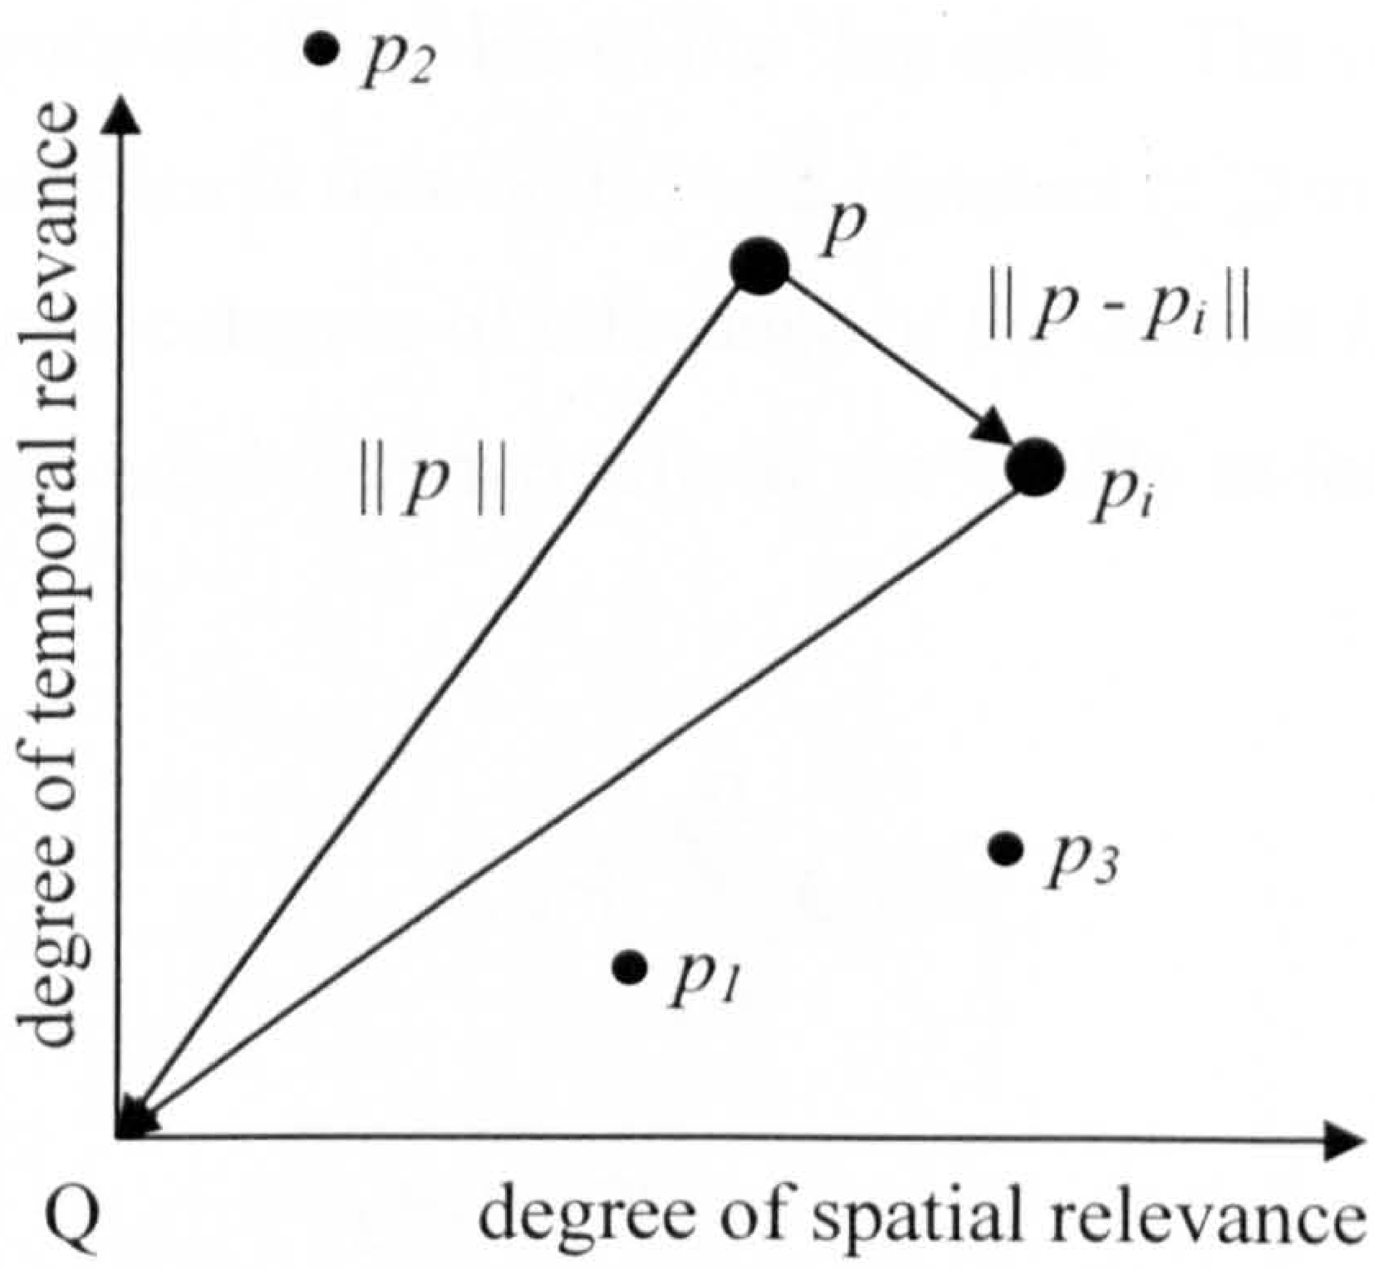

Figure 2.2 Documents scattered in vector space according to their spatio-temporal relevance (adapted from van Kreveld, 2004)

Therefore, within vector space the rank is,

$$S(p,R) = \min_{p_i \in R} \left((1 - e^{-\lambda \bullet \| p - p_i \|}) \bullet \frac{1}{1 + \| p \|} \right)$$

where:

- p is the modulus of the distance from the query to the unranked document
- p_i is the distance from the query to the closest ranked document to p

- $p - p_i$ is the modulus of the distance from an unranked document to its closest ranked document
- $||x||$ is the modulus of a number x
- R is the set of ranked documents
- λ is a constant, greater than zero, that defines the base e^{λ} of the exponential function
- Q represents the query in vector space

Geographic data may also be ranked based on the amount of spatial overlap as suggested by Larson and Frontiera (2004). Their approach uses Logistic Regression (LR) to rank datasets according to their spatial footprints. In this LR model the ratios of overlap of the query region and the candidate region are assigned to each independent variable (X_i) of the 'log odds', *log O(Q, D)*. By multiplying the independent variables by weighting coefficients (C_i), each variable can be given a degree of influence on the value of the 'log odds'. The sum of the weighted independent variables is then added to a constant (C_o) to produce the 'log odds' which represents the degree of relevance of the dataset D to the query Q. The 'log odds' is then converted to a normalised probability to form the measure of spatial similarity.

$$\log O(Q,D) = C_0 + \sum_{i=1}^{m} C_i X_i$$

$$P(Q, D) = \frac{e^{\log O(Q,D)}}{1 + e^{\log O(Q,D)}}$$

where:

- m is the number of ratios of overlap (in our case, $m = 2$)
- X_1 = area of overlap (query region, candidate region) / area of query region
- X_2 = area of overlap (query region, candidate region) / area of candidate region
- C_i are coefficients for assigning weights to variables
- P is the probability of relevance

The concept of overlap is also adopted for time-based ranking of datasets as proposed by Yamuna and Candan (2000). Within this temporal model, the amount of overlap between the time specified in a query and that specified in the metadata record is determined. From the illustration in Figure 2.3, the similarity between periods A and B is the overlapping period C. Hence two periods starting and ending at the same time 'are completely similar'. The implementation of this model within a geographic data discovery framework is facilitated by the fact that geographic metadata specify the 'time period of content' as a mandatory metadata field. This field may be represented either as a single date or as two dates representing the beginning and end of dataset creation (ISO, 2003). Yamuna and Candan (2000) define the temporal similarity (*ts*) of two periods of time A and B as:

$$ts(A,B) = \frac{|C|}{|A| + |B| + |C|}$$

where:

- A, B are two independent periods of time
- C is the proportion of overlap between A and B

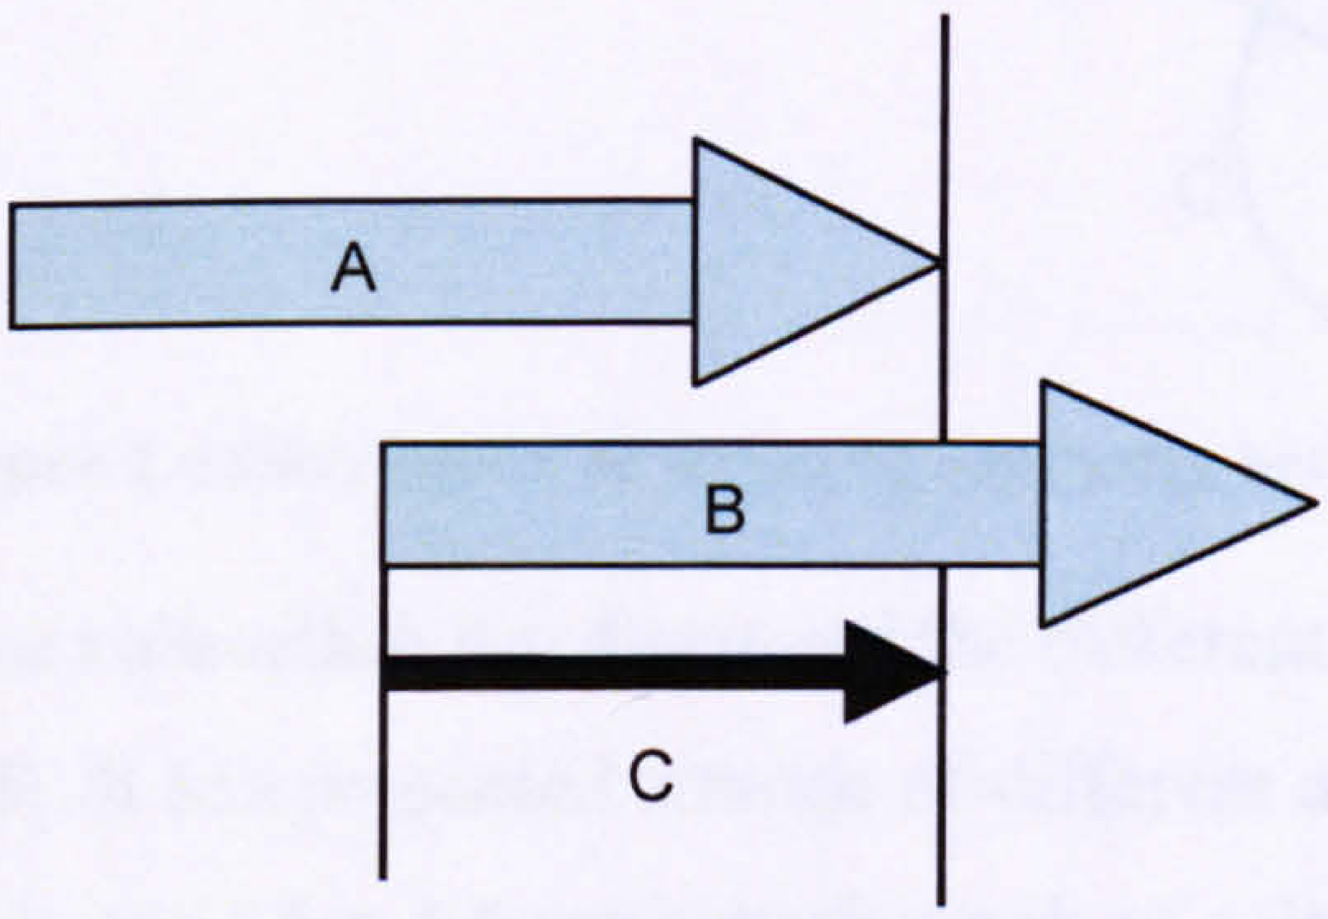

Figure 2.3 Temporal Ranges and overlap

Early studies in the psychology of similarity developed models for quantitatively estimating the similarity between two entities. Tversky (1977) proposes a set-based ratio model for measuring similarity. Rodriguez and Egenhofer (2004) adopted and successfully applied this model to geographic entities. From the illustration in Figure 2.4, set A is more similar to set B than set C, because they have more similar properties. Each similar property (element) is deemed as being a vote for the similarity between sets. Each set in this example represents an individual entity. Weightings are assigned to reflect the asymmetric similarity of A to B. This means that the similarity of A to B may not be the same as the similarity of B to A. In set theory notation, the ratio model for calculating semantic similarity *S* is:

$$S(A,B) = \frac{|f(A \cap B)|}{|f(A \cap B)| + \alpha |f(A-B)| + (1-\alpha)|f(B-A)|} \quad \text{for } 0 \leq \alpha \leq 1$$

where:

- A and B are two separate sets
- the set of all features that are common to both A and B is (A∩B)
- the set of all features that belong to A and not to B is (A−B)
- α and ($1-\alpha$) are weightings reflecting the asymmetric influence of A and B on the similarity value

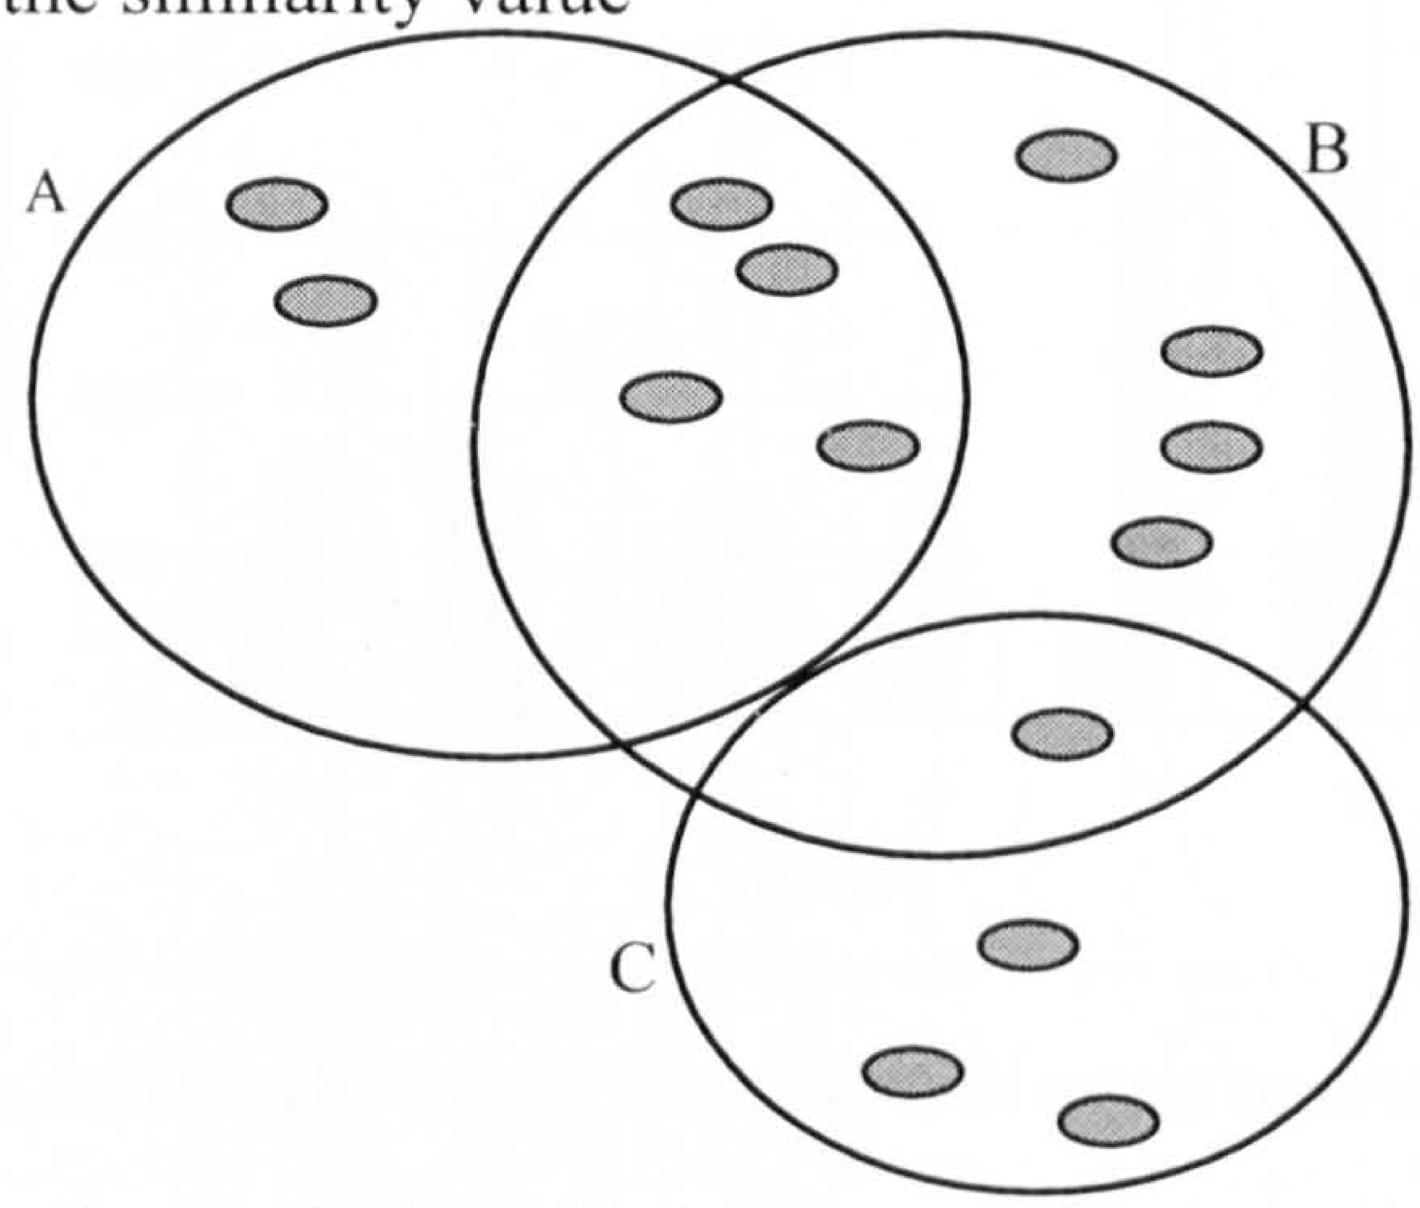

Figure 2.4 Illustration of semantic similarity between entities using set notation

This subsection has discussed the different aspects of ranking with regard to IR and GIR. It has presented a range of different approaches for computing the relevance score used for determining the rank of a dataset. Once datasets have been ranked, they are then presented to a user through a visual interface. The next section discusses the role of visual interfaces in the process of IR.

2.2 Visual Interface

One of the stated research questions is "what are the limitations of existing approaches in the visualisation of candidate datasets during geographic data discovery?", this section will address this question. As already observed in Figure 2.1, the visual interface is responsible for communicating the results of the search to the user. Most IR user interfaces offer forms for entering search parameters. After a search has been processed, the results are presented to the user. Traditional approaches for the presentation of search results were based on the ranked textual list; however, a textual list offers a one dimensional view of relevance. Therefore, in a GIR environment the spatial, temporal or semantic relevance scores are merged into a single relevance score when presented on a ranked list. We illustrate this limitation with the following example.

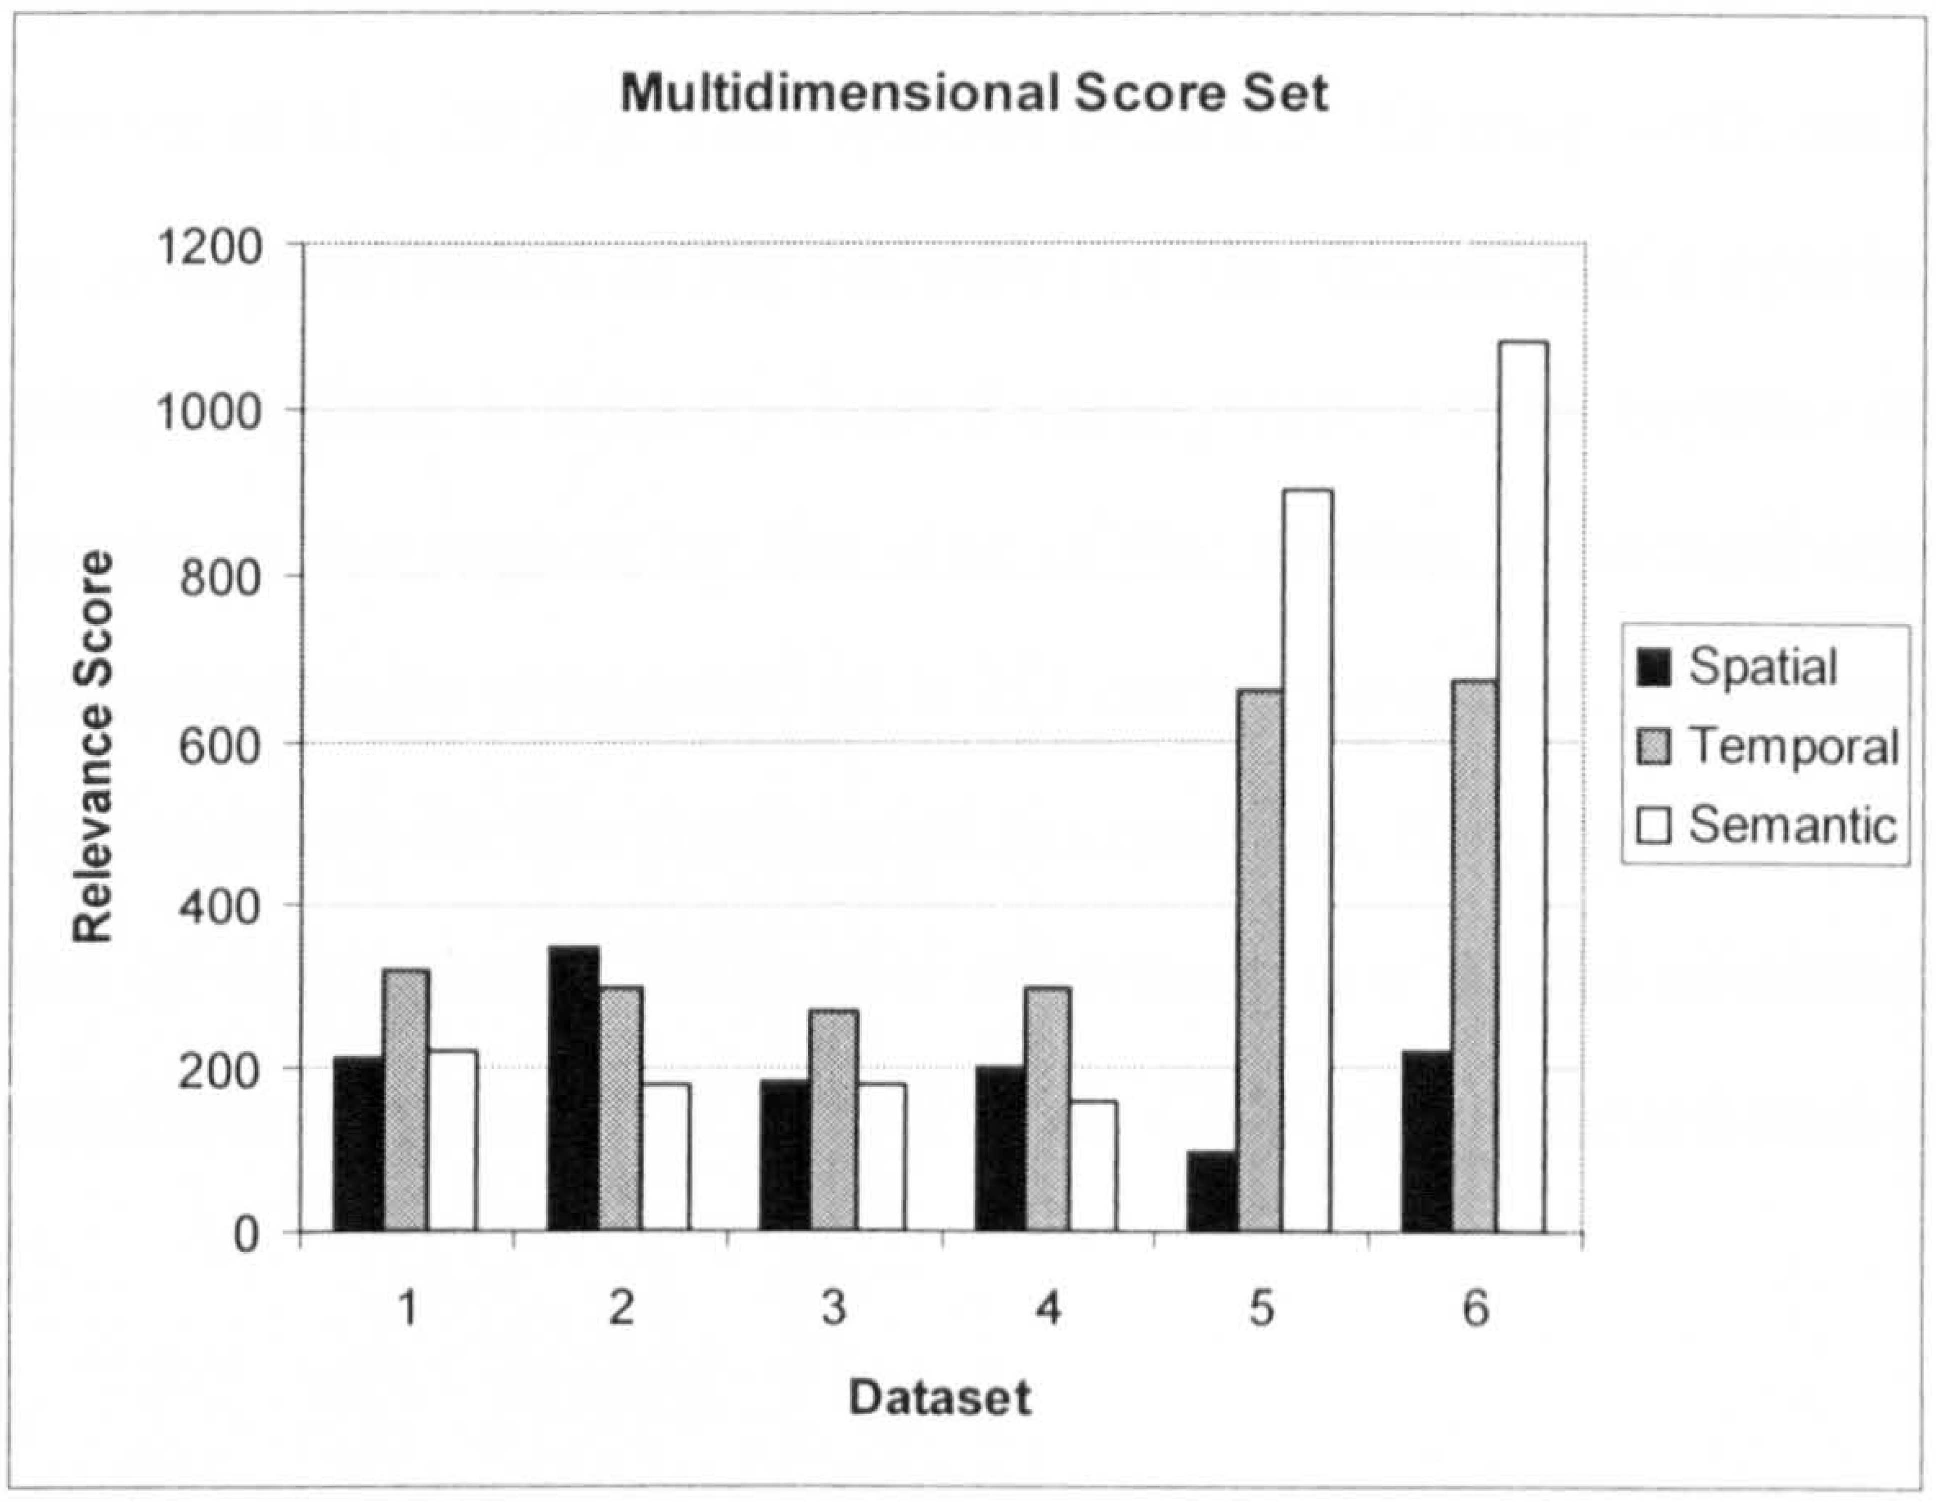

Figure 2.5 An example set of multidimensional relevance scores (associated with Table 2.1)

As an example, Figure 2.5 and table 2.1 present the spatial, temporal and semantic relevance scores of six datasets. Table 2.1 is a linear addition of values from the graph in Figure 2.5, where the spatial, temporal and semantic relevance scores are multiplied by 5, 10 and 20 respectively. From the illustration, it can be seen that a one-dimensional representation of values (presented in the table) can lead to a misinterpretation of the individual relevance measures of the spatial, temporal and semantic properties of the datasets (illustrated in Figure 2.5). For example, dataset 5 has the lowest spatial relevance according to Figure 2.5. However, because of the weightings used in the linear model it appears to have the second highest integrated degree of relevance. This becomes a limitation when a user is particularly interested in individual relevance scores rather than the composite relevance score.

Title	Score
Dataset 6	1970
Dataset 5	1660
Dataset 1	755
Dataset 2	830
Dataset 4	660
Dataset 3	635

Table 2.1 Results of adding values from Figure 2.5 linearly, presented as textual listings

Alternative cartographic presentations of search results are offered by the SPIRIT search engine (Purves et al., 2005). The system offers a 2D map with each document represented by an icon positioned at the location of the document's spatial footprint. Additional metaphors include a density-based cartogram, which represents the number of documents in the region by the size of the circles. Alternatively the numbers of documents can be presented in a 2D density surface. Although these approaches clearly improve on the traditional textual list, they however ignore the temporal relevance of documents. Temporal relevance is a useful constraint to use in GIR, as has already been observed by Göbel et al. (2002) and Beard and Sharma (1997).

A study that considers all three properties of geographic data (space, time and theme) is by Beard and Sharma (1997). They propose calculating the temporal relevance from a ratio of the period defined by the candidate dataset and the period defined by the query. This approach to calculation of temporal relevance is consistent with that by Yamuna and Candan (2000) discussed earlier in this chapter, however Beard and Sharma (1997) propose three separate formulae for handling different possible scenarios. In contrast, Yamuna and Candan (2000) propose a single similarity formula for handling different temporal scenarios. Similarly, Beard and Sharma (1997) propose an approach for calculating the spatial rank of candidate datasets based on the ratio of overlap between the spatial extent of datasets and a query. To address thematic relevance, they propose ranking several thematic fields separately and presenting a Boolean value to indicate whether a thematic field was matched or not.

Consistent with proposals by Göbel et al. (2002), their study proposes ranking each property individually and thereafter combining the rankings graphically into a three-part glyph, illustrated in Figure 2.6. Each part is colour-coded according to the calculated degree of relevance for that property, such that the darker the shade of colour, the higher the degree of relevance. However, related studies warn that the Red, Green, Blue (RGB) colour scheme of modern computer graphics "offers 256^3 or 16 777 216 possible colour combinations, far more than the eye can distinguish" (Robinson et al., 1995: pp. 356). To overcome this, Beard and Sharma (1997) calibrate the glyphs into four parts starting from zero (i.e. 0, 0.25, 0.5, 0.75 and 1). We highlight that this significantly reduces the resolution of each scale. However, we support their argument that their approach "avoids a meaningless mathematical combination of different attributes" (Beard and Sharma, 1997: pp. 158) and highlight that it is consistent with our discussion of the one-dimensionality of textual ranked lists presented earlier in this section.

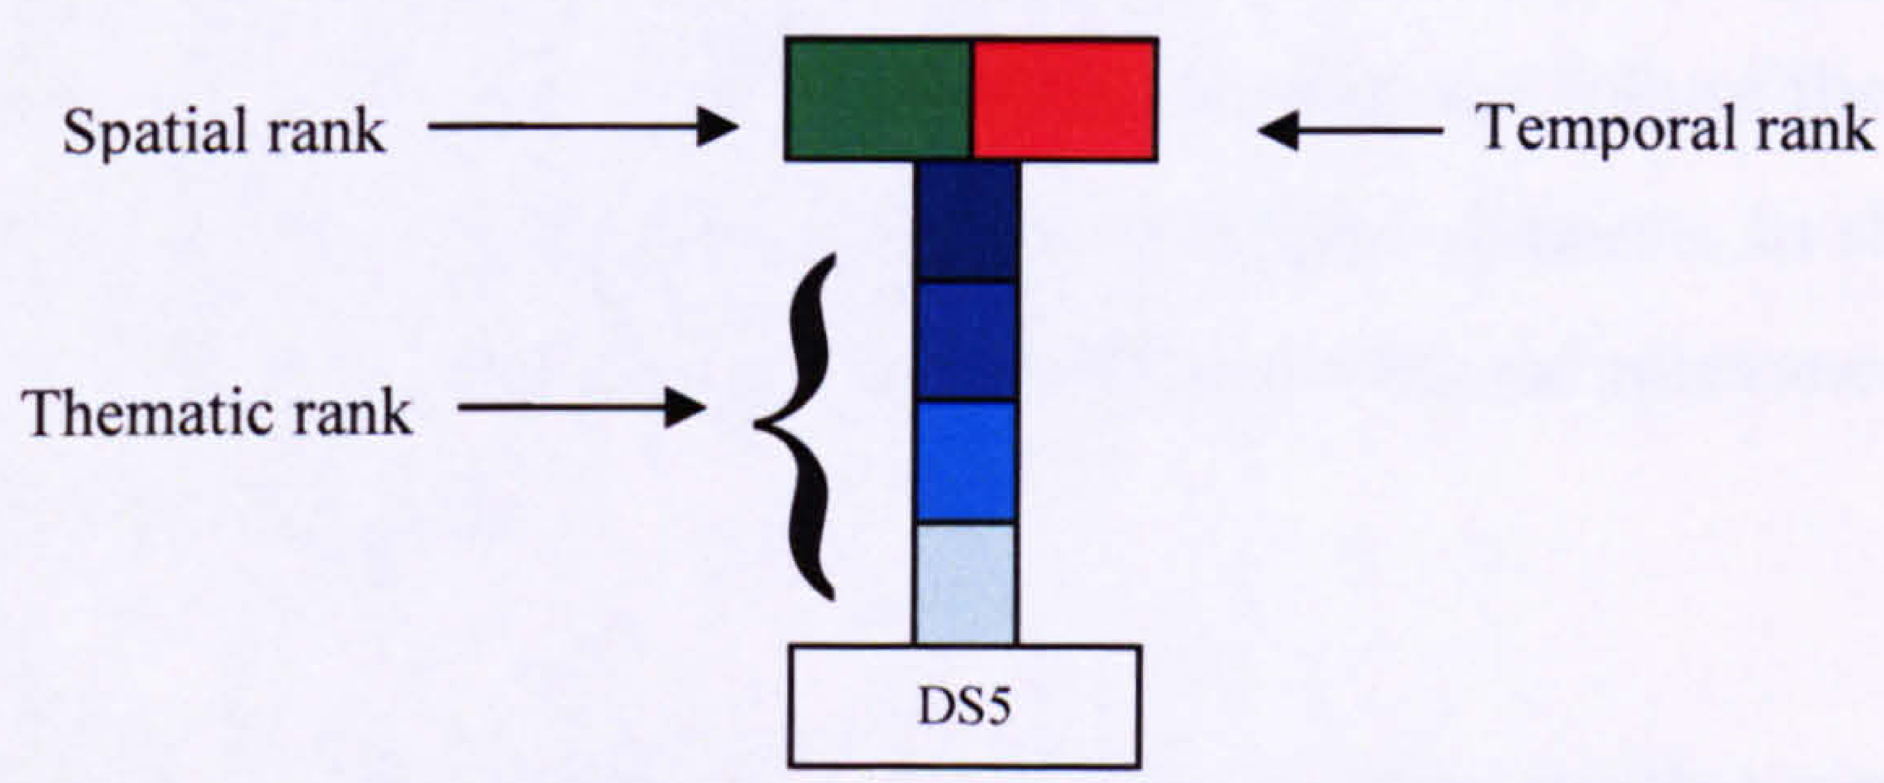

Figure 2.6 Multidimensional ranking glyph proposed by Beard and Sharma (1997)

An alternative study by Rainio (2005) considers the visualisation of geographic metadata through a scatterplot matrix and parallel coordinate plots. A parallel coordinate plot arranges the axes of a multidimensional phenomenon into parallel bars and thus allowing several axes to be viewed concurrently. A scatterplot matrix presents all pairs of axes from a multidimensional set of matrixes in smaller scatterplots. From an experiment comparing the suitability of scatterplots to parallel coordinate plots for purposes of data discovery, Rainio (2005) observed that users found scatterplots more complex than parallel coordinate plots, even though they could detect the best datasets through using scatterplots. The users also spent more time viewing textual metadata than the parallel coordinate plots. The study concludes that although these forms of multivariate visualisations provide useful insight into the characteristics of geographic (geospatial) metadata, they however "cannot support users of geospatial metadata in the comparison of spatial elements of geographic elements" (Rainio, 2005: pp. 92). We disagree with the generalisation of this conclusion, as comparison of spatial similarity (indicated through overlapping or intersection geometries) has been demonstrated in studies by Beard and Sharma (1997), and Larson and Frontiera (2004). The study does not consider important concepts of Information Retrieval such as similarity, relevance and ranking.

It is necessary to enable a user to view and evaluate the 'fitness for use' of a candidate dataset. Raper et al. (2002) propose a framework for evaluating geographic information, through the visualisation of the datasets and associated metadata. They acknowledge the importance of allowing the user to assess a candidate dataset before acquisition, by examining a visualisation and the associated metadata. Through the implementation of a prototype, called PanoraMap, they demonstrate the feasibility of implementing a web-based visualisation application for assessing geographic datasets. Unfortunately, their study does not consider the role of the data discovery system in assisting the user with evaluating candidate datasets. In short, their study does not consider the ranking of datasets and how inferred relevance is communicated to the user.

This subsection has highlighted a major limitation of presenting ranked geographic search results as textual lists; specifically, the one-dimensionality of ranked textual lists was discussed. An alternative 2D cartographic approach was discussed as well; however, the 2D approach ignores temporal properties of geographic data. A 3D visualisation approach was also discussed; however, the approach relies on the ability of users to distinguish between different shades of colour which has been argued in related studies to be a difficult cognitive task. It should be noted that there has been increased activity within the web mapping community with the introduction of web maps on mainstream portals such as Google. This could be an indication of a future uptake of 2D and possibly 3D visualisation within future GIR applications. A further and more detailed discussion of visualisation approaches is presented in the next chapter.

2.3 Indexing and Text-Operations

Also shown in Figure 2.1 is the role played by the Indexing mechanisms. The main benefit of creating an index is the significant improvement in retrieval time (that is, a shorter response time) (PostgreSQL Global Development Group, 2005).
Unlike traditional IR systems, a GIR system has the added task of handling geometric information. Consequently, a GIR system ideally offers both textual and spatial indexing facilities. An example of a text indexing method is the B-tree approach, which breaks up text into nodes, then creates records of keys and pointers to each node or groups of nodes. A spatial indexing structure organises spatial objects into sets of buckets, which normally correspond to pages of secondary memory (Güting 1994). An example of a spatial indexing approach is the R-tree —a B^+-tree like structure which indexes groups of points, lines and polygons based on which MBRs would require the least expansion to accommodate a new object (Longley et al., 2001, Beckmann et al., 1990). An optimised variant of this approach is the R*-tree, which minimises the MBR of each enclosing node. Of all the R-tree variants, the R*-tree offers the best performance (Papadias and Theodoridis, 1997). Other spatial indexing methods include quad-trees, k-d trees and the Z-order index (Gaede and Günther, 1998). A detailed description of each of these index methods is beyond the scope of this thesis; however Gaede and Günther (1998) offer an extensive discussion of the approaches.

To improve the retrieval performance of an IR or GIR system, the index is built over normalised text (Baeza-Yates, 1992). Text normalisation transforms metadata using a number of pre-defined operations, for example removal of stop-words such as “the” or “and”, synonym mapping, stemming, standardisation of date formats and change of character case (from upper to lowercase or vice versa). A query to the IR system is also normalised, removing stop-words and specific characters, before an index is searched. Baeza-Yates (1992) cautions that text normalisation has the disadvantage that documents cannot be searched using common words as these may have been identified as stop-words by the system developer.

2.4 Ontology in GIR

2.4.1 Ontology-assisted Query Operations

Queries on a GIR system may be invoked in a number of ways, for example, by pointing at a map, typing in a question, pulling down a menu or clicking some buttons. The queries used by GIR systems are often expressed in controlled natural language to improve usability (Larson, 1996). It is easier for a user to remember the name of a place rather than the coordinates of a place. For this reason, geoportals have, historically, consisted of gazetteers that are relational tables consisting of tuples of *<PlaceName, SpatialFootprint>* (de Oliveira et al., 1998, Larson, 1996, Brown, 1999, Tait, 2005). As the spatial footprints of abstract phenomenon, such as administrative or political boundaries are always changing, some gazetteers are built from *<PlaceName, Centroid>* tuples instead. However, recent developments in ontology-assisted GIR have extended the ability of applications to use semantics that are more expressive than thesaurus and gazetteer databases (Abdelmoty et al., 2005).

Ontology is a branch of metaphysics that is concerned with the nature or essence of being or existence (Soanes and Hawkes, 2005). Although originally an area of study within philosophical and metaphysical research, the study of ontology has been adopted by the computational sciences as a result of the increasing recognition of the potential of semantically-aware computation. From a computing perspective, "an ontology is the manifestation of a shared understanding of a domain, that is agreed between a number of agents, and such agreement facilitates accurate and effective communications of meaning, which in turn leads to other benefits such as interoperability, reuse and sharing" (Agarwal, 2005: pp. 504). The reference to a 'shared understanding', in the definition proposed by Agarwal (2005), highlights the importance of an ontology being agreed upon by different stakeholders within a domain. The definition also highlights the 'communication of meaning' which refers to the exchange of semantic information between different agents.

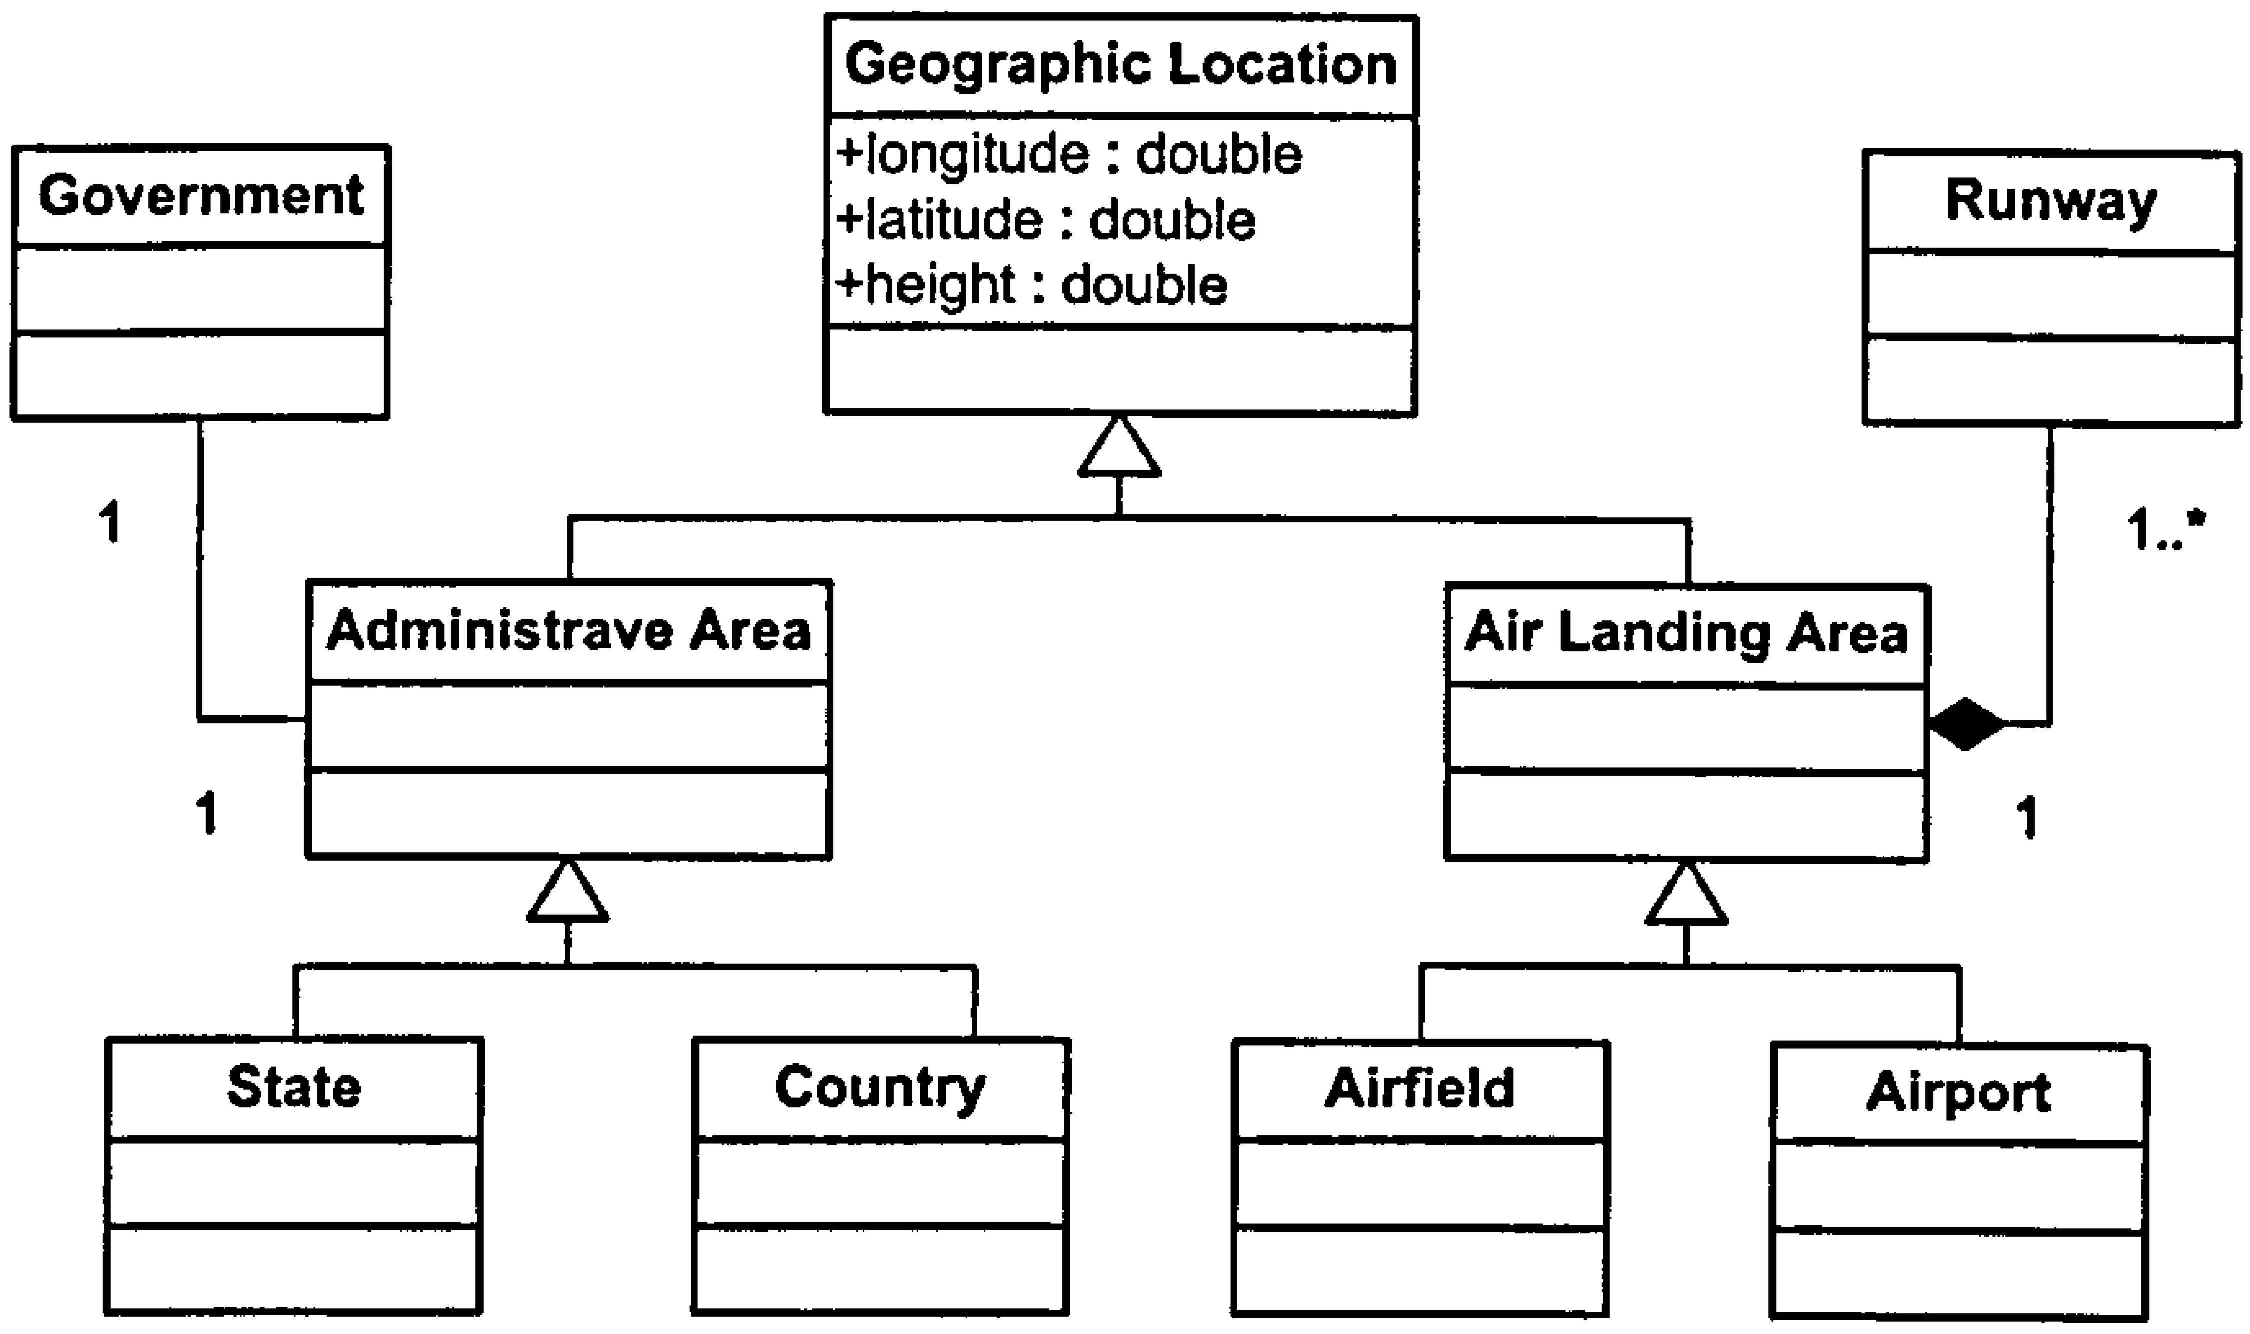

Figure 2.7 A UML illustration of a simple Ontology (adapted from DARPA, 2001)

Figure 2.7 shows a class hierarchy extracted from an ontology and illustrated in the Unified Modelling Language (UML). UML is a design notation used primarily for modelling object-oriented data models. According to the illustrated, an *Administrative Area* and an *Air Landing Area* are *Geographic Locations*, they therefore inherit properties from *Geographic Location* (namely Longitude, Latitude and height). *Each State* and each *Country*, by virtue of being an administrative area, is associated with exactly one *Government*. However, an *Air Landing Area* may be composed of one or more *Runways.* Within an ontological framework, each property is also defined such that relationships between classes can also be encoded, for example, a 'municipality governs a city'.

The development of ontologies has presented the possibility of semantically-aware information processing on the web through the Semantic Web —"a vision for the future of the World Wide Web in which information is given explicit meaning, making it easier for machines to automatically process and integrate information available on the Web" (World Wide Web Consortium, 2004). The geographic component of the semantic web has been the subject of several studies (Egenhofer, 2002, Raskin and Pan, 2005, Rodriguez and Egenhofer, 2004). In his work Egenhofer (2002) identifies four methods for presenting geographic meaning within the Semantic Web; these include natural language through basic hypertext, simple metadata encoded in structured human-readable markup, data models encoding triples of (entity, relationship, attribute) and logical semantics which provide relationships between terms and real-world entities. These methods of presenting geographic meaning in the Semantic Web form the basis of the 'Geospatial Semantic Web' (Kolas et al., 2005), referred to in some texts as the 'Semantic Geospatial Web' (Egenhofer, 2002).

2.4.2 Sharing Semantics

The World Wide Web Consortium (W3C) recently published the Web Ontology language (OWL) as a means of sharing semantics of entities and relationships between applications. Although it is now a W3C standard, OWL was preceded by the DARPA Agent Markup Language and Ontology Inference Layer (DAML+OIL). As the name suggests, DAML+OIL was originally a project of the United States Defense Advanced Research Projects Agency (DARPA). OWL has been developed encompassing lessons learnt from the design and application DAML + OIL. Figure 2.8 offers an illustration of an OWL fragment.

```
<rdfs:Class rdf:ID="GeographicLocation">
<rdfs:subClassOf rdf:resource="#Location"/>
<rdfs:comment>
                A Geographic Location is an item from the GEOFILE.
                It has a name and code.
</rdfs:comment>
</rdfs:Class>
<rdfs:Class rdf:ID="AirLandingArea">
<rdfs:subClassOf rdf:resource="#GeographicLocation"/>
</rdfs:Class>
<owl:Class rdf:ID="Airport">
<rdfs:subClassOf rdf:resource="#AirLandingArea"/>
<owl:equivalentClass>
<owl:Restriction>
 <owl:onProperty rdf:resource="#installationTypeCode"/>
 <owl:hasValue>
                APT
</owl:hasValue>
</owl:Restriction>
</owl:equivalentClass>
</owl:Class>
```

Figure 2.8 A subset of the GEOFILE Ontology encoded in OWL (DARPA, 2001).

The W3C defines three derivatives (also known as "sublanguages") of OWL that make up the complete OWL specification; OWL Lite, OWL DL and OWL Full. The derivatives implement the OWL specification in increasing degrees of expressiveness that is, *Full* > *DL* > *Lite*. A detailed discussion of all the derivatives is beyond the scope of this thesis; however, a full discussion is offered by Antoniou and van Harmelen (2004). A summary of the main characteristics of the derivates is as follows:

- OWL Lite: Is the least complex of the three derivatives. It is mainly suited for users requiring a classification hierarchy and simple restrictions. For example, it imposes a maximum cardinality of one for every property defined. This is in contrast to the other derivatives that allow higher cardinality values. One of the benefits of having a simple version, despite its restrictions, is that it allows for easy migration of thesauri and dictionaries.
- OWL DL: This sublanguage offers more expressiveness than OWL Lite. For example, it allows maximum cardinality. It includes all OWL constructs, however it is subject to some restrictions to ensure that all conclusions are computable and can be computed in finite time. An example of a restriction imposed on using this sublanguage is that a class cannot also be an individual or a property.
- OWL Full: Both OWL Full and OWL DL use the same complete vocabulary. However, OWL Full does not have the same restrictions imposed on OWL DL. This language is basically the most expressive of the derivatives.

Abdelmoty *et al.* (2005) present a critical comparison of OWL with the Geography Markup Language (GML). The latter language is developed by the Open Geospatial Consortium and is based on the object-oriented geographic data model (Open Geospatial Consortium, 2004). Their study concludes that OWL offers better mechanisms for representing relationships including relationships not included within the GML specification such as transitive, symmetric and inverse relationships. However, they acknowledge that GML offers a mature geographic data model that includes a large and rich vocabulary for encoding geographic concepts (Abdelmoty et al., 2005). This thesis supports the view that OWL offers more expressiveness of relationships than GML; however, we highlight that the GML specification allows attributes of any type to be included within a geographic object. This means that it is possible to include a property in a GML object that references an ontology class encoded in OWL. It should be noted that though OWL is the current standard, there are other ontology languages in existence, some of which are considered more expressive that OWL itself (Agarwal, 2005).

2.4.3 Ontology Implementations in GIScience

One of the earliest investigations into the use of ontologies in GIR is by Hübner et al. (2004). They present a system called GeoShare that allows users to search for geographic datasets and retrieve them from geographic web services. The retrieved datasets are then integrated into a 2D cartographic presentation. GeoShare computes the spatial, temporal and semantic relevance between a query and a document and integrates them into a single measure. The shortcomings of presenting single-measured relevance scores of geographic data have already been discussed earlier in this chapter. The linear function used by GeoShare is:

$$Dc(A,B) = \alpha\, Dh(A,B) + (1-\alpha)\, Dv(A,B)$$

where :

- A is a query constraint
- B is a candidate document property
- Dh is the spatial distance
- Dv is the hierarchical distance in the ontology
- Dc is the integrated degree of relevance

Another study involving the use of ontologies for the discovery of geographic web services is described in Klien et al. (2006). They present a scenario where data has been archived using the keywords "storm hazard model" and a user unknowingly searches for "estimate storm damage". They address this problem of heterogeneity through the implementation of a system that matches the outputs of geographic web services according to their semantics. They conclude that a method for constructing application ontology for service discovery should reflect the multitude of properties that may be used during service discovery (Klien et al., 2006). A further discussion of their approaches is presented in Lutz and Klien (2006), where they propose a query language for invoking geographic searches. The query language allows users to select properties of specific ontology concepts using a variety of constraints.

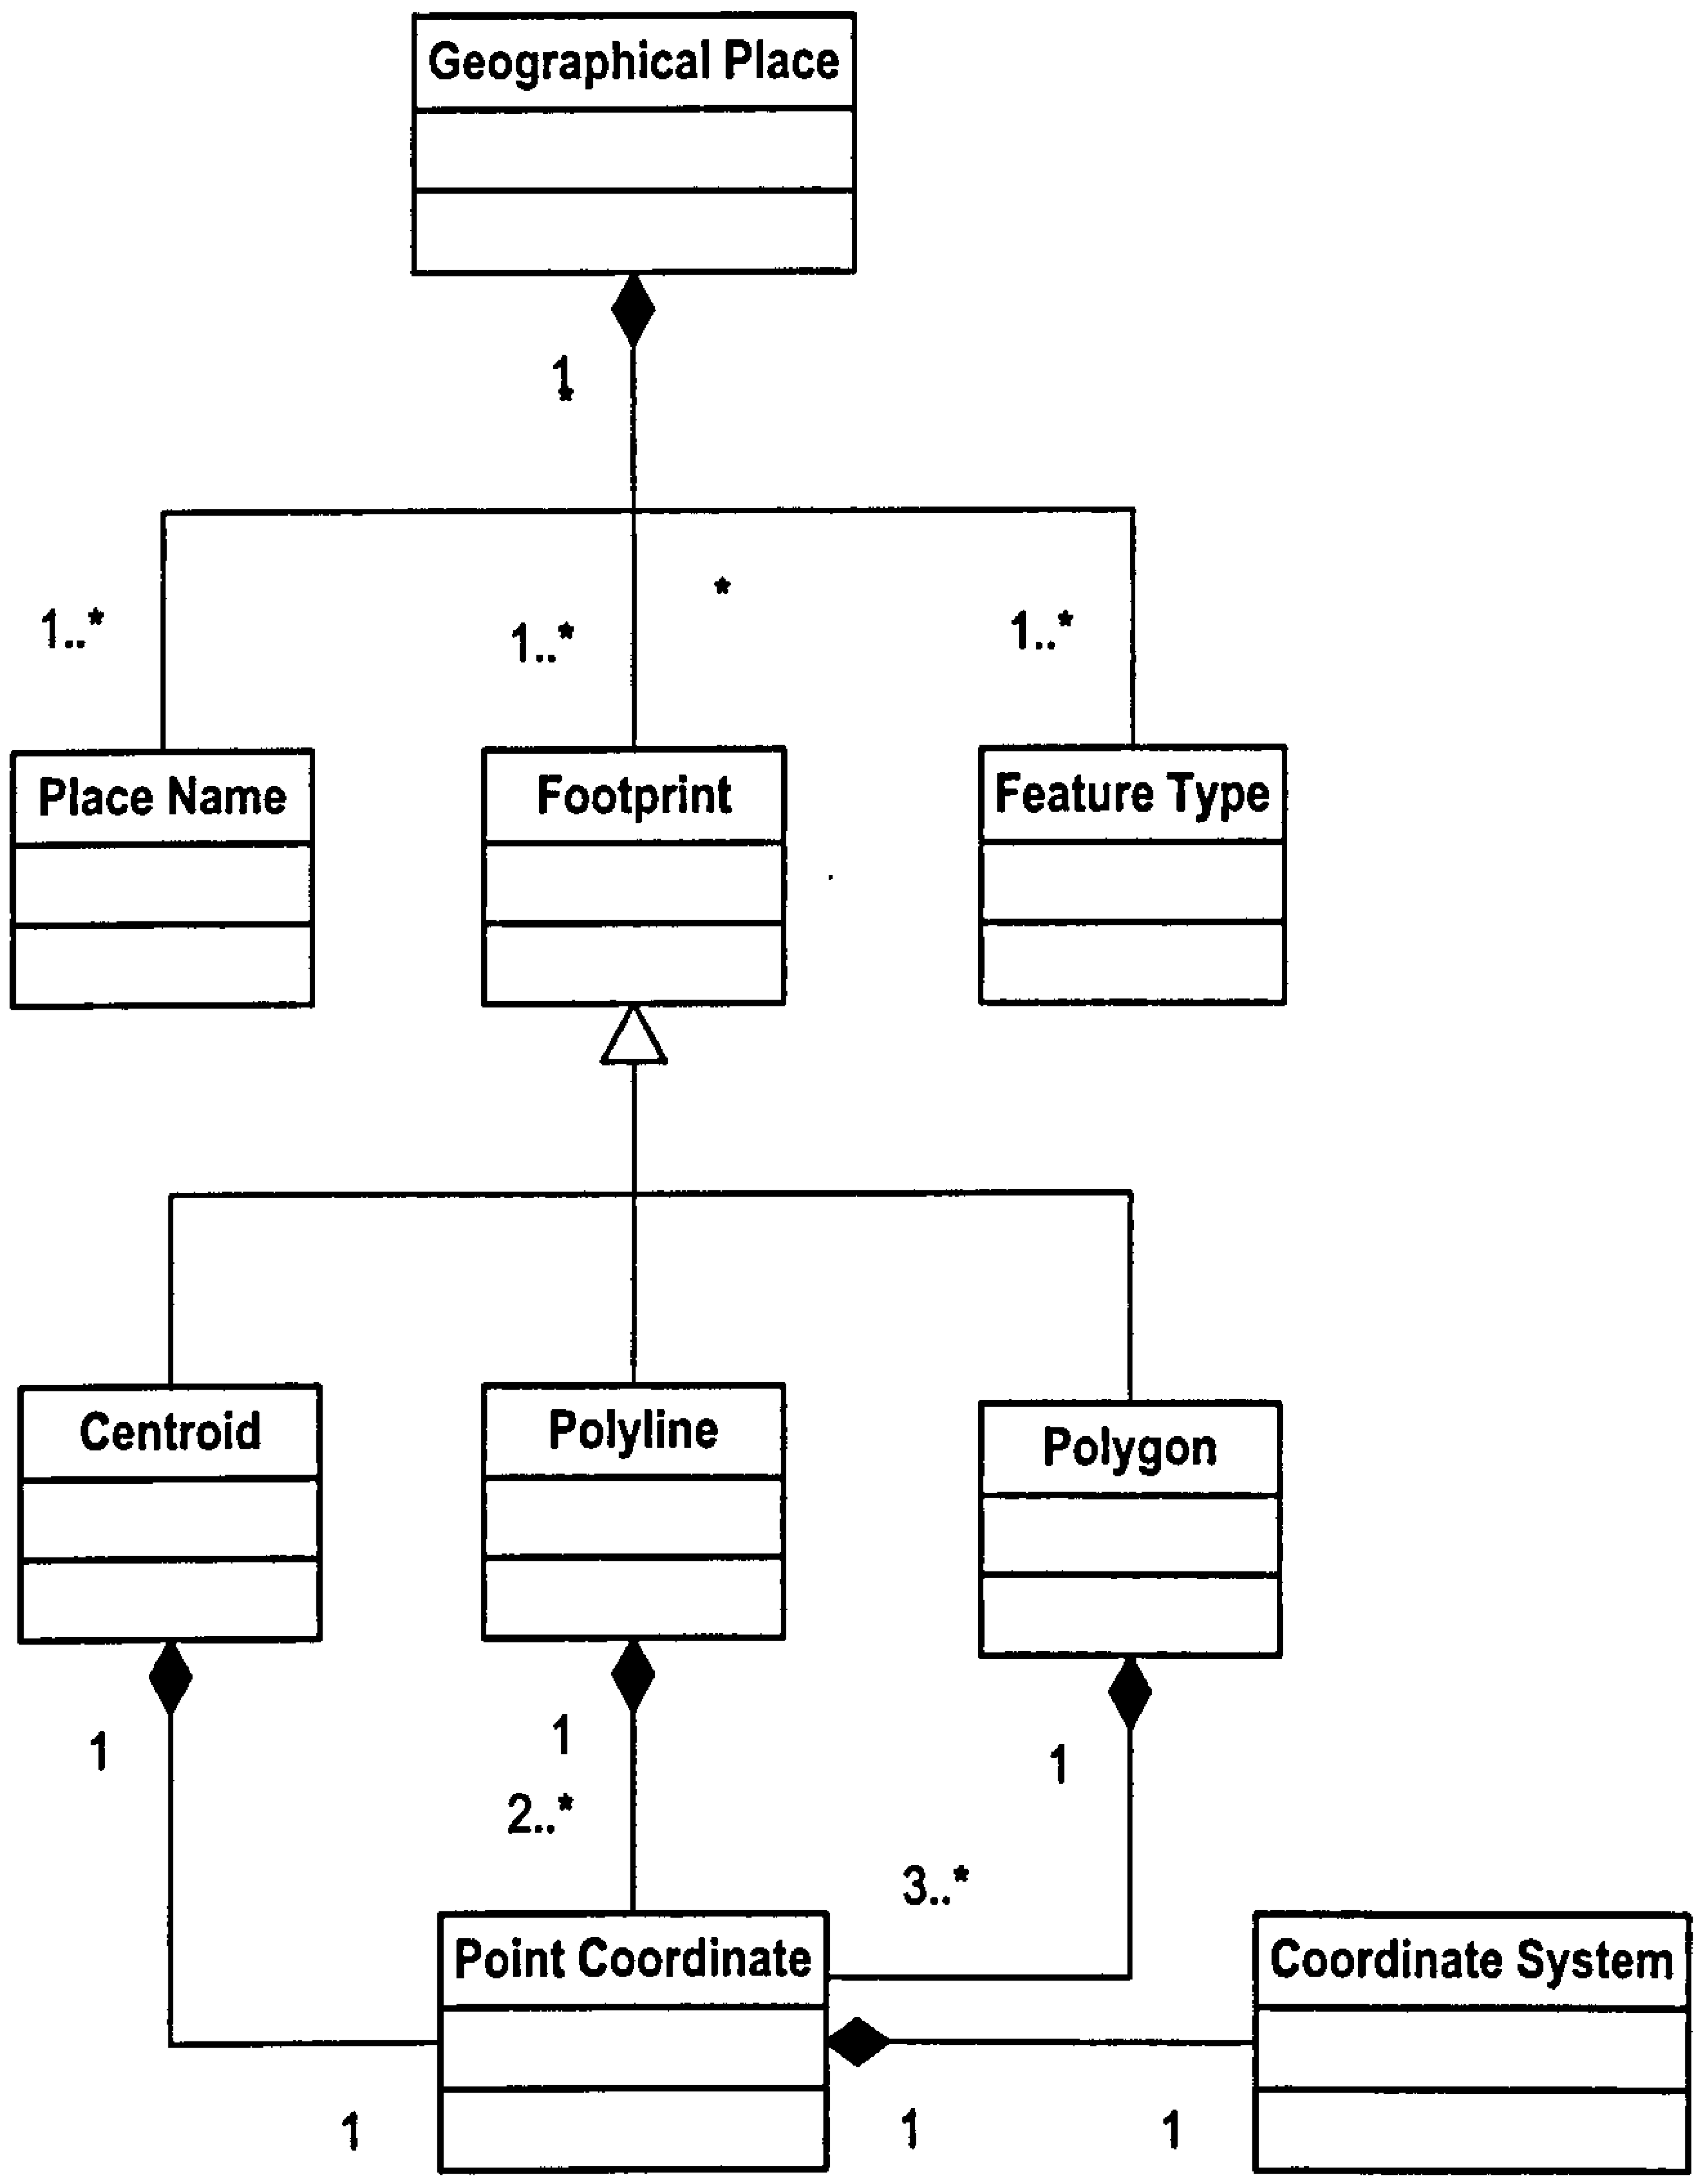

Figure 2.9 Conceptual design of the SPIRIT Geo-ontology (Jones et al., 2004).

An example of a study that integrates concepts of semantics and GIR is the European Union-funded project called SPIRIT (Spatially-Aware Information Retrieval on the Internet) (Jones et al., 2002). The geographical ontology used within SPIRIT abstracts both the geographical and thematic structure of places for use within a GIR infrastructure, as illustrated in Figure 2.9. Unlike conventional geoportals, SPIRIT is unique in that it ranks unstructured content, specifically web pages. This differs to other geoportals based on relational database engines and formal metadata as these hold structured data (formal data types and specified attribute fields). In contrast, content within web pages can be of any alphanumeric form and thus poses an additional challenge for IR. Using a domain ontology based on tourism and a geographic ontology for gazetteer purposes, the SPIRIT search engine is able to rank documents according to their textual and spatial properties (Jones et al., 2004). The role of ontologies within the SPIRIT system is illustrated in Figure 2.10.

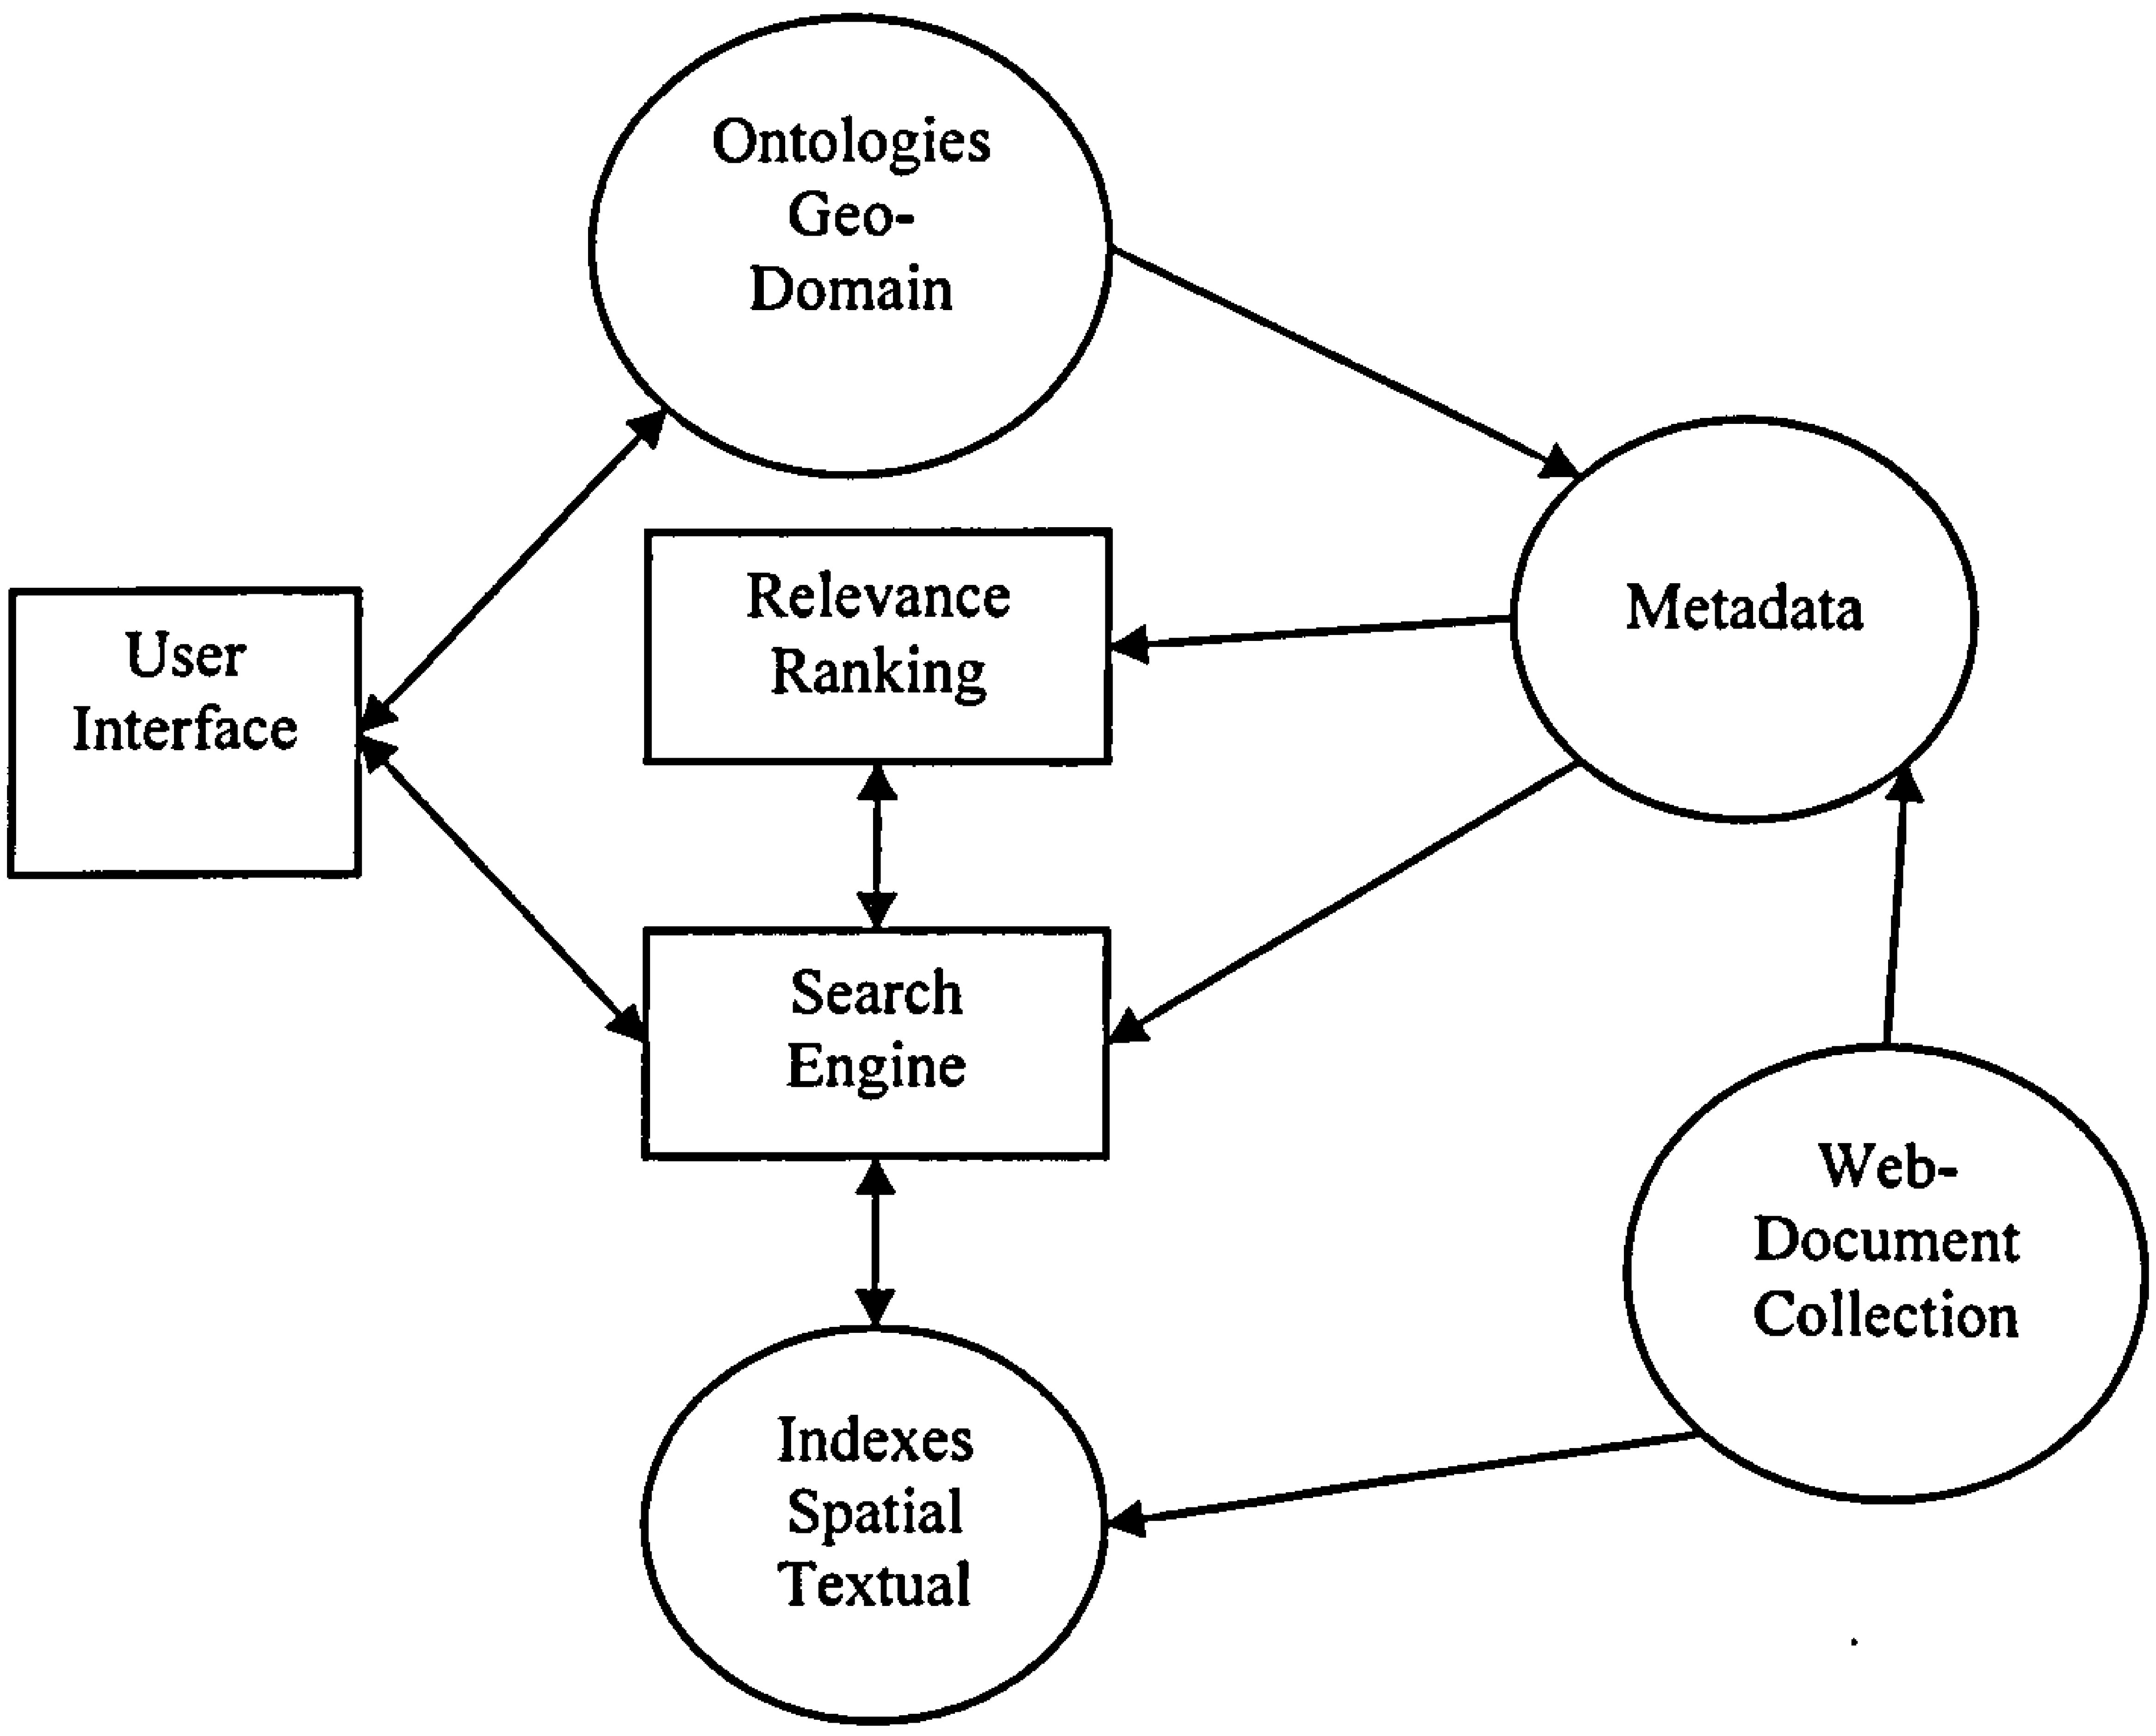

Figure 2.10 The Architecture of the SPIRIT search engine (adapted from Jones et al., 2004)

As part of our study, an application was developed that expands a query with similar concepts extracted from an ontology. The application is called the Spatio-Temporal Ontological Relevance Model (STORM) and is discussed in more detail later in this thesis. Given the wide range of ontologies available, two main ontologies were considered for supporting STORM, namely WordNet (Miller, 1995) and Cyc (CycCorp, 2003). WordNet is a lexical database offering approximately 100 000 terms including nouns, verbs, adjectives and adverbs; statistics of which are illustrated in Table 2.2. A key concept of the WordNet ontology is the concept of a synset — a set of terms with similar meaning (i.e. synonyms). Semantic relations implemented in WordNet include hyponyms and hypernyms, known as 'ISA' relationships because they indicate that a class is a subclass of another. Other semantic relationships include meronyms and holonyms, known as 'HASA' relationships because they indicate that a concept is a component of another.

As WordNet defines terms in common and scientific use, it is often regarded as a lexical ontology or a linguistic ontology. Agarwal (2005) highlights that linguistic ontologies are sometimes used for translating between philosophical and engineering concepts by relating concepts to natural language. This suggests that a linguistic ontology such as WordNet is ideal for geographic data discovery as queries in GIR are often expressed in natural language (Larson, 1996). Further, WordNet includes geographical categories (Kavouras et al., 2005). Lastly, due to its relational model, it can be stored in both an OWL document and a relational database engine, the latter being significantly faster and more efficient when handling millions of records. The popularity of WordNet for IR applications, is also recognised by Mihalcea (2003), whose study proposes a series of transformations for improving the effectiveness of WordNet for IR.

Part of speech	**Word forms**	**No. synsets**	**Total senses**
Noun	107929	74487	132408
Verb	10805	12753	23256
Adjectives	21364	18522	31078
adverbs	4582	3611	5722

Table 2.2 Statistics on WordNet (Mihalcea, 2003)

Another example of a popular ontology is Cyc, which offers a knowledge base (KB) and an inference engine. "The full version of the KB contains over 2.5 million assertions (facts and rules), interrelating more than 155,000 concepts" Siegel et al (2004: pp. 3). The assertions are derived from expert knowledge in the domains of chemistry, biology, defence, diseases, grammar, lexicons and several others. Using a predicate-based language, called CycL, it is possible to add to, modify, delete from or query the KB. An API is provided for integrating Cyc with third-party clients. An additional feature, called the Semantic Knowledge Source Integration facility (SKSI), allows Cyc to read semantic descriptions from a relational database. Until July 14th, 2006, only a subset of the entire Cyc ontology was available for free.

Since then the entire KB has been made freely available and is expected to encourage further ontology-based research.

Comparisons between these and other ontologies have been the subject of related studies by Lenat et al., (1995) and Kavouras et al.(2005). Their discussions suggest that Cyc offers more relationships per concept, whereas WordNet offers more concepts but with fewer relationships per concept. However, the successful use of WordNet in GIS is noted by Agarwal (2005) and its popularity is noted by Mihalcea (2003). Consequently, WordNet was selected for integration into the STORM prototype discussed later in this thesis. This subsection has presented some examples of studies involving the use ontologies in geographic or IR applications. The studies discussed in this subsection were selected because they discuss ontology-supported discovery specifically; other studies discussing the use of ontologies in geographic information science can be found at Albertoni et al. (2005), Arpinar et al. (2005) and Raskin and Pan (2005).

2.4.4 Calculating Semantic Similarity Using WordNet

Several studies have proposed approaches for calculating semantic similarity using Wordnet. Varelas et al.(2005) classify the approaches into three groups: edge counting, information content and feature-based approaches. Edge counting methods calculate semantic similarity based on the length of the path mapping one concept to another (Richardson and Smeaton, 1995). The shortest path from one concept to another represents the highest similarity. Resnik (1999) observed that a limitation of an edge-based approach is that it assumes equivalent distances between nodes in the path. However, uniform distances between nodes are difficult to define as some parts of the taxonomy may have a higher density of concepts than others. This irregularity in densities of links between concepts can lead to "unexpected conceptual distance measures" (Richardson and Smeaton, 1995: pp. 6). Feature-based methods calculate semantic similarity based on the properties of terms or on their relationships to other similar terms in the taxonomy. An example of a feature-based approach, the Tversky (1977) similarity model, was discussed earlier in this chapter. As Richardson and Smeaton (1995: pp. 4) note, "Tversky's feature based similarity model is arguably the most powerful similarity model to date". However; they also note that the existence of relatively few semantic relations in Wordnet, limits its effectiveness for IR purposes.

Information content-based methods calculate the difference in information content of two concepts based on the probability of occurrence of the concept in a corpus (Resnik, 1999). Mathematically, the information content is the negative of the log of the likelihood of occurrence. This suggests that a similarity measure based on an information content-model can adapt well to different applications, provided the application is adequately represented in a corpus. Within a GIS-domain for example, the terms 'geographic' and 'spatial' would likely have high probabilities of occurrence in a geospatial metadata collection, as they are often used to describe geospatial datasets. However, Richardson and Smeaton (1995) observed that the information content model only considers synonym and ISA relations, thus ignoring other semantic relations such as HASA. Another limitation of this approach is that it considers only the existence of a lexical equivalent on a term, ignoring the semantic definition of a term; for example, the probability of occurrence of the term bank would be influenced by references to 'river bank' and 'commercial bank'.

2.5 Chapter Summary

A geo-centric modification of the process of IR was presented in Figure 2.1 and the sections within this chapter discuss each stage of that process. Some of the issues discussed included indexing, ranking, query operations and the visual interface. An examination of the literature suggests that GIR-based research is an integration of concepts from IR research and Geographic Information Science. Considerations in IR tend to be directly applicable to geographic data, as geographic data includes attributive properties that may not be of a spatial nature. However, as has been illustrated from the referenced texts, IR ignores spatial properties and hence geo-centric methodologies are necessary for efficient GIR.

The concept of relevance in IR was also examined. Of particular importance is the high-dimensionality of relevance and how it leads to variations in the perception of relevance amongst different users. A view characterising relevance according to system and human-based relevance, proposed by Borlund (2003a), is discussed. We contend that this view implies that evaluation of an IR or GIR system should reflect whether relevance is being assessed by an application or the human. The discussion on relevance presented in the chapter leads to an examination of ranking approaches used within the field of IR and GIR. Different relevance computation approaches are presented based on similarity, proximity and hyper-textual references. Most of the ranking models presented result in normalised probabilistic values offering a standardised interface for implementation in any GIR system.

The role of ontology and semantics within query operations is also examined. Within a GIR infrastructure ontologies promise to extend the services originally offered by gazetteers and thesauri. A description of the current standard used for encoding ontology, OWL, is also presented. Related research comparing OWL with GML (a standard for encoding geographic data) is also presented. Also noted is the concept of the Semantic Web and its geographic component, the Geospatial Semantic Web. Approaches for presenting meaning within a Geospatial Semantic Web are presented. This thesis emphasizes that these methods of presenting semantics, coupled with Ontology encoding standards, offer the possibility of multiple GIR systems sharing the same definitions of semantics.

The chapter contended that ranked lists are one-dimensional in nature. An illustration is offered showing how three series are merged into a single value when presented in a textual list. Related studies that support the view that ranked lists are one-dimensional were presented. One of the presented studies proposes a 2D cartographic approach for presenting results on a GIR system. The approach, implemented on the SPIRIT search engine, ignores temporal properties of the candidate datasets. An alternative approach, discussed in this chapter, proposes a multidimensional visual metaphor for presenting degrees of relevance through a multi-part glyph. Each part of the glyph is shaded according to a four-step scale representing the degrees of relevance. Although the approach improves on the textual ranked list, a scale of four shades reduces the resolution of a similarity-based relevance score (a fraction between zero and one). However, a four-step scale ensures that each degree of relevance is distinguishable from the next on the scale. The next chapter considers the visualisation and computer graphics with regard to geographic data discovery.

Chapter 3 Visualisation Support for Geographic Data Discovery

The appearance and presentation of geographic information has been the subject of several studies in geographic visualisation (geovisualisation) — defined as the integration of approaches from scientific visualisation, cartography, image analysis, information visualisation, exploratory data analysis and GIS to provide theory, methods and tools for visual exploration, analysis, synthesis and presentation of geographic data (MacEachren and Kraak, 2001). Concepts developed in geovisualisation research are important for geographic data discovery as they assist in communicating descriptions of datasets and associated metadata to users. This chapter starts by tracing the conceptual origins of geovisualisation with respect to data and information and goes on to discuss the visualisation of geographic data and the results of a geographic search.

3.1 An Introduction to Visualisation

The maxim "a picture is worth a thousand words" represents traditional thinking regarding the benefits of visualisation. The fundamental objective of visualisation is to allow humans to get insight into data and draw conclusions through the detection of relationships between the portrayed objects. Visualisation as a means to data exploration is particularly useful for extracting information from large amounts of potentially useful datasets. However, it is widely suggested within information systems research that the availability of data does not necessarily imply the existence of information, knowledge or wisdom. It is the role of computer assisted systems such as visualisation applications to attempt to make information apparent to a human. This raises the question, "when does data become information?" Longley et al., (2001) suggest an answer to this question in their portrayal of decision-making support infrastructure(Figure 3.1):

- Data: The quantities or characters operated on by a machine such as a computer. For example, recordings of wind direction, rainfall and temperature measurements in the Kalahari Desert may be considered as data.
- Information: It is data that has been selected, organised and prepared according to a given purpose or a posed question. For example, the rainfall measurements aforementioned offer an answer to the question "is the Kalahari Desert a dry terrain?". In this case the weather conditions of the desert (dryness in particular) are the context within which the question is posed.
- Evidence: A conflation of validated (and sometimes contradictory) information from different sources related to specific problems. For example, rainfall measurements from the desert for a period of a week could suggest that the desert is wet territory. However, a series of rainfall datasets over longer periods would offer statistical evidence to the contrary.
- Knowledge: It is information to which value has been added by interpretation based on context, skill, experience, education or purpose. For example, it is widespread knowledge that deserts are generally dry territories.
- Wisdom: It is the most difficult to define as it is highly individualised and depends on the context within which a decision is made based on all the available evidence and knowledge (Longley et al., 2001).

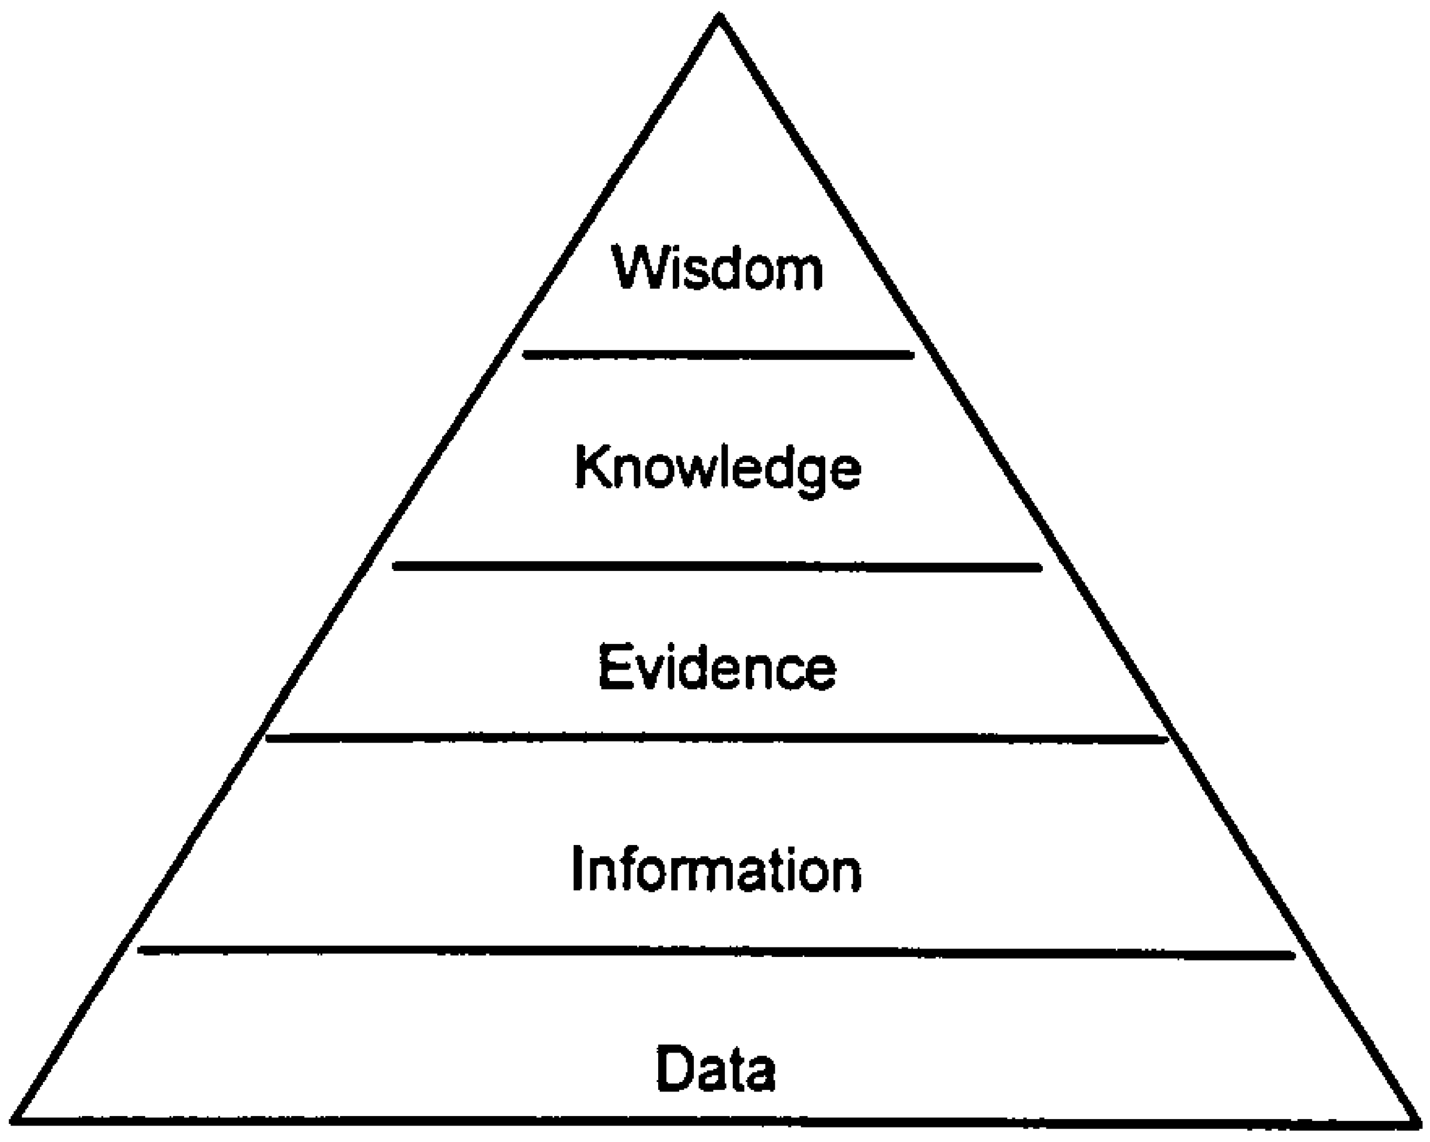

Figure 3.1 Decision-making support infrastructure (Longley et al., 2001)

3.1.1 Analysis through Visualisation

Having distinguished between data and information we next introduce information visualisation. Information Visualisation is that aspect of visualisation that deals with the representation of abstract and non-spatial concepts and relationships that may not have physical counterparts in the real world, for example, the concept of relevance discussed in Chapter 2 (de Oliveira and Levkowit, 2003, Rohrer and Swing, 1997). Scientific visualisation, sometimes referred to as Visualisation in Scientific Computing, involves the representation of measured or simulated data representing objects or concepts associated with phenomenon from the physical world (de Oliveira and Levkowit, 2003). Both these categories of visualisation create graphical representations from data and facilitate the detection of patterns and relationships between objects. This thesis adopts the view that geographic visualisations span the definitions of both scientific and information visualisation as the attributive component of a geographic datum can hold varying forms of data or information.

An important concept in Information Visualisation is spatialization, which is defined as "data transformation method based on spatial metaphors, with the aim of generating a cognitively adequate graphic representation (e.g. a depiction that matches human's internal visualisation capabilities) for data exploration and knowledge discovery" (Fabrikant and Skupin, 2005: pp. 668). To enable a user to effectively interpret an information visualisation, Fabrikant and Skupin (2005) identify three main challenges as being important for spatialization, i) encoding the semantics of data into appropriate spatial metaphors, ii) employing appropriate spatial structures for visually depicting the semantics of the data, iii) controlling the potentially experiential effects spatialized visualisations may have on users through, for example, controlled navigation and browsing. They propose the use of ontology to generalise the semantic descriptions, by identifying the essential characteristics of the source, and formalizing the appropriate source-target referencing rules. This form of semantic generalization is then followed by geometric generalization, which involves the reduction of the geometric detail of a representation. Historically, the reduction of the amount of semantic or geometric detail from spatial information has been an area of study within cartographic generalization (Robinson et al., 1995, Kraak and Ormeling, 2003).

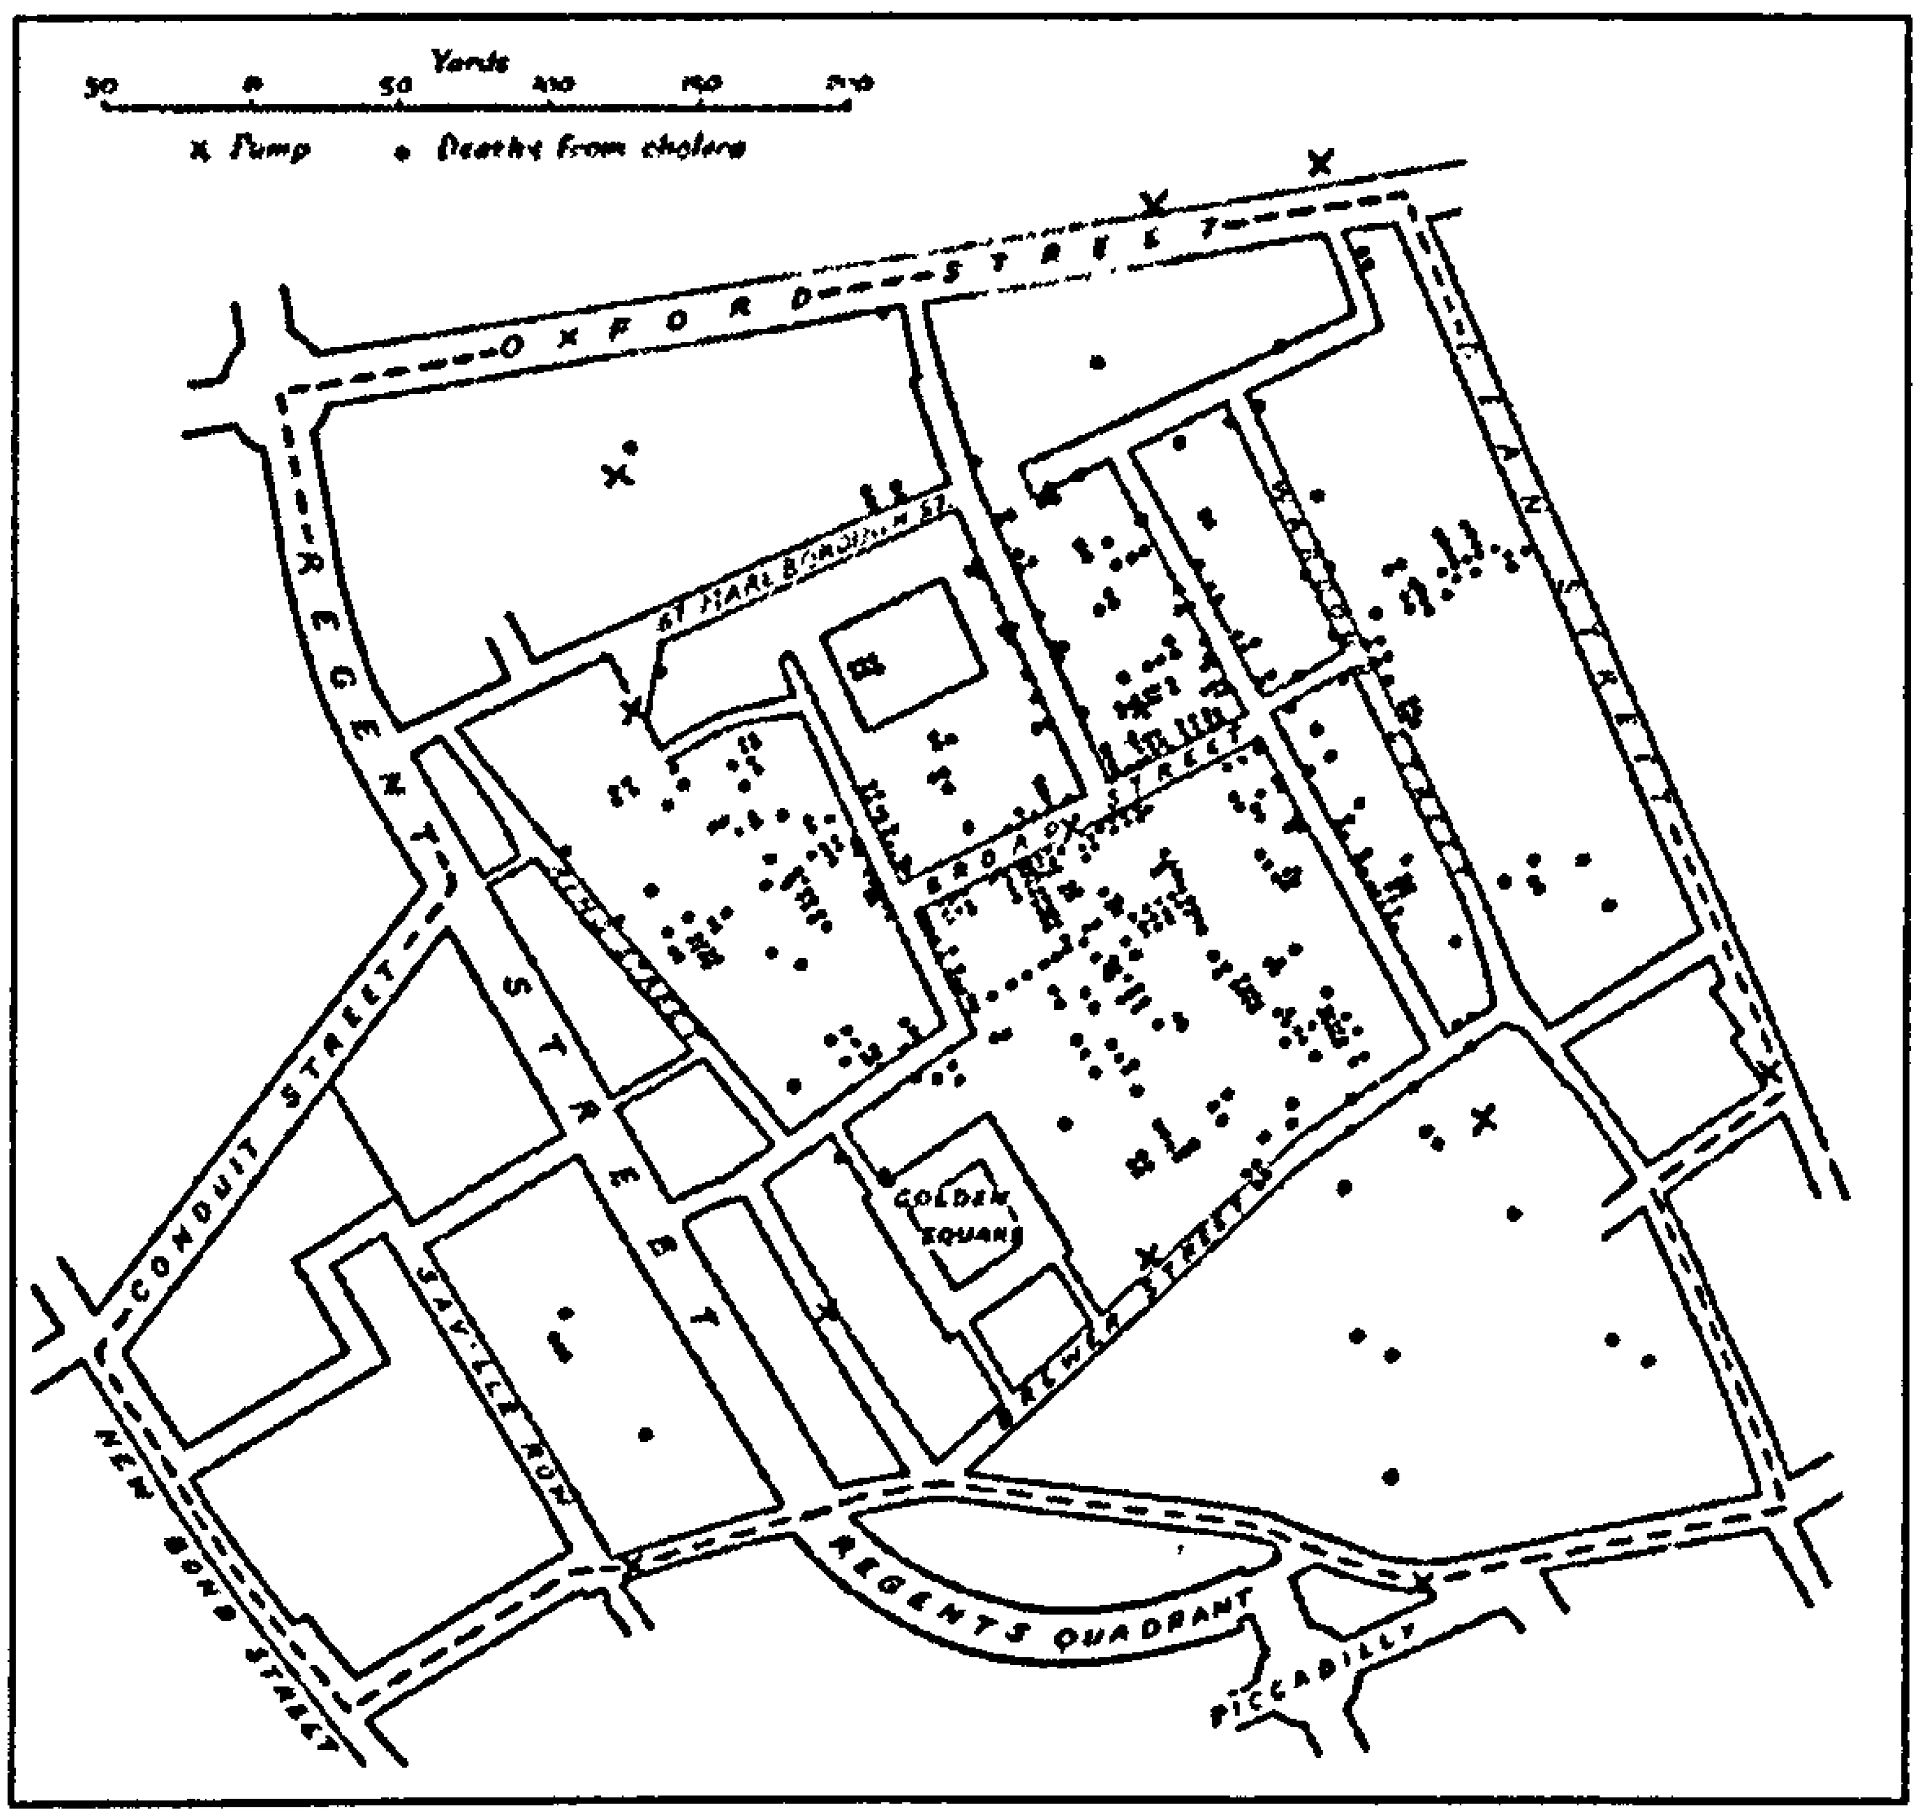

Figure 3.2 The original map of the 1850s cholera outbreak in London by Dr John Snow (Gilbert, 1958)

One of the earliest applications of visualisation within the geographic domain was in 1855 by Dr John Snow. In the second edition of his paper 'On the mode of communication of cholera' Dr Snow presented a map, shown in Figure 3.2, illustrating the locations of the first five hundred deaths from cholera in the Broad Street area of London in September 1854 (Gilbert, 1958, Longley et al., 2001). The deaths had occurred over a period of ten days. The 'cholera field', as Snow referred to it, had its centre at a pump in Broad Street. This led to Snow formulating a hypothesis that "the incidence of cholera was only amongst people who drank from the Broad Street pump". At Snow's request, the handle of the pump at Broad Street was removed and the incidence of new cases ceased almost immediately. A modern day example of visualisation within spatial analysis includes a study, by Proctor et al. (2005), on the spatial autocorrelation between troop locations and health symptoms in the 1991 Gulf War. Their study employs visualisation techniques to arrive at results suggesting "significant local spatial clustering" in terms of where veterans were located during their deployment in relation to severe postwar health conditions.

Based on these examples of visualisation-supported spatial analysis and on the aforementioned definition of geovisualisation by MacEachren and Kraak (2001), we take the view that geovisualisation is not only about how the world looks but also about how the world works. As geographical information, by definition, always includes the dimensions of location and attribute, the rest of this chapter therefore discusses 2D and multidimensional geovisualisation.

3.1.2 Comparative Studies between 2D and 3D

Before discussing geovisualisation later in this chapter, it is important to note that the usefulness of multidimensional visualisations is an active area of research and controversy. Nielsen (2000) offers a criticism of three-dimensional (3D) visualisation centred on a difficulty to navigate a virtual world. We appreciate that the views were published in 2000 before the current generation of 3D applications, such as Google Earth, that have been acquired by the wider public (Butler, 2006). The high download statistics for all the aforementioned Digital Earth applications suggest that the public is becoming more accustomed to 3D navigation, possibly due to the increasing popularity of videogames. The role of videogames in improving users' navigation skills is suggested by results from a study by Sjölinder et al (2005), whose study showed that older users needed more time to solve tasks in 3D virtual environments as compared to the younger users — who make up most of the videogaming community. Second, we contend that navigation cannot be considered an appropriate indicator of the usability of all 3D applications as it varies from application to application. Variability of navigation controls is not unique to 3D visualisation applications but is also evident in two-dimensional (2D) visualisation applications as well, for example between ESRI ArcMap and Autodesk Map. Instead we suggest that the ability of users to discern accurate and useful information from a visualisation is a more appropriate indicator.

In a study by Ware and Frank (1996) improvements of a factor of three in support of stereo-enabled 3D visualisations over their 2D alternatives were observed. These observations are supported by a quantitative comparison of 2D and 3D visualisation by Carvajal (2005) who observed an increase in understanding when users were exposed to 3D models as compared to traditional 2D drawings. In another study, Cockburn and McKenzie (2004) observed that there was no significant difference between the effectiveness of spatial memory when using 2D and 3D computer-supported systems; However, effectiveness decreased when using static 3D physical models (made from wires, papers cards and strings). They expressed reservations over the role of static-perspective views in 3D. This thesis supports their view that 3D physical models are limited in their usability, however we highlight that real-time computer graphics are dynamic and not subject to the same limitations. This is evidenced by an evaluation study by Risden et al., (2000) in which they observed that users performed search tasks faster, without jeopardising quality of response, on dynamic 3D visualisations in comparison to 2D visualisations. Based on the arguments posed by these earlier studies we conclude that 3D visualisation has a significant role to play in Information and Scientific visualisation by making apparent relationships between objects; however, its role is alongside its 2D counterpart.

3.2 Geovisualisation

Being able to convey information using different metaphors of visualisation is important for data discovery as users may be accustomed to viewing variable metaphors depending on their context, skill, experience, education or purpose. These parameters have already been noted as being important factors of knowledge in Section 3.1. We illustrate the importance of variable metaphors with the following example involving civil engineers and hydrologists: civil engineers often use surface models when designing new roads or construction sites. Hydrologists often use contour maps to determine river catchments and other topographic features. This thesis contends that given some geographic data, it is appropriate to allow each user to be able to view the data using a metaphor they are accustomed to. Hence given a grid of spot heights, a civil engineer may choose to visualise it as a surface model, shown in Figure 3.3a; and a hydrologist may visualise it using a contour plot, shown in Figure 3.3b.

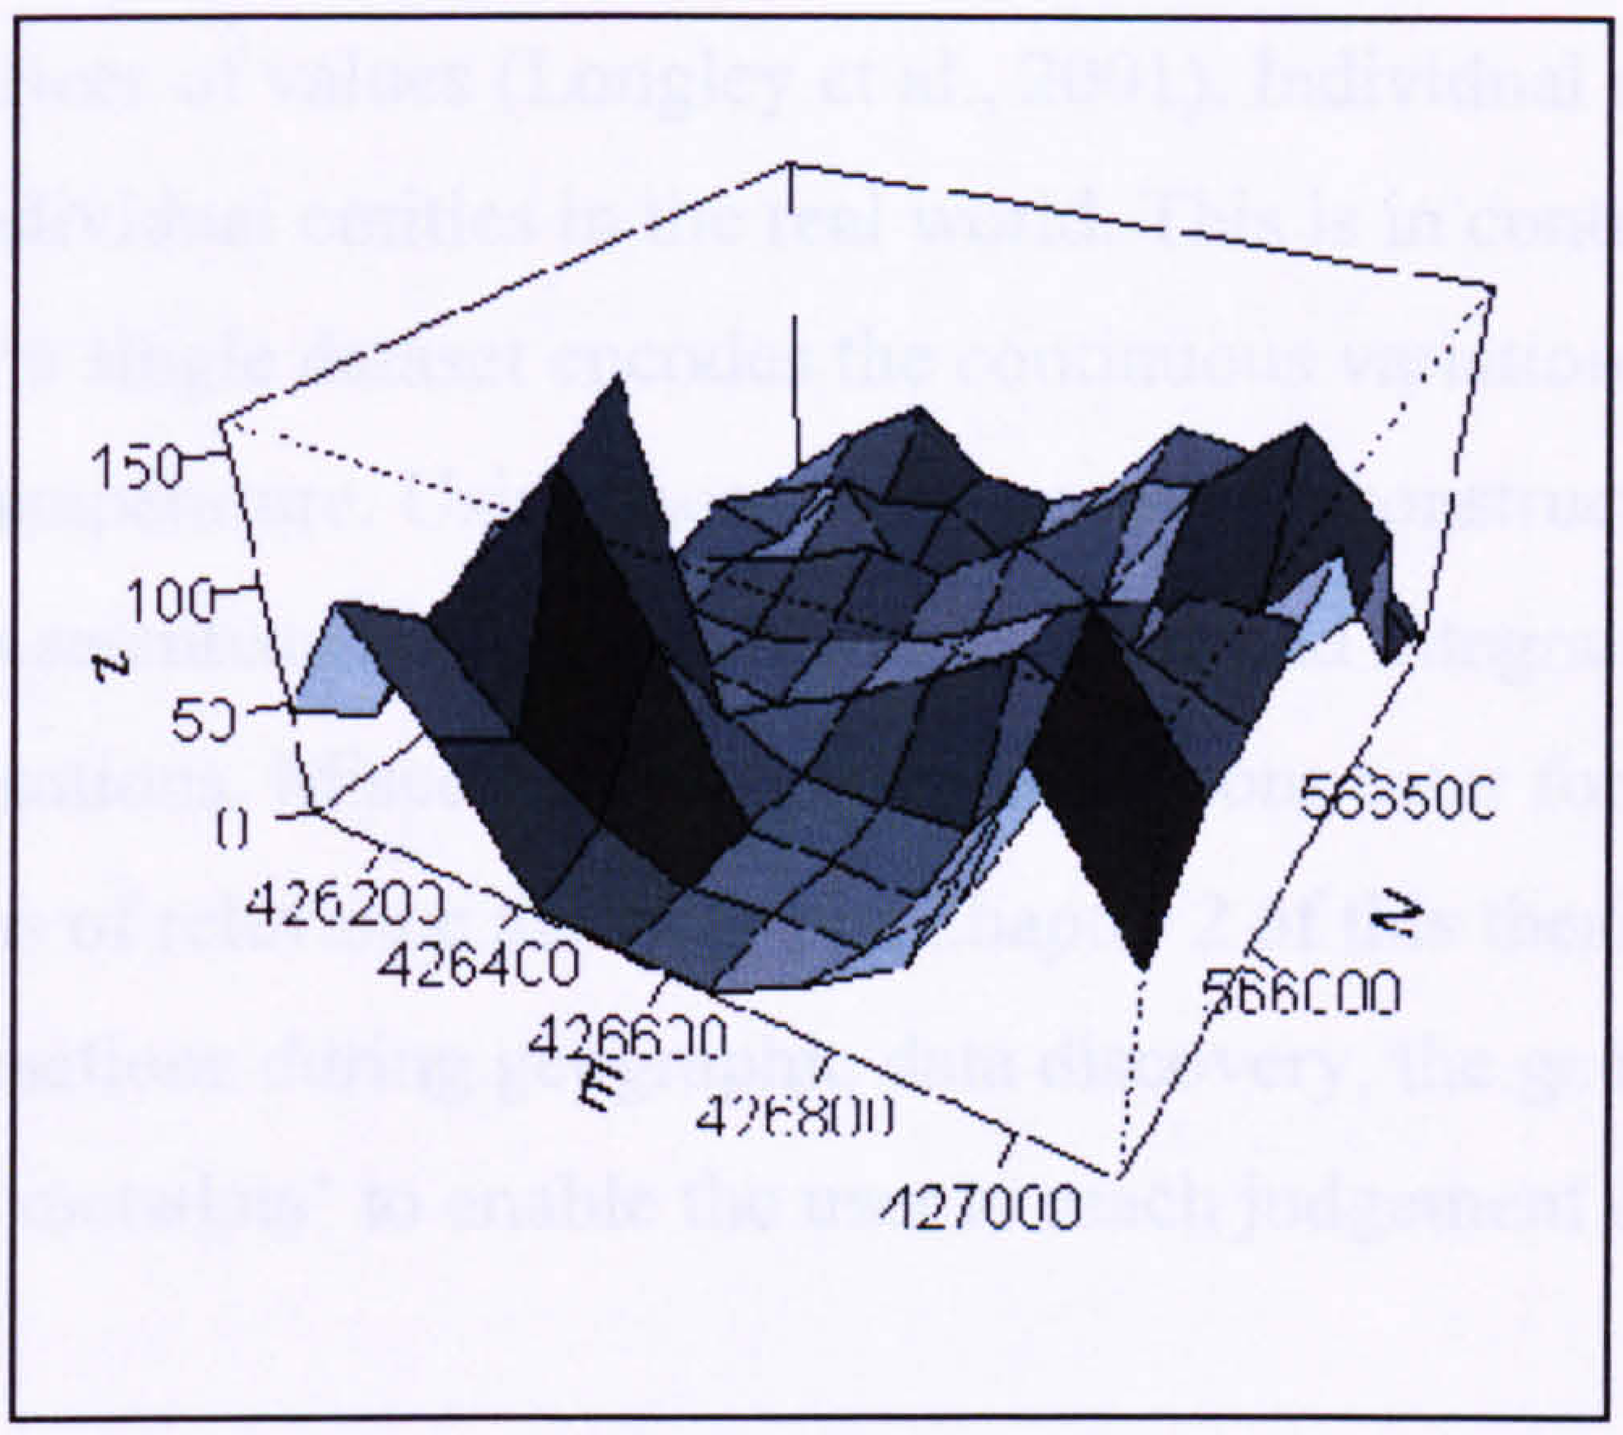

(a)

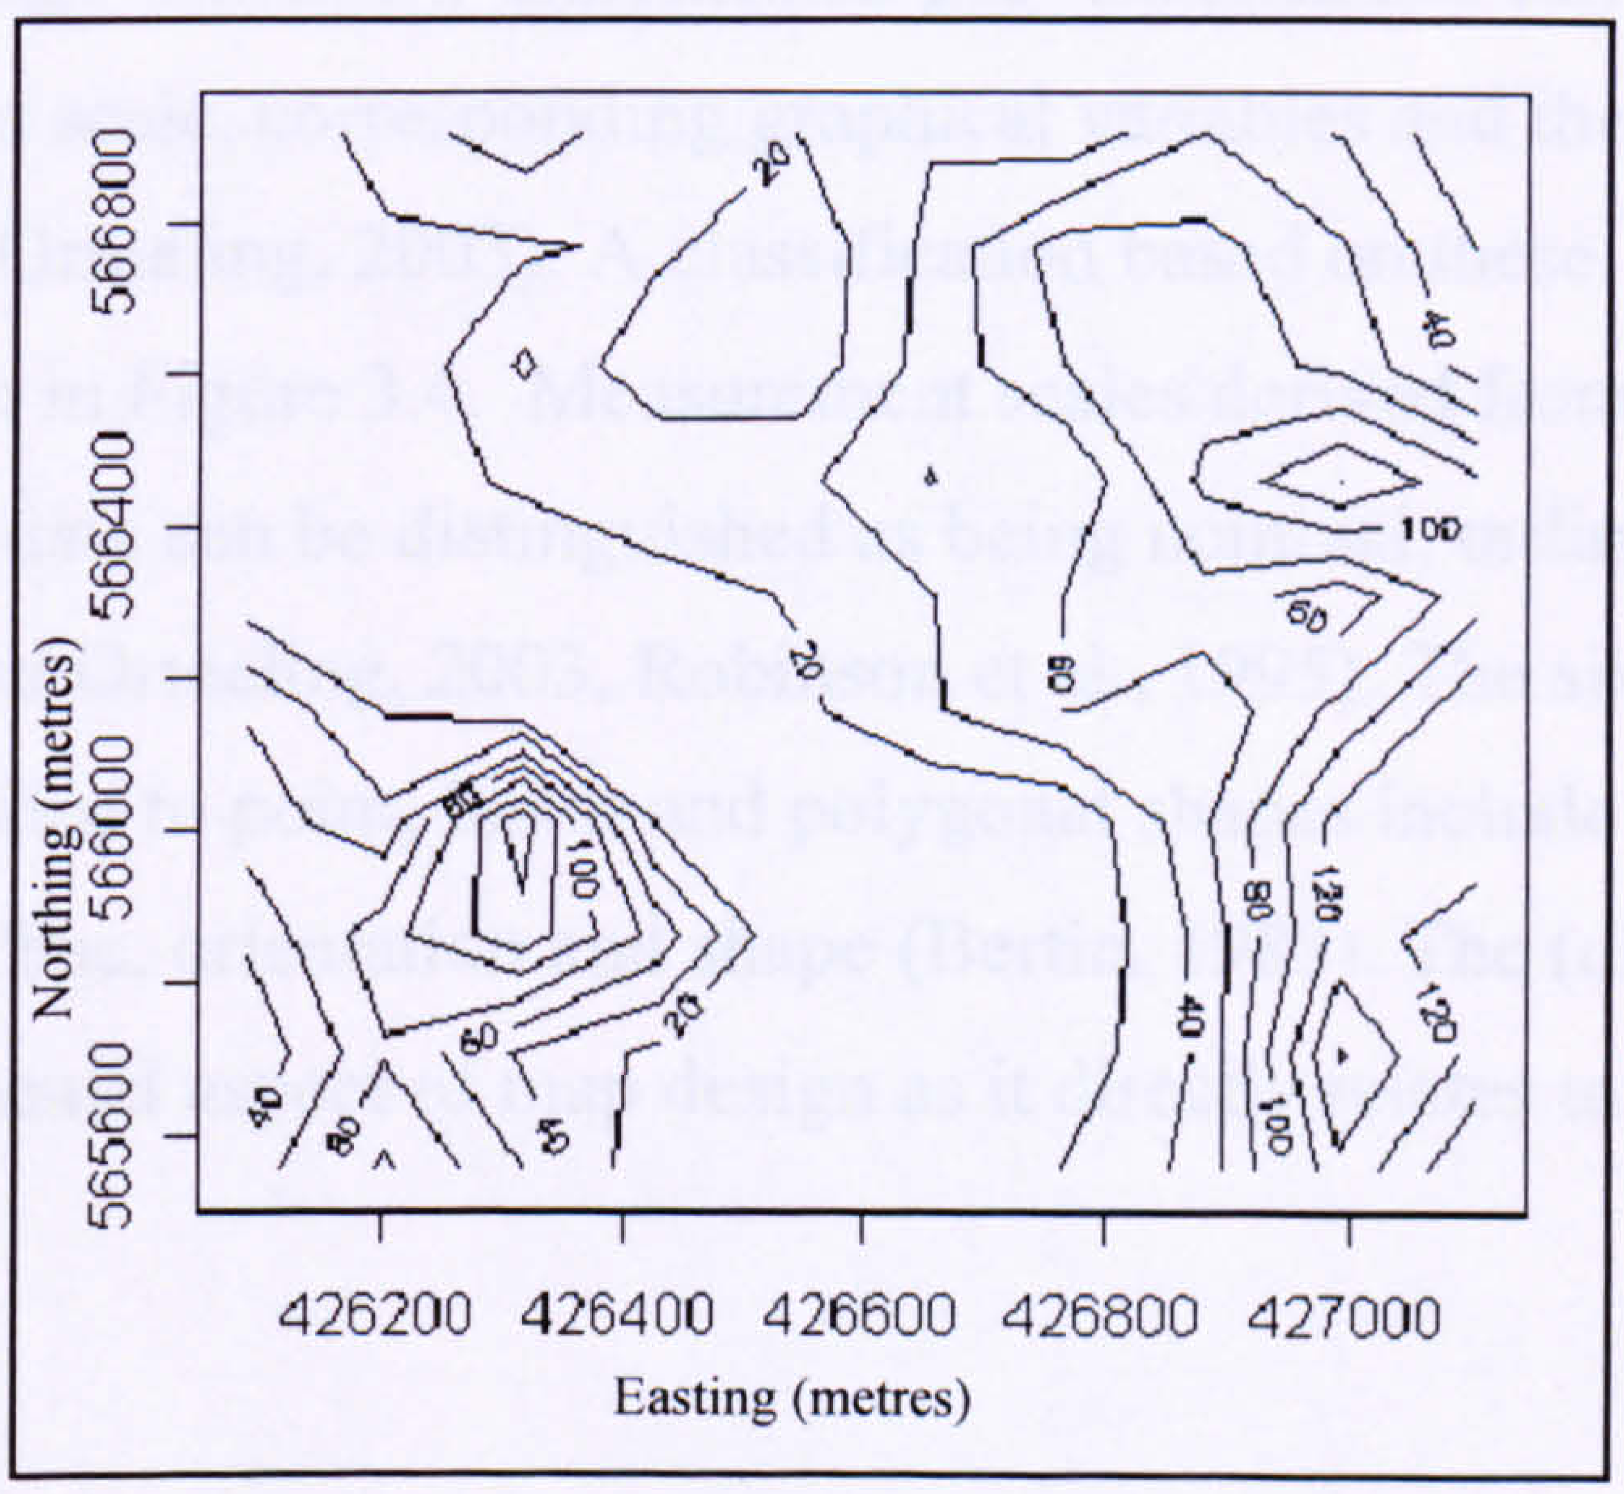

(b)

Figure 3.3 Different visualisations created from the same grid of spot heights: (a) a surface model (b) a contour plot

3.2.1 Two-Dimensional Geovisualisation

Geovisualisations are normally constructed from two forms of data namely, *vector* data encoded as points, lines, polygonal or volumetric objects and *raster* data encoded as regular matrices of values (Longley et al., 2001). Individual objects in vector data represent individual entities in the real world. This is in contrast to *raster* geographic data, where a single dataset encodes the continuous variation of a single phenomenon, such as temperature. Using these basic geometric constructs, geographic information scientists have been able to compile and integrate datasets into meaningful visualisations. Miscellaneous geovisualisations cater for the variations in perceptions of relevance discussed in Chapter 2 of this thesis. By presenting these visualisations during geographic data discovery, the geoportal is providing some 'visual metadata' to enable the user to reach judgement on the relevance of a dataset.

Some of the ways through which two-dimensional geovisualisations can be classified is through measurement scale, corresponding graphical variables and the continuity of the data (Kraak and Ormeling, 2003). A classification based on these characteristics is shown in Figure 3.4. Measurement scales derived from the values of the properties in the data can be distinguished as being nominal, ordinal, interval or ratio scale (Kraak and Ormeling, 2003, Robinson et al., 1995). The six basic graphical variables applied to point, linear and polygonal shapes include differences in size, colour, texture, hue, orientation and shape (Bertin, 1983). The (dis)continuity of data forms a fundamental aspect of map design as it directly relates to vector and raster data models.

	nominal	**ordinal/interval/ratio**		**composite**
graphic variables	**variation of hue, orientation, form**	**Repetition**	**variation of grain, size, grey value**	**variation of size, segmenta-tion**
point data	nominal point 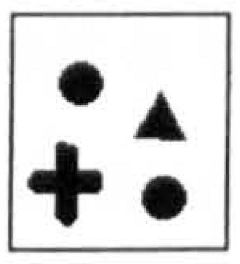	dot maps 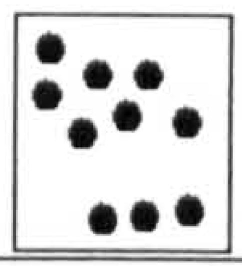	proportional maps 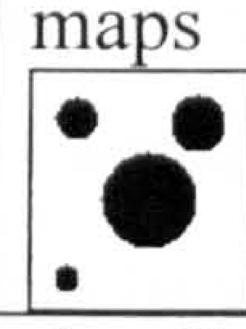	point diagram 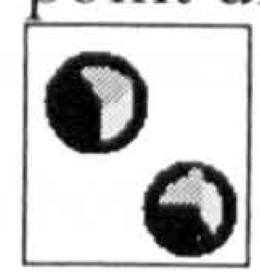
linear data a)lines	nominal line symbol maps 	—	flowline maps 	line diagram 
b)vectors	—	standard vector maps	graduated vector maps	vector diagram maps
polygonal data a)regular distribute-on	land use maps	regular grid symbol maps	proportional symbol grid maps and choropleth	areal diagram grid
b)irregular boundaries	chorochro-matic mosaic maps	—	choropleth	areal diagram
volume data	—	—	stepped statistical surface	—
surface data	—	isoline map	filled-in isoline map	—
volume data	—	—	smooth statistical surface	—

Figure 3.4 A classification of different mapping methods based on measurement scale, corresponding graphical variables and continuity of data (Kraak & Ormeling 2003: pp127). Grey cells show continuous data and white cells show discrete data.

Ordinal scales of measurement, as applied to cartography, distinguish by class and rank objects within a class according to a quantitative measure. These allow for visual comparisons between maps as they encode values of attributes by colour, shade or other visual variables (Andrienko and Andrienko, 1999). For example, dot maps where a high density of maps represents a concentration of a phenomenon.

Additional dimensionality can be added by varying two or more graphic variables such as hue with the sizes of points, lines or polygons. Within a geographic data discovery framework, quantitative representation of geographic information offers immediate visual data exploration as compared to presenting the user with relational tables of alphanumeric data.

Similarly, *interval* scales of measurement introduce a specified quantity between ranks. The interval is specified in terms of a standardised unit of measurement and remains equal between all consecutive ranks. For example, elevation above mean sea level can be scaled into 256 equal parts, allowing it to be presented on a Red, Green, Blue (RGB) 8-bit screen. *Ratio* scales are an extension of interval scales; however, ratios have a non-arbitrary absolute zero. Ratios are, therefore, independent of the unit of measurement. For example, an object that is 4 metres long remains twice as long as that which is 2 metres long, even if the length is measured in inches. Other examples of ratio scales include weight and population.

In contrast, nominal scales distinguish objects on a map according to qualitative considerations. These, therefore, do not imply any quantity of a represented geographic phenomenon. For example, a chorochromatic map presents differences between classes of phenomenon, such as land-cover. No order is assumed between each class and hence they are regarded as presenting qualitative information. Another example of a qualitative map is a tourist map that simply identifies names and locations of places of interest. Within a geographic data discovery framework, nominal maps can offer information about the semantics of the map. An integration of two geographic datasets and their presentation on a nominal map could allow users to discover composite datasets that are much richer in information than their aggregate subsets.

This subsection has presented some of the different approaches used in 2D geovisualisation. It has highlighted the challenge of matching a method of representation to the needs and experience of a user. It should be noted that not all presentations of geographic information are map-based, for example, bar charts of rainfall measurements sampled at different locations present geographic information but are not maps. Whereas, this subsection discussed 2D representation, the next subsection discusses 3D and higher dimensional representation of geographic information.

3.2.2 Multidimensional Geovisualisation

One of the most common types of 3D geovisualisation is the Digital Terrain Model (DTM). A DTM is a digital three dimensional representation of a terrain surface and of selected 1D, 2D and 3D geographic objects that are related to this surface (Kraak and Ormeling, 2003). The geographic objects referred to in the definition may include manmade infrastructure such as buildings, roads, bridges and others. When the digital representation shows only altimetry then it is referred to as a Digital Elevation Model (DEM). "The DEM is the most useful representation of relief in GIS" (Longley et al., 2001: pp. 289). DTMs have been widely used in geology to illustrate subsurface constructs and their relationships, for example, underground mineral deposits in relation to settlements above surface. Other applications of DTMs have been in demonstrating the effects of extreme weather conditions such as hurricanes and floods. Figure 3.5 presents a DTM showing proximity of buildings to a river.

Figure 3.5 A DTM containing buildings, trees and a DEM

Just as there are different methods of 2D geovisualisation, 3D geovisualisations can also be significantly dissimilar. This dissimilarity in realism and content is sometimes due to the extra degrees of freedom and the significantly powerful graphics capabilities available in present day computer platforms. An examination of the geovisualisation literature suggests that web-served 3D terrain models towards the end of the millennium were often localised subsets of the whole earth, for example, terrain and buildings within the city limits of Newcastle upon Tyne. One of the major reasons for this limitation was that graphics rendering operations were computationally intensive and early hardware capabilities had not advanced enough to handle the computations effectively and efficiently. However, improvements in the affordability and availability of high-performance computer graphics technology have resulted in the creation of dynamic and seamless DTMs. A metaphor that has resulted from the availability of improved computer graphics is the Digital Earth concept, which presents geographic information on a representation of the earth centred in 3D Cartesian space as illustrated in Figure 3.6.

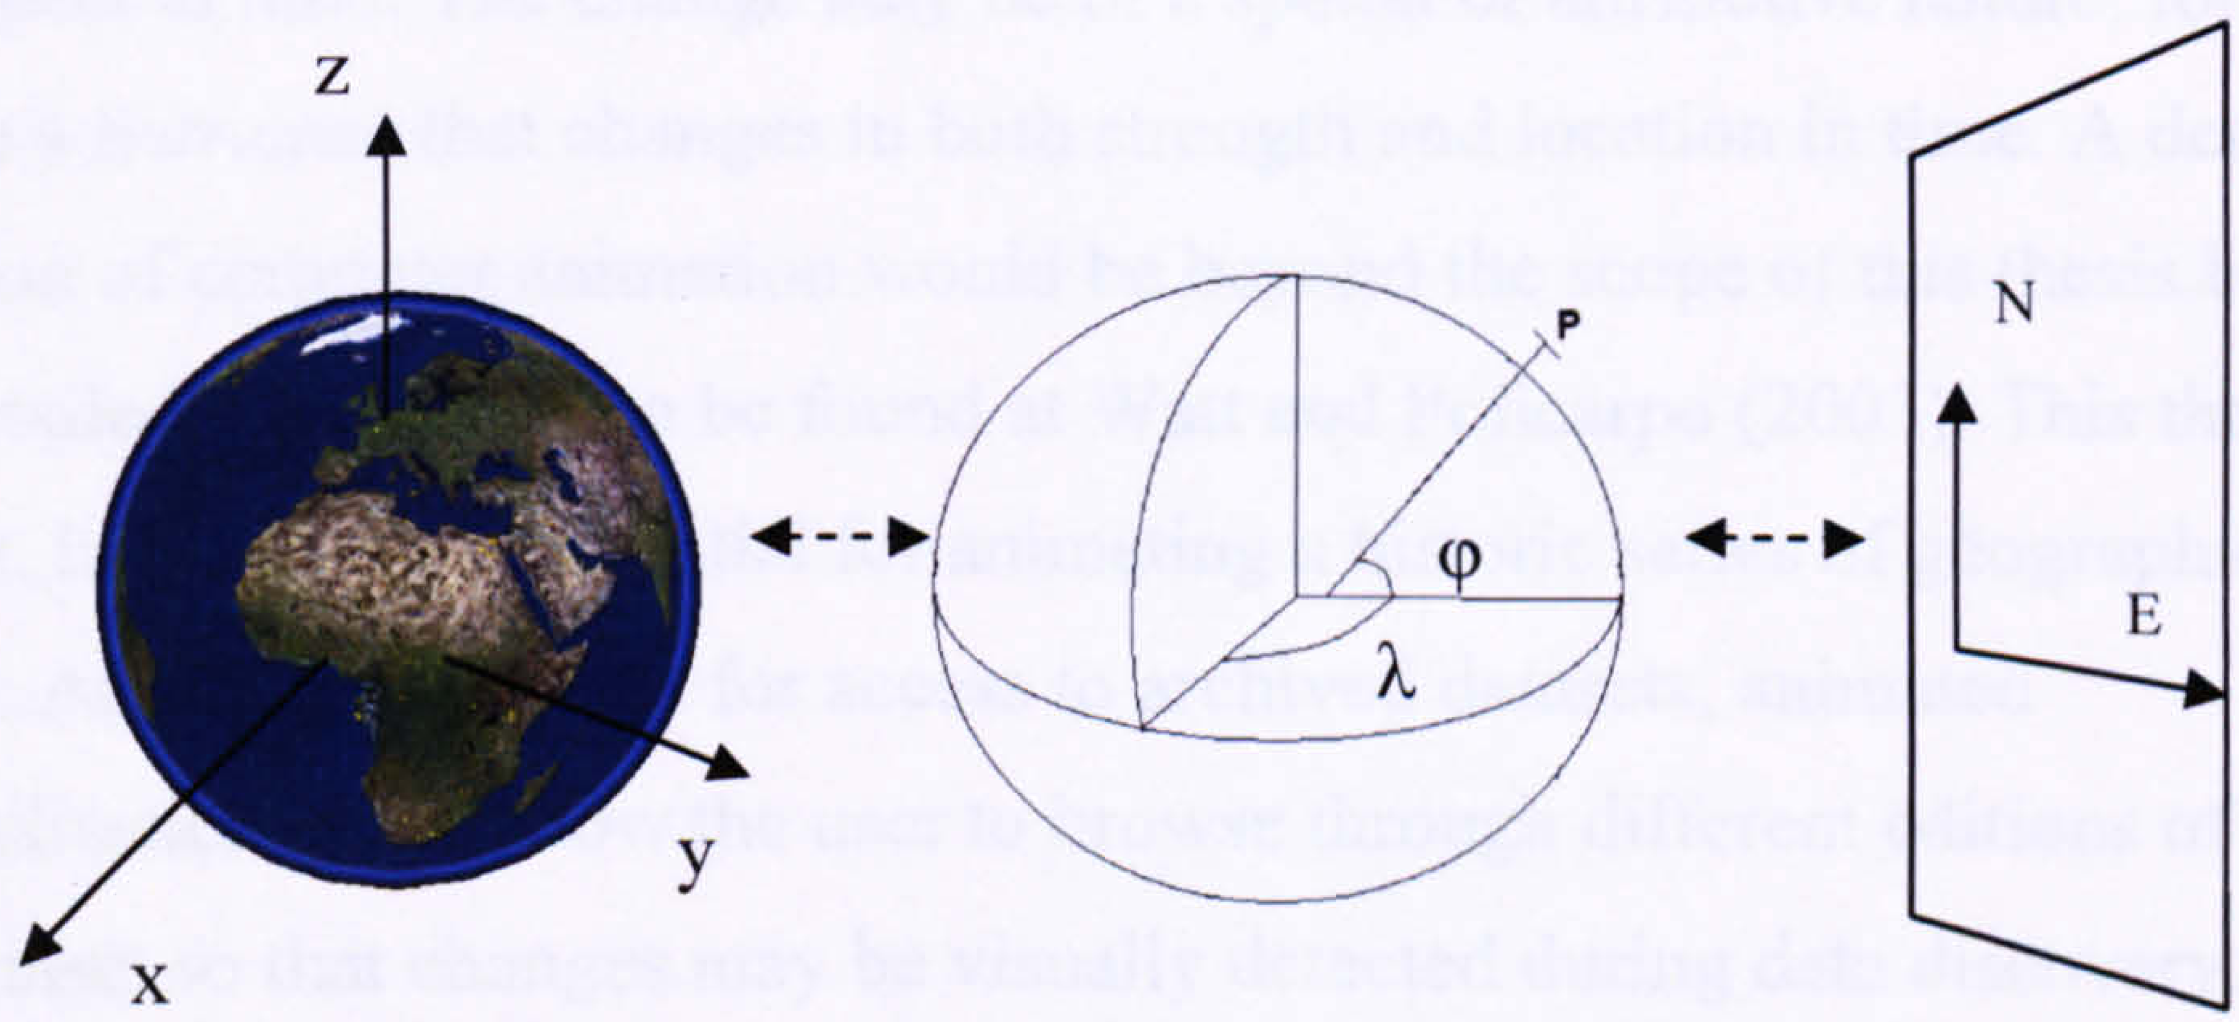

where:
- φ is geographic latitude
- λ is geographic longitude
- **P** is a point at a location on earth
- **X, Y and Z** are axes in geocentric Cartesian space
- **N** is North on a 2D map
- **E** is East on a 2D map

Figure 3.6 Coordinate Transformations between 2D, 3D geographic and Cartesian space

The Digital Earth metaphor, originally proposed by former US Vice President Al Gore (Gore, 1998), is drawn from the vision of 'zooming' in from space to the exact location of a dataset or information; effectively moving from very small-scale to very large-scale geovisualisations. For geoportals this could allow users to view atmospheric, astronomical, cadastral and meterological data within the same visualisation. Software implementing this concept include Google Earth, NASA

World Wind and ESRI ArcGlobe (Butler, 2006). The most popular of these packages is Google Earth, with over 1 500 000 downloads, followed by World Wind with over 500 000 downloaded (download.com, 2005). However, as Butler (2006: pp. 777) notes "Google Earth has no analytical functions and is not meant to replace professional GIS software". This could suggest that the public is appreciating the benefits of geovisualisation-supported IR, as Google Earth is primarily an IR application and does not offer spatial analytical capabilities.

With the axes in 3D Cartesian space representing spatial coordinates, the fourth dimension is commonly implemented to represent time in geovisualisation. Due to the inherent temporal nature of animations, they are often used to represent change with respect to time. The change may be of a spatial or attributive nature, for example a hurricane that changes in both strength and location in time. A detailed discussion of computer animation would be beyond the scope of this thesis but a more detailed discussion can be found at Watt and Policarpo (2003). This thesis, however, highlights the potential for animating a historic series of geographic datasets. As geoportals allow for access to archived datasets, animated geovisualisation could allow the user to browse through different editions of the same dataset so that changes may be visually detected during data discovery.

This subsection has examined the different approaches to 3D geovisualisation, differing in levels of realism. Also presented is the Digital Earth vision which forms the basis of popular software such as Google Earth. This thesis notes the possible benefit of the Digital Earth concept for geoportals serving earth observation information. Also highlighted in this subsection is the benefit of animating higher dimensional (4D and beyond) data during data discovery. The next section discusses the graphics technologies that make these visualisations possible.

3.3 Interactive Real-time Computer Graphics

The technologies enabling geovisualisation include interactive real-time 3D computer graphics hardware and software. Döllner (2005) offers a three level classification of computer graphics technologies, the classes are low-level, high-level and description languages. An illustration of their architecture is shown in Figure 3.7. Low-level computer graphics systems generally include two main technologies, OpenGL and DirectX. The former is an open technology, originally by Silicon Graphics Inc (SGI) and "is the most widely-adopted device-independent 3D graphics API" (Chen, 2002: pp. 111). On the other hand, Direct3D, the 3D rendering engine of DirectX, is a closed product of the Microsoft Corporation and is the de-facto standard 3D graphics engine for the Windows platform (Chen, 2002).

These low-level technologies offer a functional approach to the developer. This means that rendering follows a 'bottom-up' or 'step-by-step' approach, in contrast to an object-oriented approach which allows some functions to be grouped into a class. For an object-oriented approach, developers have implemented higher-level scenegraph Application Programmers Interfaces (APIs) that encapsulate algorithms and data structures that operate on lower-level OpenGL or Direct3D functions. "Although there are a few scenegraph APIs available, none have yet become dominant"(Angel, 2006:pp. 523).

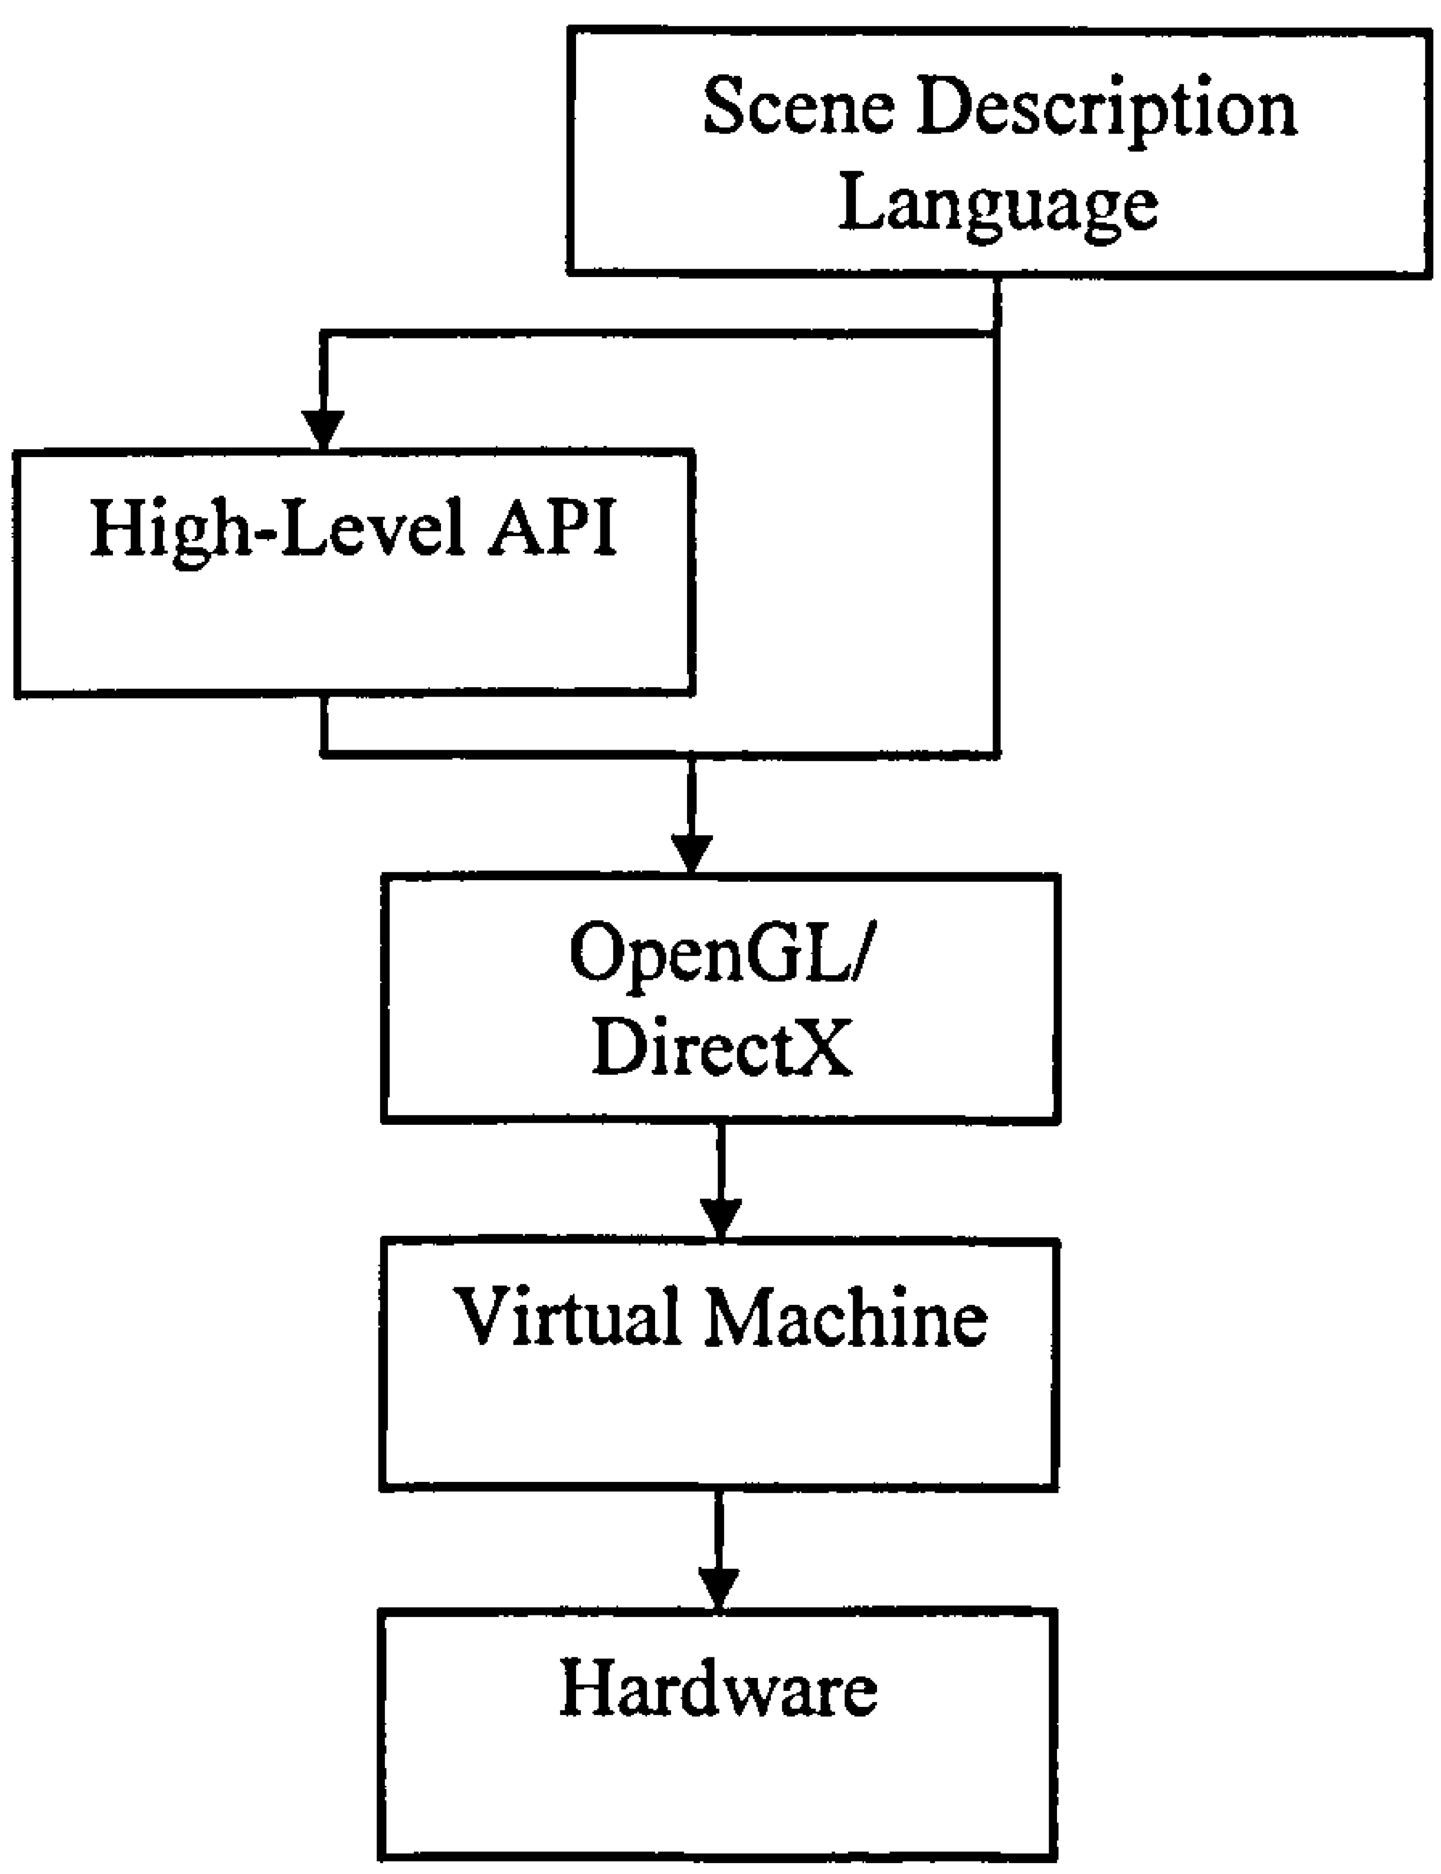

Figure 3.7 Modern Graphics Architecture (adapted from Angel, 2006)

Two examples of high-level graphics APIs are OpenGL Performer and Java3D. OpenGL Performer, originally by SGI, offers a scenegraph-based data model over OpenGL functions. Its programming platform is based on the C and C++ programming languages. A major benefit of OpenGL Performer is that it offers multiprocessing (Rohlf and Helman, 1994). This means that it can split rendering amongst different Central Processing Units (CPU) and ensure that they are all synchronised. Another benefit is that it monitors and modifies the frame rate according to user requirements when initiating a simulation (Seron et al., 2002). It also manages available memory resources by using auto-detecting high-level calls and reserving necessary memory for execution before entering the main simulation loop. A major disadvantage of using OpenGL Performer for our study is that it is not web-deployable meaning it must be installed on the client machine prior to use.

Another widely adopted API is Java3D, originally developed by Sun Microsystems. A user is offered the choice of installing either an OpenGL or DirectX wrapper. By virtue of being a Java API, it allows the developer to create 3D applications using a single language, Java, and embed additional Java functionality such as relational database connectivity. Similar to other Java-based applications, the disposing of created objects is handled by the 'garbage collector' of the Java Runtime Environment (JRE). Although this is a benefit for Java applications, for Java3D it results in slower frame rates due to higher memory consumption (Selman, 2002). In contrast, native-based APIs, such as OpenGL Performer, allow the developer to create 'deconstructors' that can be invoked to dispose of objects at any time during execution. Selman (2002) also highlights the tendency of Java3D applications to stall during garbage collection. Despite this disadvantage, Java3D offers several benefits for example, access to other Java APIs, platform-independence, client-side computation, web-deployment through Java webstart and applets. Although both are deployed via the web, Java applets are embedded on a webpage and webstart applications run as standalone applications.

At the highest level of the architectures illustrated in Figure 3.7 are the scene description languages (formats) such as the Virtual Reality Modelling Language (VRML), OpenFlight, eXtensible 3D (X3D), 3D Studio Max and others. These languages allow for the encoding of a 3D virtual world within a text or binary file. Of the aforementioned languages, VRML is the most widely adopted for web-based visualisation. The latest version, called VRML97, was formally adopted by the International Organisation for Standardisation as ISO14772. One of its main benefits is that it is encoded in ASCII text and hence is human-readable. Another benefit is that, because it is an ISO standard it is open and supported by several viewers. It also adopts a scenegraph data model similar to that of OpenGL Inventor and its successor OpenGL Performer (Chen, 2002).

However, there are disadvantages of using VRML, or its geographic modelling profile GeoVRML. First, it only offers single-precision floating point values; which means it is not possible to present geographic data resolutions greater than 10 to 100 meters (Rhyne, 1999, Moore et al., 1999). Second, it supports a fixed number of spatial referencing systems, specifically 3 coordinate systems, 21 ellipsoids and 1 geoid (GeoVRML WG, 2002). Third, at present most video cards will support texture sizes of 1024 by 1024 pixel dimensions therefore terrain representations larger than this require the geometry to be 'decomposed' into several smaller aggregates. This results in a heavier load on bandwidth, when using VRML for web-based visualisation, as all the aggregate textures have to be transmitted over a network (GeoVRML WG, 2002). Fourth, VRML is encoded in ASCII text which is verbose, in comparison to binary file formats (Horstmann and Cornell, 2001). VRML is therefore, slower to parse than its binary counterparts. Last, linking a VRML model to an external database required a use of the External Authoring Interface (EAI), however the EAI was particularly unstable as vendors developed their own custom implementations (Moore et al., 1999). X3D, VRML's successor, has addressed the issue of floating point numbers by offering double-typed numbers. However, other issues observed with VRML remain unaddressed by the X3D specification (Web3D Consortium, 2004).

3.4 Techniques of Web-based Visualisation

Within this section we consider current approaches for web-based visualisation with respect to geographic data and geographic searches. The section describes how a web-deployable high-level graphics API such as Java3D overcomes some of these limitations. The section focuses on web-based visualisation as geoportals and other GIR systems are similarly web-based.

3.4.1 Visualising Results of a Geographic Search

One of the stated research questions is "how can ontology be used to support visualisations of geographic search results?", this question will be addressed in this subsection. The previous chapter presented the view that geographic relevance is made up of at least three main properties – spatial, temporal and thematic (semantic). Although geographic data as described in metadata is highly dimensional, users find it difficult to visualise spaces of more than three dimensions (de Oliveira and Levkowit, 2003, Fabrikant and Buttenfield, 2001). We therefore limit this discussion to the three aforementioned properties of geographic relevance. A challenge for presenting geographic relevance based on three-criteria ranking is that textual ranked lists, as already discussed in the previous chapter, are in effect one dimensional as they list search results from top-bottom or vice versa. Consequently, some studies have investigated the possibility of presenting search results in multidimensional visualisations. We discuss some of these in this subsection.

In their study, Leuski and Allan (2004) propose spring-embedded visualisation for visualising clusters of ranked documents. Their approach, named Lighthouse, scales the multidimensional similarities of documents down to Euclidean space and maps representative spheres in Euclidean space such that the distances between the spheres correspond to the dissimilarities between the original documents as closely as possible (Leuski and Allan, 2004). Figure 3.8 shows a simplified depiction of the Lighthouse system. They observed that users constantly pointed out that accurate identification of the inter-document distances in 3D required frequent rotations of the structure thus making the visualisation task more difficult. They hence concluded that although 3D spring-embedded visualisations were more accurate than 2D visualisations in their portrayals of inter-document similarity, users still found it easier to use 2D visualisations for their IR tasks. Although users appeared to prefer 2D representations in comparison to the Lighthouse system, we highlight that the axes are not calibrated; therefore they do not offer a visual aid for detecting distances or depth. Further, the Lighthouse system ignores the spatio-temporal properties of documents as it is designed for traditional IR, therefore it does not address spatial data discovery.

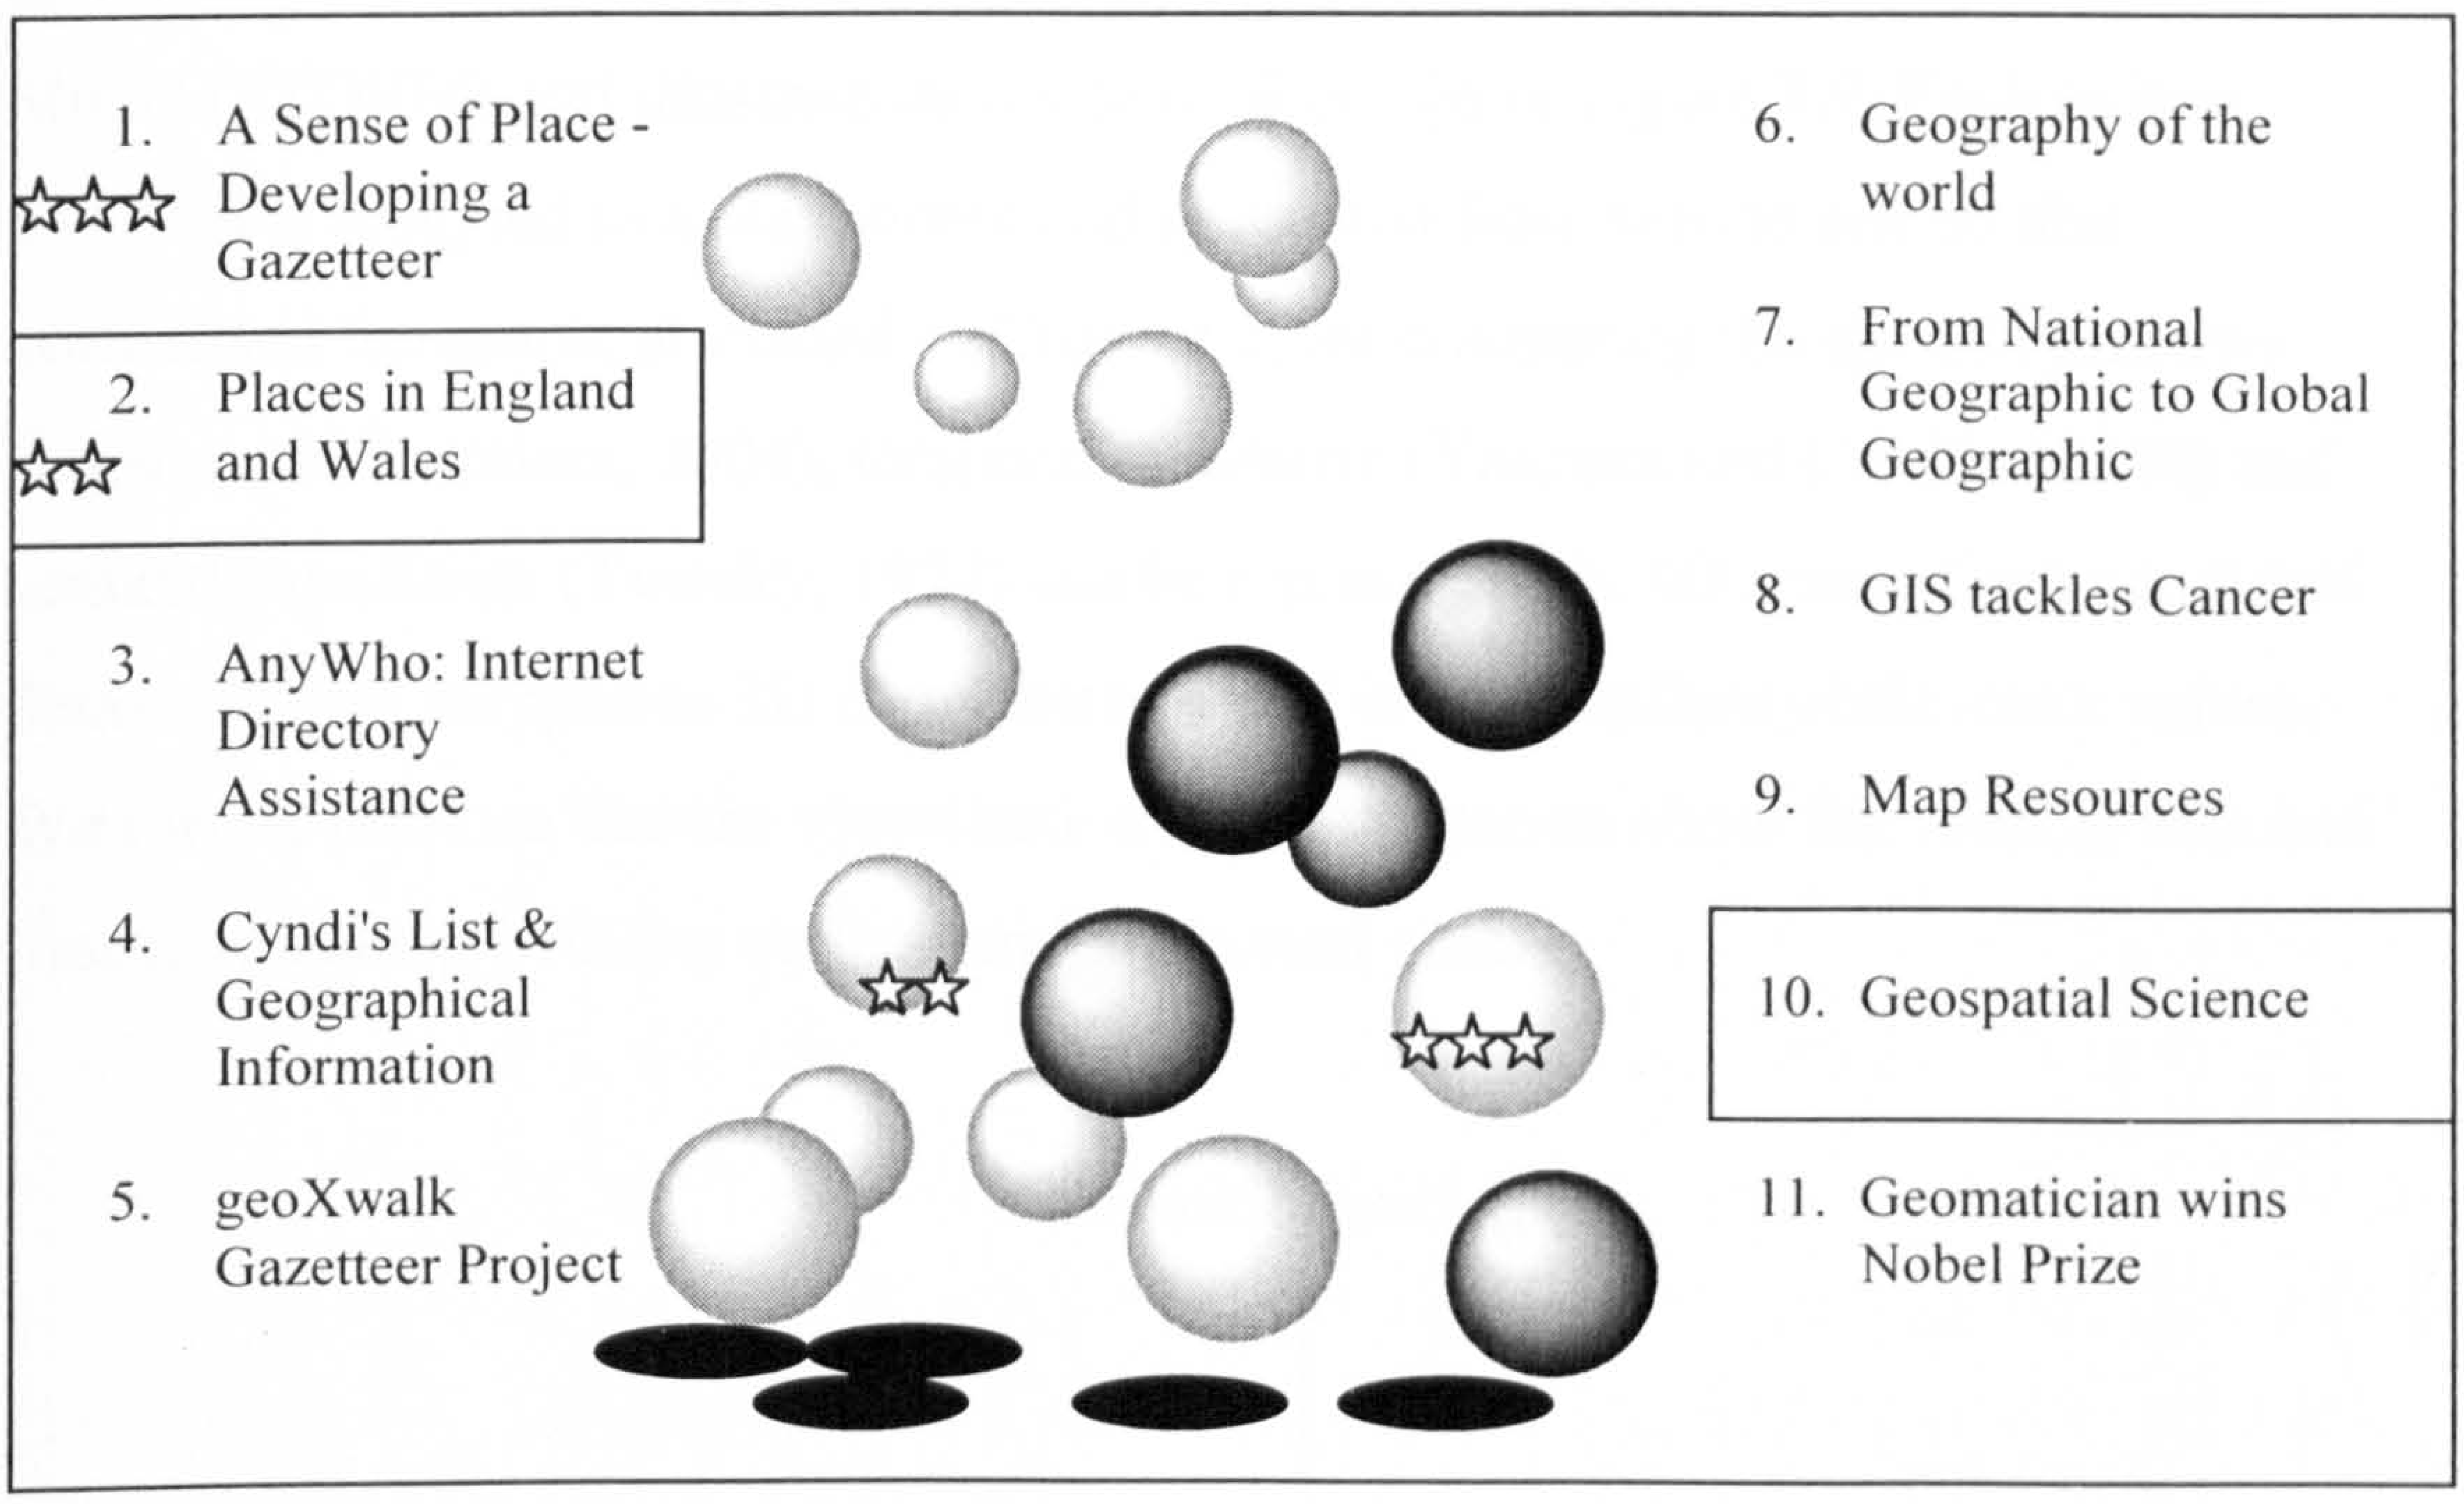

Figure 3.8 Depiction of the Lighthouse visualisation (Leuski and Allan, 2004)

Similarly, Göbel et al., (2002) use a 3D visualisation metaphor to present search results on their GeoCrystal system. Although the visualisation is presented in a 3D environment, a vector map occupies the xy plane and the spatial footprints of each document are rendered on this vector map. This means that only a single dimension is left for both the temporal and the thematic properties. Further, GeoCrystal ignores the proportion of the area covered by the candidate document to that covered by the query, a consideration recognised by other studies by Larson and Frontiera (2004) and Beard and Sharma (1997). Another disadvantage of the GeoCrystal visualisation metaphor is that it does not provide a visual aid for interpreting inter-document proximity within the 3D presentation; this makes interpretation difficult when not in stereo-view. However, an advantage of the system is that it acknowledges the spatial, temporal and thematic characteristics of geographic metadata, thereby supporting heterogeneous perceptions of geographic relevance.

This thesis proposes the use of 3D space to illustrate the degrees of relevance of geographic data based on the three components of geographic relevance already identified. We name the approach the Spatio-Temporal Ontological Relevance Model (STORM) and illustrate its conceptual design in Figure 3.9. Each axis in STORM is assigned to a component and graduated from zero to one so that normalised functions, discussed in Chapter 2, for computing the spatial similarity (Larson and Frontiera, 2004), temporal similarity (Yamuna and Candan, 2000) and semantic similarity (Tversky, 1977) can be represented in 3D space. Each candidate dataset is then mapped to 3D coordinates based on the similarity/relevance values. We further propose that the visualisation support rotation about the vertical axis and translation along both the vertical and horizontal planes.

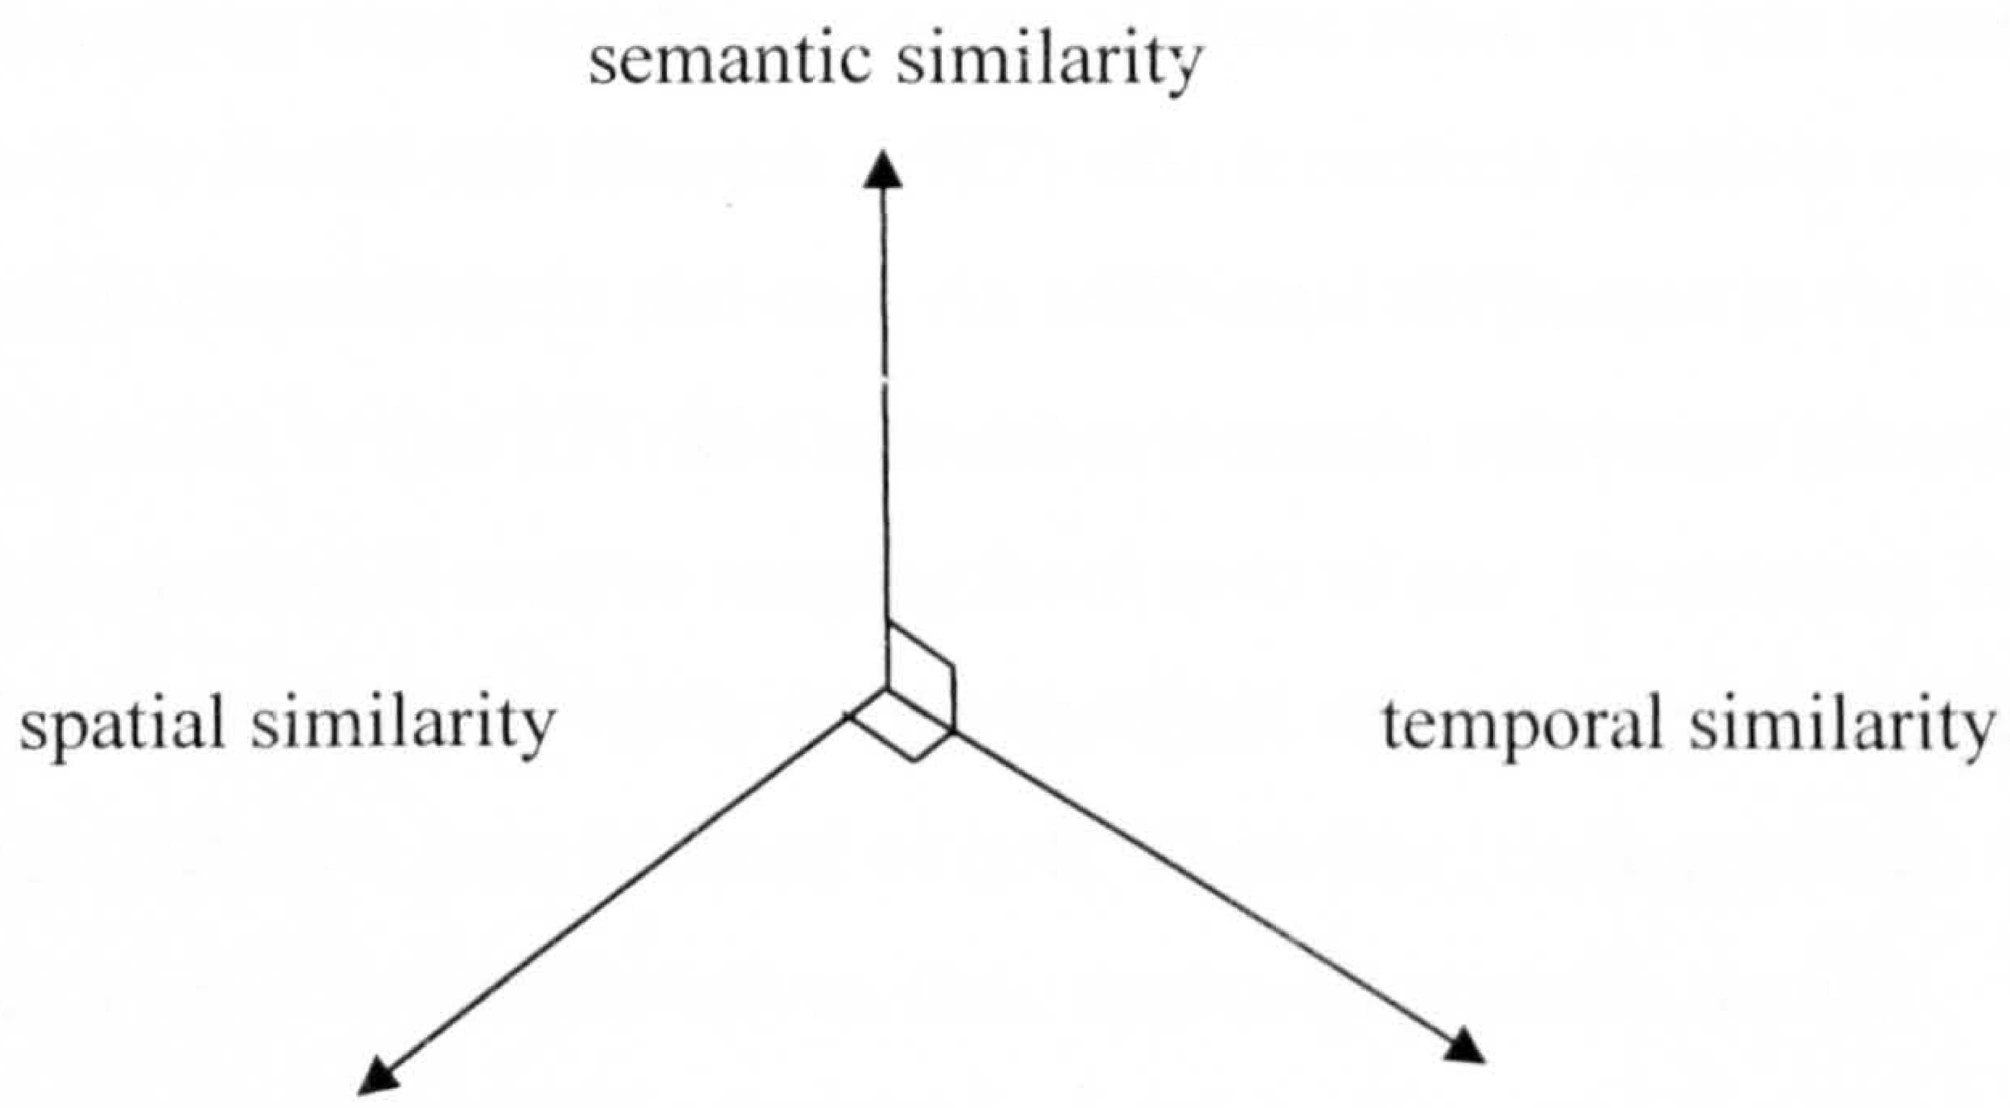

Figure 3.9 The STORM vector space model for representing the relevance of geographic data

The similarity measures are normalised such that the least relevant documents are positioned closer to coordinate (0,0,0) and the most relevant documents closer to (1,1,1). The normalised axes mean that any similarity function that returns values ranging from zero to one can be assigned to an axis within the STORM environment. The sensitivity of each of these measures is represented by the scaling of all axes from zero to one, thereby allowing positions to increase or decrease along each axis through equal units. With documents presented as thumbnails, the geometric size of the thumbnails is intentionally kept constant so as to help the user detect distance easily through perspective views (that is, thumbnails closer to the viewer's position appear larger than those that are further away). Further, when viewed from coordinate (1,1,1), documents that occlude others are more relevant than those occluded. This feature demonstrates a benefit of both perspective and occlusion (blocking of the visibility of objects in 3D space) for detecting the relevance of represented documents relative to others. "Occlusion is the single most important depth cue, over-riding others and because of occlusion we can see far less information in the *z* direction" (Ware and Plumlee, 2005: pp. 570).

The STORM approach is unique to other presentation approaches adopted in GIR in a number of ways. Unlike ranked textual lists, STORM offers a non-linear view of relevance. This is because STORM offers support for rotation and translation of the visualisation such that a user can select varying views of relevance. Therefore, even though the axes are linear with respect to each geographic property (space, time and theme), the overall illustration of relevance is non-linear. STORM presents relevance scores as a fraction ranging from zero to one along all three axes; this is in contrast to the shading approach by Beard and Sharma (1997) which restricts possible relevance values to four intervals between zero and one. An additional difference to the Beard and Sharma (1997) model, is that STORM calculates thematic relevance according to semantic similarity, resulting in a value ranging from zero to one. In contrast, Beard and Sharma (1997) offer a Boolean value for the thematic relevance of each thematic property (indicating whether a term is found or not). Therefore, their approach ignores semantically similar terms. However, their approach presents several thematic properties separately, whereas STORM combines them into a single relevance measure. STORM is different to the Lighthouse approach (Leuski and Allan, 2004) as all three axes in STORM vector space are formally assigned to specific relevance measures based on spatial, temporal and semantic properties. The assignment of axes to distinct properties also demonstrates a difference to the GeoCrystal approach (Göbel et al., 2002), which assigns two out of three dimensions to a 2D cartographic map.

This subsection has discussed various 3D metaphors for visualising search results on both an IR and GIR system. A metaphor for the 3D visualisation of search results based on spatial, temporal and semantic properties is proposed. The example metaphors show that there is active research looking into 3D visualisations for offering visual interfaces for GIR systems. However, different compromises are made due to clarity, dimensionality and usability. We conclude from these examples that due to the varying information needs of users and their perceptions of geographic relevance, it is important for GIR research to continue to develop different metaphors of visualisation.

3.4.2 Visualisation of Geographic Data

One of the research questions that was stated is "from search results, how can a multidimensional geographic dataset be visualised in greater detail to determine its relevance?", this question will be addressed in this subsection. In the previous subsection we discussed the visualisation of degrees of relevance of geographic data. Unfortunately, relevance ranking algorithms cannot provide the final judgement on relevance as searchers themselves must be the ultimate judges as to the relevance of a listed document (Chen, 2004). It is necessary to provide the searcher with a preview of the data to enable her to reach a final decision of relevance. This subsection discusses geovisualisation techniques and proposes an approach for creating previews of multidimensional geographic data.

3.4.2.1 Vector based Modelling

Vector-based modelling follows two main approaches: geometry and texture-based mapping. The former involves the creation of a 3D model based on representative geometries of real-world objects. Traditional approaches used VRML for describing the geometries of objects within a virtual world. This meant that the solutions were subject to the same limitations of VRML already identified earlier in this chapter, namely single floating point numbers, few coordinate systems, limited terrain representation and vendor-specific constraints. We propose an approach for web-based 3D geovisualisation through the direct retrieval of geographic objects from a Relation Database Management System (RDBMS). The geographic objects are based on geometry primitives conforming to the OGC Simple Features specification (SFS)

(Open Geospatial Consortium, 2005). The SFS model is illustrated in Figure 3.10 and includes the basic geometric primitives of points, lines and polygons. The approach is implemented within the Geospatial Database Online Visualisation Environment (GeoDOVE), a web-based GIS.

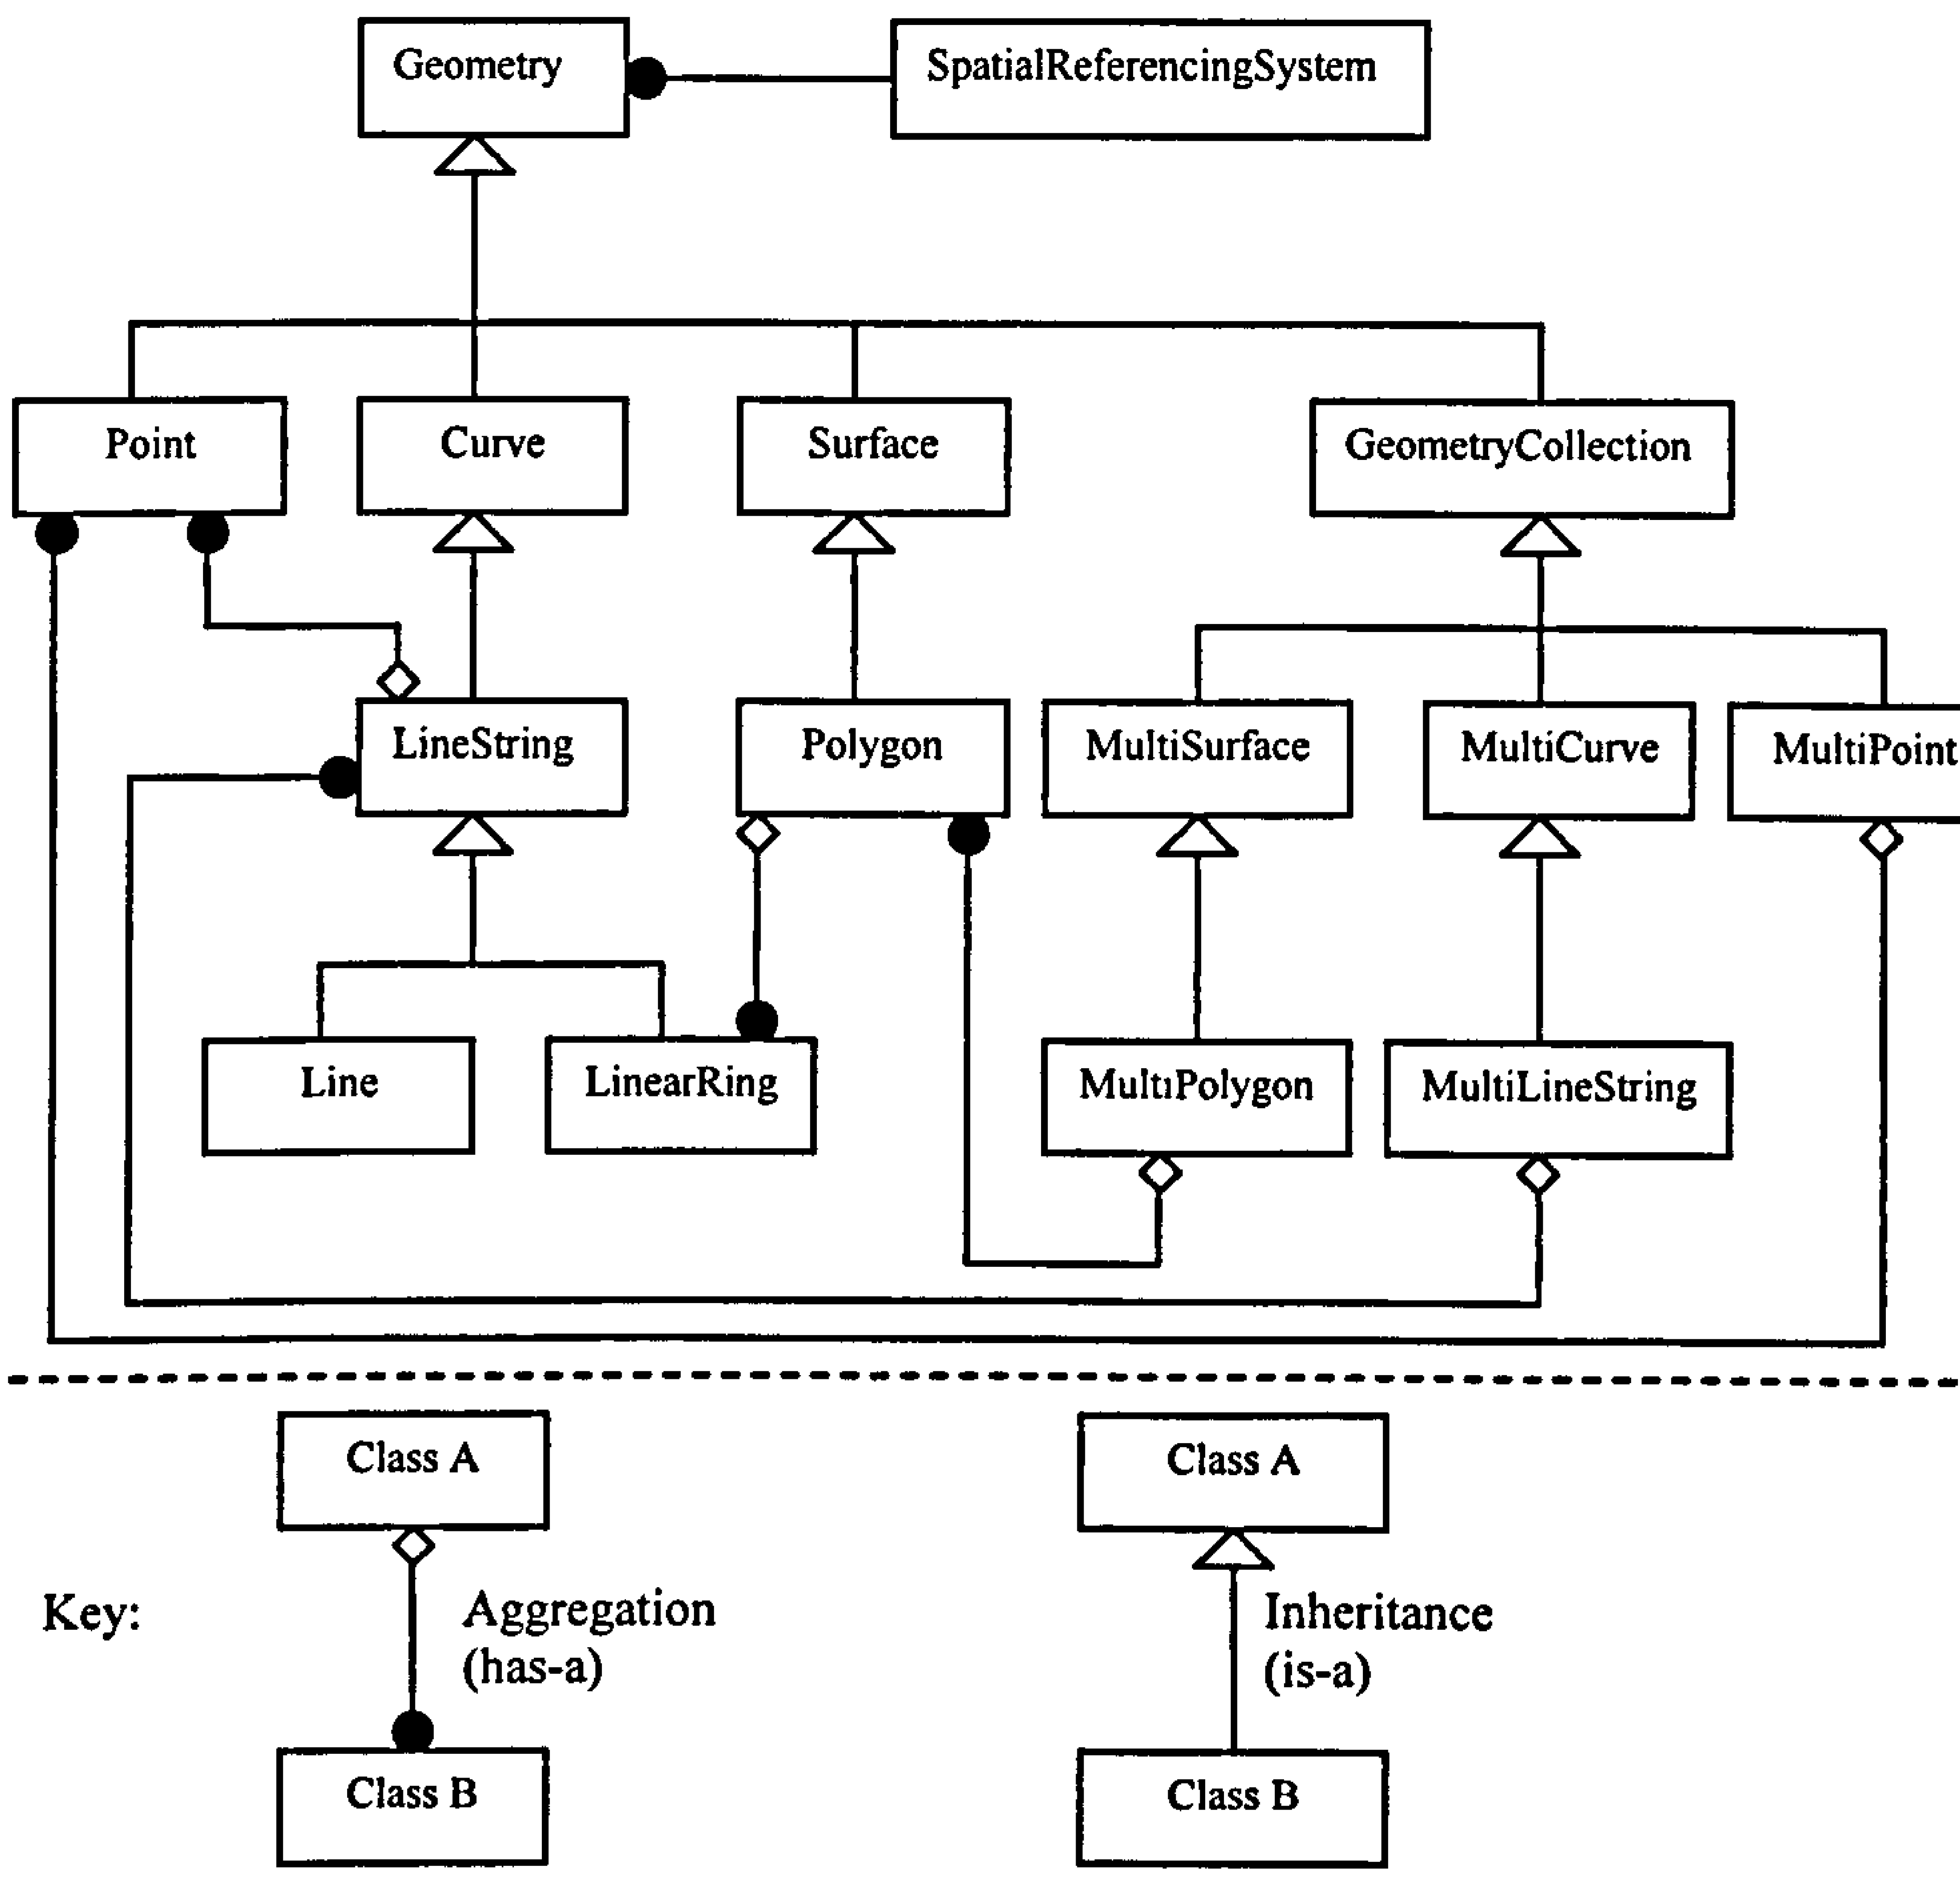

Figure 3.10 The Simple Features model(adapted from OGC, 2005)

Using Java3D as an example, the proposed approach can be implemented in the following way. Java3D uses the `Shape3D` class for the creation of three dimensional shapes. Each instantiation of the `Shape3D` class contains instantiations of `Geometry` and `Appearance` classes. Within GeoDOVE, SFS `MultiLineString` and `MultiPolygon` geometries are represented through the extrusion of their aggregate SFS `LineString` objects. Figure 3.11 illustrates the extrusion process. The wall resulting from the extrusion is built by instantiating a Java3D `QuadArray` object from a pair of adjacent vertices on the line and their associated vertical vertices. Each pair of adjacent vertices therefore results in an individual rectangular plane. Creation of extruded `LinearRing` walls follows the same process as `LineStrings`. The representation of SFS `Point` features is achieved through extrusion of a circle (geometrically a cylinder centred at the SFS `Point`) defined by the following formula.

$$y = y_0 \pm \sqrt{r^2 - (x - x_0)^2} \qquad \text{for } (-r < x \leq r);$$

Where:

- x_0, y_0 are coordinates of the SFS Point
- r is the radius of the circle
- y is the corresponding coordinate for each x coordinate

When representing objects such as buildings from polygonal geometries, it is necessary to create a roof over the extruded polygon. Therefore extruded `LinearRing` objects (defining polygonal geometries) are roofed through the modelling of a roof that adopts the shape of its polygonal base. In 3D geovisualisation, roofs have traditionally been modelled as flat roofs. However, computational geometry offers an algorithm called the 'Straight Skeleton' for modelling pitched roofs. The algorithm translates each edge of a polygon towards the centre of the polygon at a fixed rate until the shrinking polygon has an area as close to zero as possible (Felkel and Obdrzalek, 1998). The polygon vertices move inward along the angular bisectors formed by the edges of the polygon. Suggestions for the use of the Straight Skeleton approach within geovisualisation have been made by Laycock and Day (2003) and Döllner and Buchholz (2005).

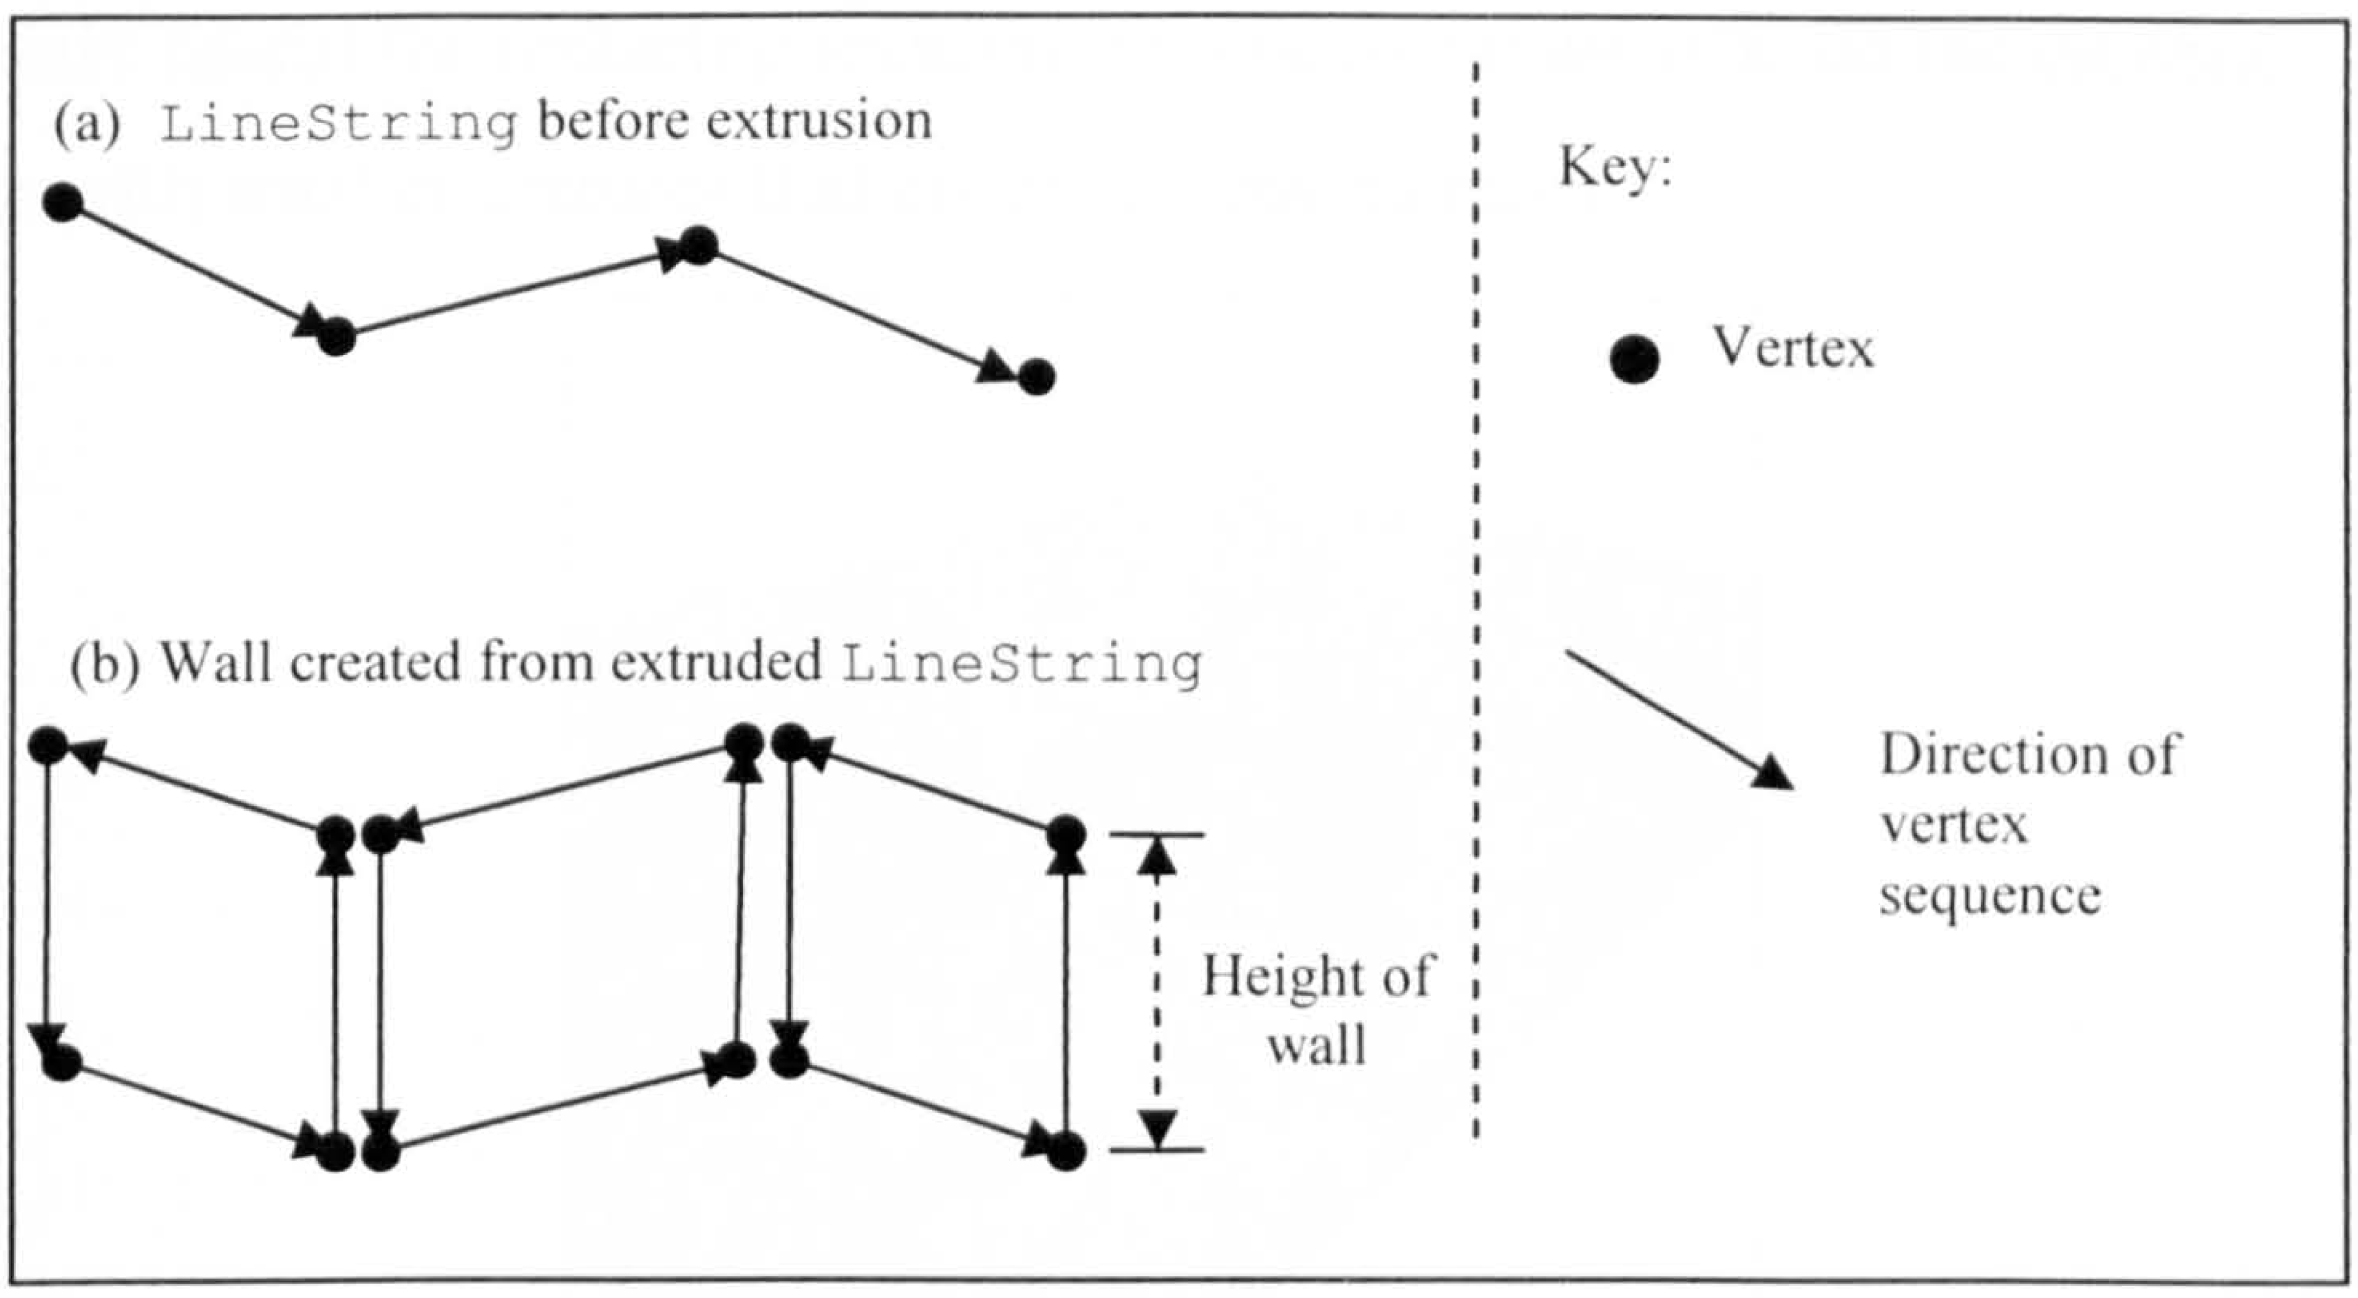

Figure 3.11 The extrusion process

An alternative approach for vector-based modelling involves texture-based mapping — rendering vector objects onto a texture draped on a DEM (Kersting and Döllner, 2002). This technique has the advantage that the complete vector representation is overlaid above the DEM. However, it produces a 2D representation on a 3D surface and hence when viewed in perspective view none of the represented objects have height. Despite this disadvantage, this thesis acknowledges that it is an ideal approach for viewing a DTM from a camera position directly above the terrain where viewing all geometric objects could otherwise consume significant memory resources. Having examined the different approaches for vector-based modelling, we next examine raster-based modelling.

3.4.2.2 Raster and TIN-based Modelling

Raster-based modelling techniques are used for the creation of a DEM from a regular grid of points or from a TIN. Often textures are draped on top of the DEM. Techniques include multi-texturing, dynamic texture generation and multi-resolution modelling (Döllner, 2005). Multi-texturing involves the assignment of multiple images to the same 3D shape. An example of multi-texturing is the merging of different thematic maps, rendered on individual images, into a single texture that is then draped onto a DEM. Multi-resolution modelling of raster data is usually implemented through the Level-Of-Detail feature offered by most scenegraph APIs.

It is mostly useful for replacing textures on objects far away from the camera position with smaller textures that consume less memory.

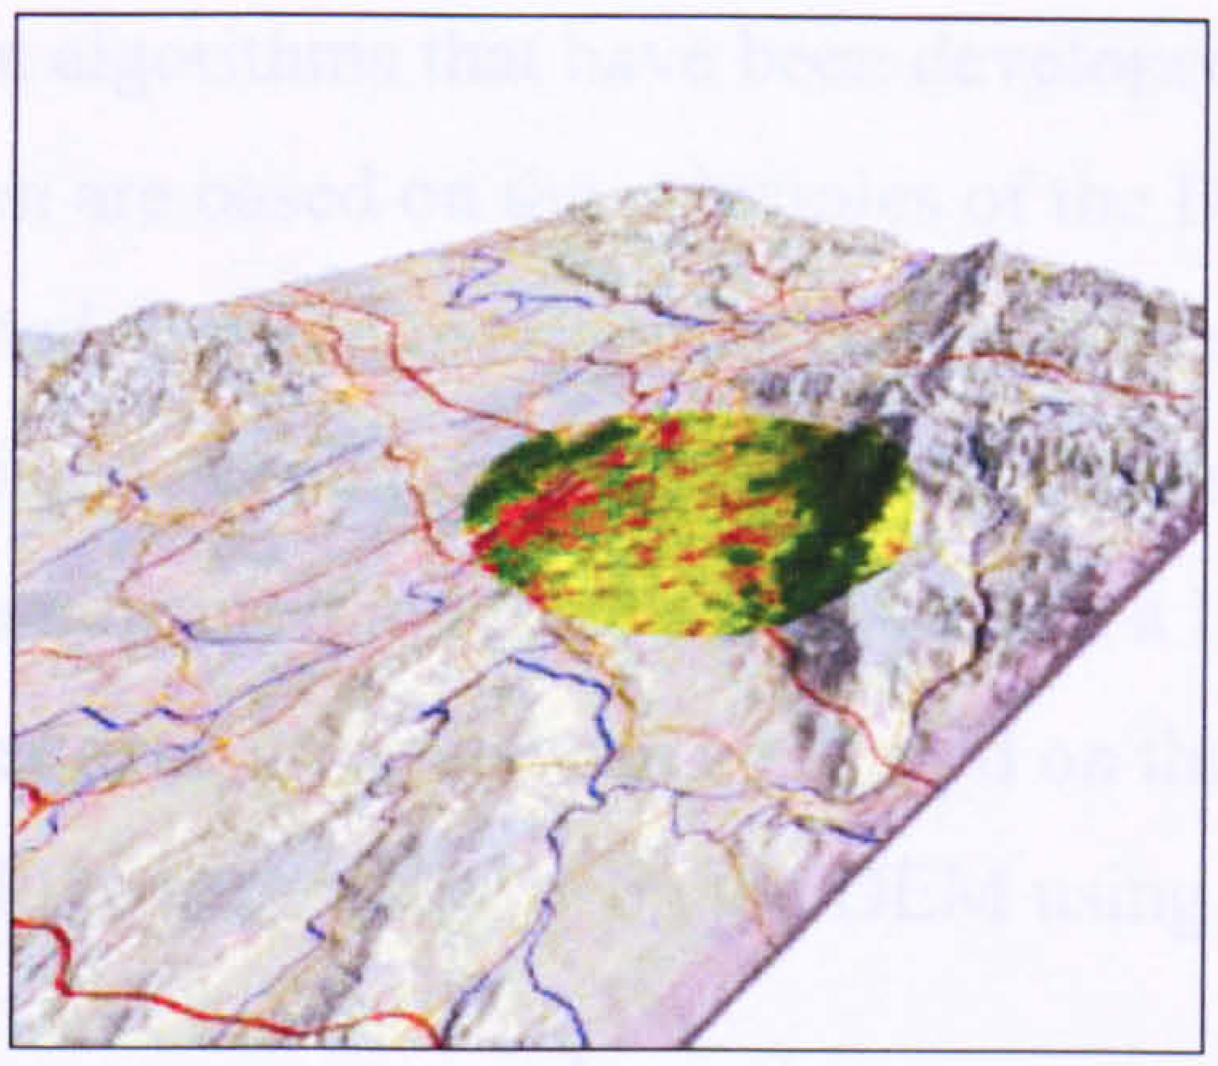

Figure 3.12 Lens creation through dynamic texture generation (Döllner, 2005)

Dynamic texture generation involves the creation of textures during runtime and often as a result of user interaction. Figure 3.12 illustrates dynamic texture generation through lenses rendered on a DEM and interactively controlled by the user. As the textures are created during run-time and in real-time, it is necessary that an application be client-side as transmission to a web server would result in longer response times for real-time visualisations. A scene description language such as VRML or X3D would be subject to such lengthy response times. In contrast, Java3D applets allow for dynamic texture generation as the textures can be rendered on the client-side. This is evidenced through GeoDOVE's gray-scale and multi-coloured texture generation created from a regular grid of spots heights. Figure 3.13 shows two textures automatically generated from the same elevation data.

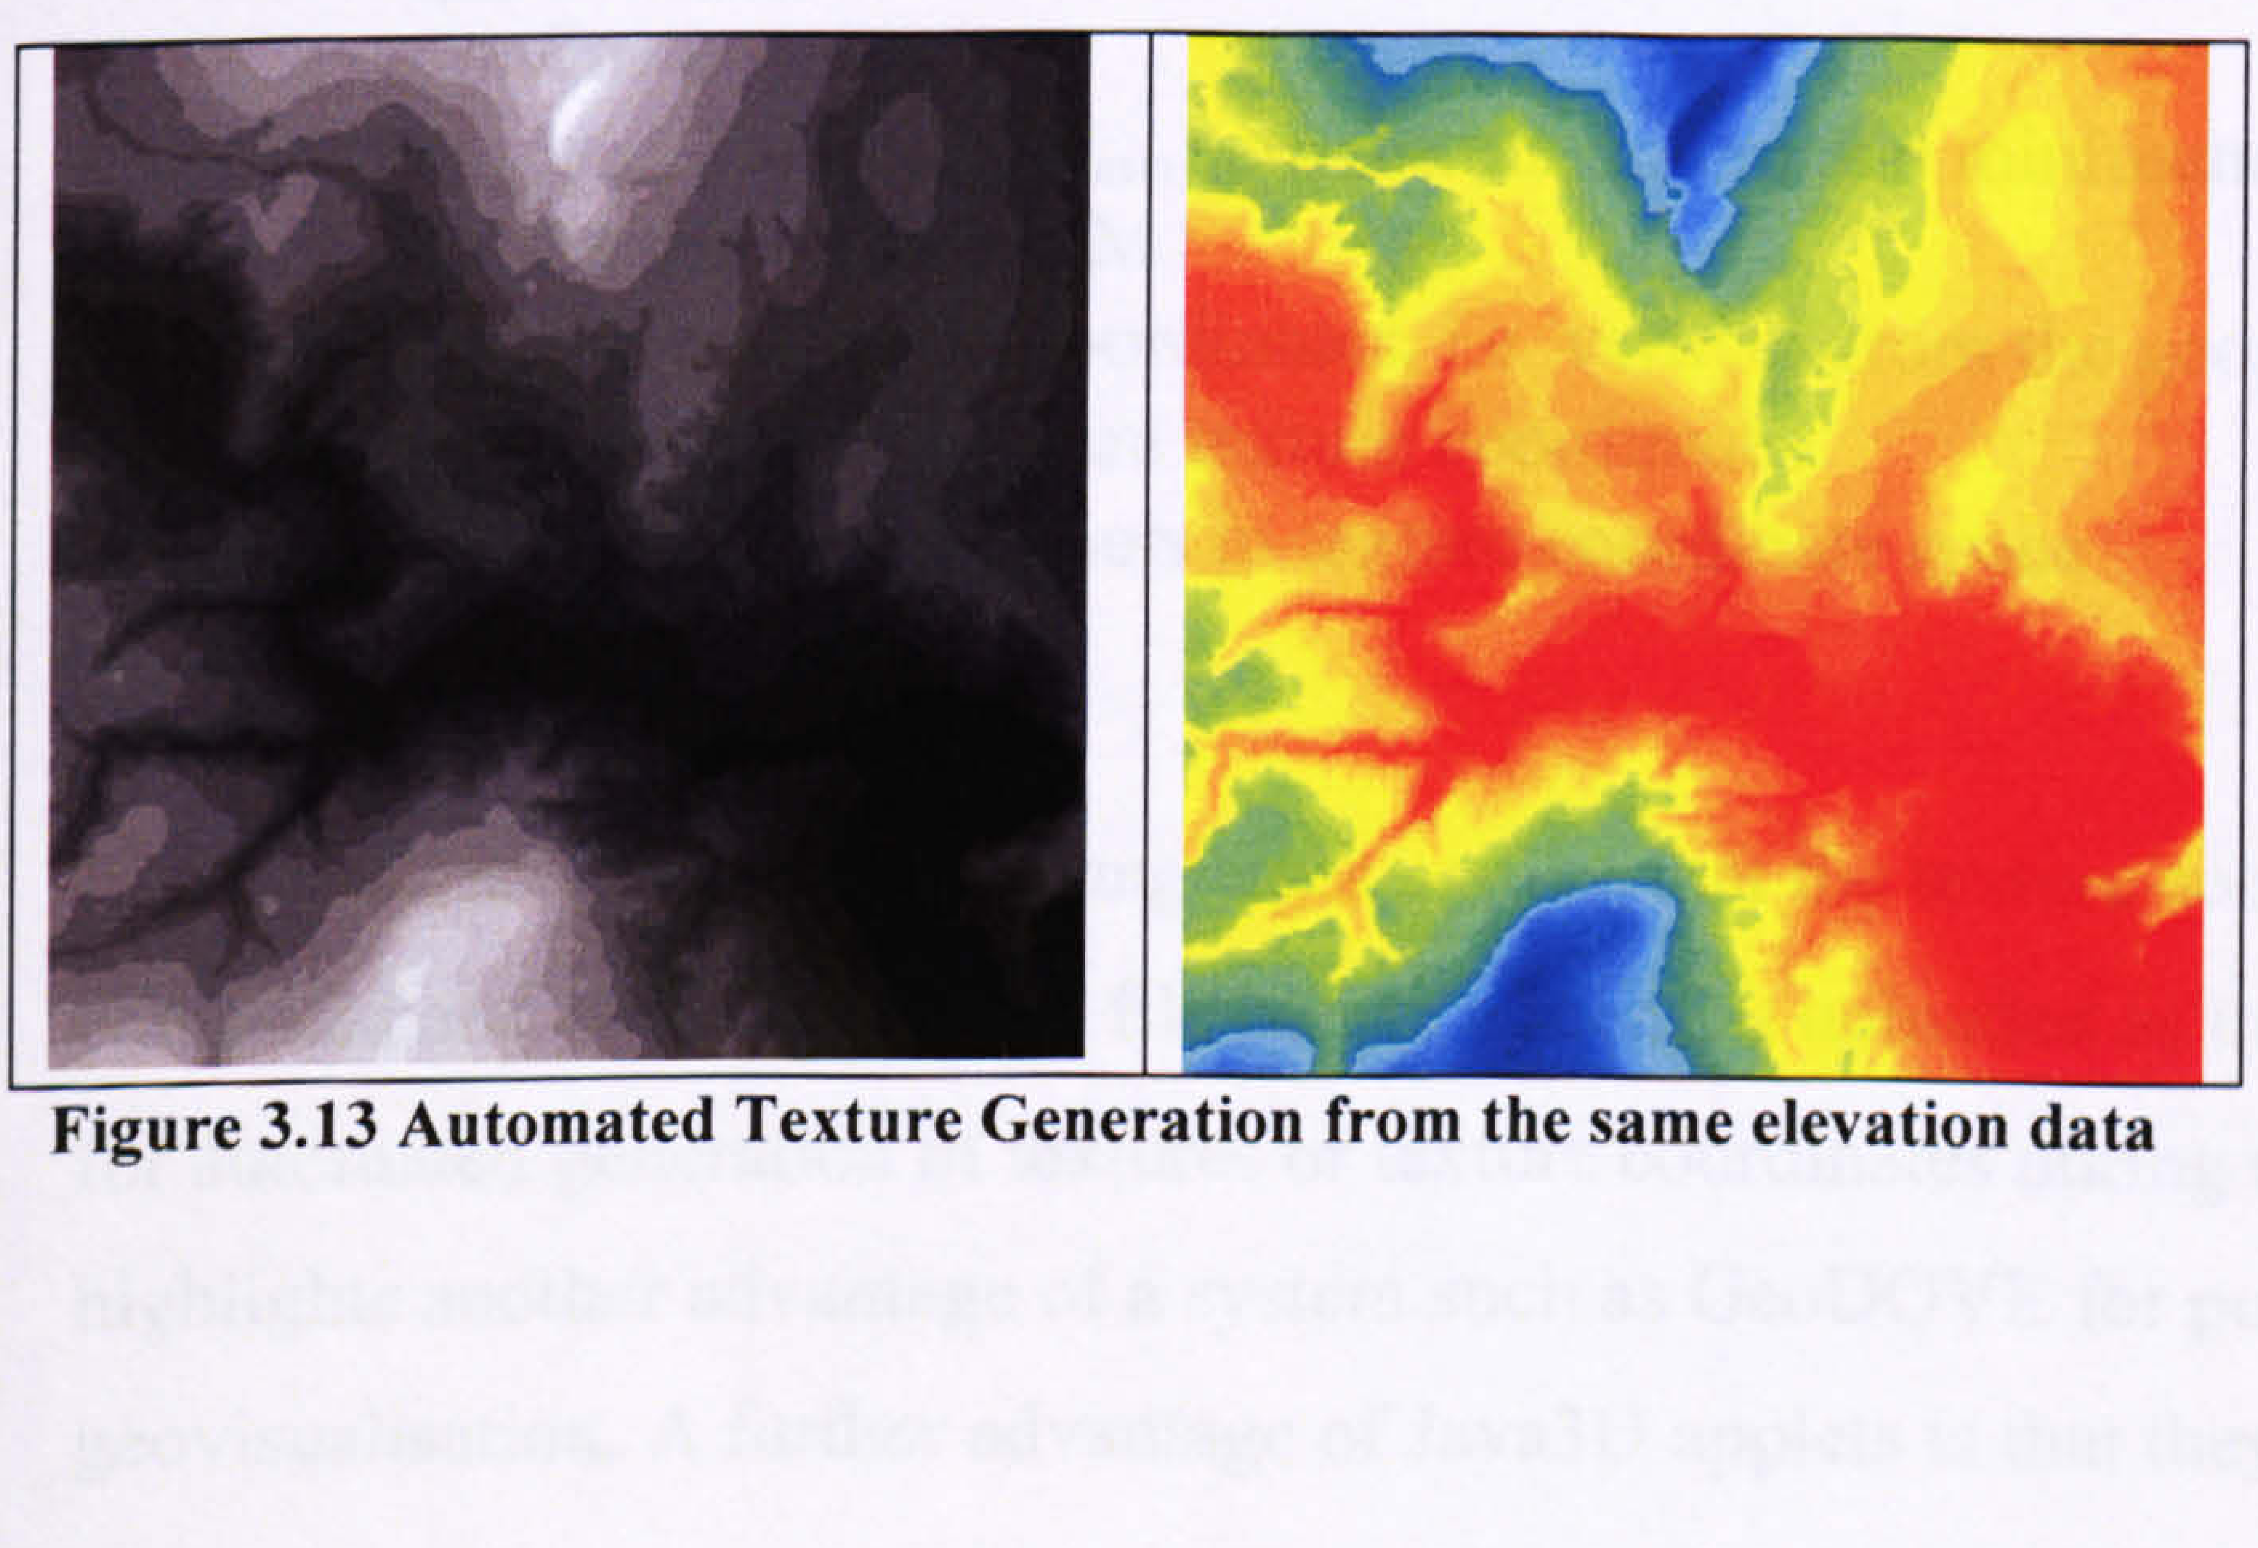

Figure 3.13 Automated Texture Generation from the same elevation data

TIN-based modelling involves the creation of a surface from irregularly positioned points usually through a triangulation algorithm, as illustrated in Figure 3.14. There are several triangulation algorithms that have been developed to date (Abdelguerfi et al., 1998), however most are based on the principles of the Delaunay triangulation — an aggregate of connected, but non overlapping triangles such that the circle defined by the points in each triangle contains no other points. Similar to raster datasets, TIN-based modelling can be used to create a DEM. Once a DEM has been created from either TIN or raster data, an image can be draped on the geometry by mapping texture coordinates (S,T) to vertices (x,y) on the DEM using the following equation.

$$s_{scale} = \frac{1}{x_{\max} - x_{\min}}$$

$$t_{scale} = \frac{1}{y_{\max} - y_{\min}}$$

$$s_{offset} = -x_{\min} \times s_{scale}$$

$$t_{offset} = -y_{\min} \times t_{scale}$$

$$s_{coordinate} = s_{scale} \times x + s_{offset}$$

$$t_{coordinate} = t_{scale} \times y + t_{offset}$$

where:

- s_{scale} and t_{scale} are scale factors from real word coordinates to texture coordinate space
- (x_{max}, y_{max}) and (x_{min}, y_{min}) are the maximum and minimum bounding coordinates of the DEM
- (x,y) are real world coordinates, i.e. vertices from the geographic data
- $s_{coordinate}$ and $t_{coordinate}$ are texture coordinates
- s_{offset} and t_{offset} are offsets in texture coordinate space

In contrast, scene description languages require that the texture coordinates be computed and packaged with a file before a scene is uploaded as they do not allow for automated generation of textures or texture coordinates during runtime. This highlights another advantage of a system such as GeoDOVE for purposes of 3D geovisualisation. A further advantage of Java3D applets is that they allow for the calculation of base heights of shapes, during run-time, using an underlying DEM as a reference. Although some non-geospatial viewers of scene description languages can alter a scene during runtime, they seldom offer geographic operations and thus lack the ability to extract height information from DEMs.

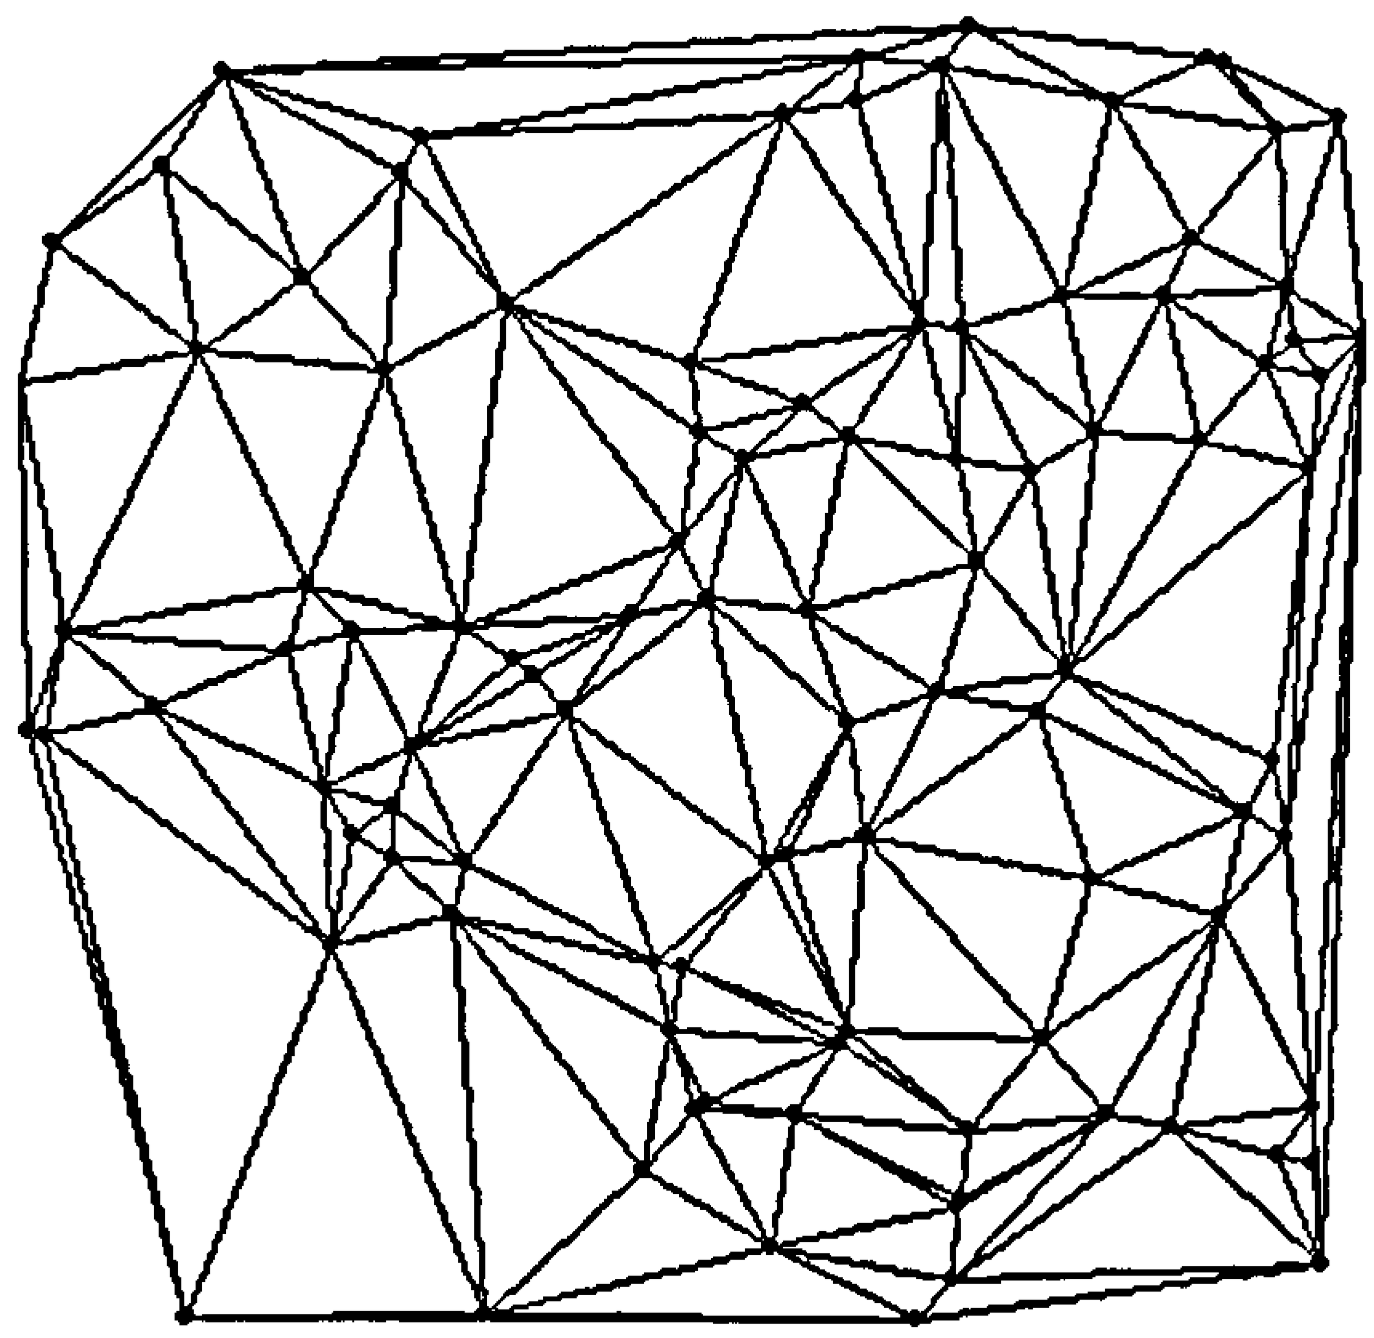

Figure 3.14 An illustration of the Delaunay Triangulation

This section has discussed different techniques for the web-based representation of objects within a 3D scene. It should be noted that the modelling techniques are directly transferable to most scenegraphs, including OpenGL Performer, as they are based on algorithms involving the same geometric primitives; however not all scenegraph APIs offer web-deployment and platform-independence. The next section presents some studies on the comparison between 2D and 3D visualisation.

3.6 Chapter Conclusion

This chapter has discussed the different types of geovisualisation approaches of both 2D and multidimensional geographic data. Geovisualisation approaches presented include small-area DTMs and global views based on the Digital Earth concept. The integration of 2D and 3D visualisation applications within mainstream search engines such as Google indicates the growing acceptance of geovisualisation-supported IR. The chapter has also presented suggestions from studies comparing 2D with 3D visualisation. From the comparative 2D and 3D studies, we conclude that

3D visualisation should be regarded as a complement of 2D visualisation approaches, rather than as a replacement.

This chapter has also emphasized the need for improved metaphors of visualisation to support the varying perceptions of geographic relevance discussed in Chapter 2. To this end, an approach for web-based visualisation of geographic data is presented. A prototype based on the proposed approach, called GeoDOVE, is introduced. Similarly, an approach called STORM, for visualising results of geographic searches during data discovery, is proposed. Whereas this chapter presented the conceptual basis of the two proposed approaches, chapter 5 will describe the developmental aspects of their prototypes. However, the next chapter describes mechanisms for delivering data to these and other visualisation applications offered on geoportals.

Chapter 4 Web-based Delivery of Geographic Information

The previous chapter discussed and proposed new methods for the visualisation of geographic data and results of a geographic search. In order to realise the models proposed in the previous chapter, considerations have to be made regarding the way in which data is delivered to the visualisation applications discussed. The computing and geographic information communities have developed two main methods for disseminating geographic data, through geographic web services and through spatially-enabled database engines. Consequently, the OGC has published specifications to standardise the implementation and operation of the two aforementioned methods of disseminating geographic information. Some of the specifications of the OGC extend on other specifications by the World Wide Web Consortium (W3C). In this chapter, we examine their role within a geoportal environment.

This chapter adopts the following geo-centric definition of interoperability "geographic interoperability is the ability of information systems to 1) freely exchange all kinds of spatial information about the Earth and about the objects and phenomena on, above, and below the Earth's surface; and 2) cooperatively, over networks, run software capable of manipulating such information."(OGC, 2001a)

4.1 Dissemination through DBMS Connectivity

Access to a database adds value to geographic data discovery as the user (searcher) can view attributive content prior to acquiring or purchasing it from a geoportal. Due to the limitations of historic web-based geovisualisation approaches, our study examined the architectures adopted in the development of desktop GIS as they offer both geovisualisation and database handling capabilities. Desktop GIS have been able to access vast amounts of data from distributed DBMS since before the turn of the millennium. In contrast, historic web-based multidimensional geovisualisation was based on dynamically generated scripts of VRML (Moore et al., 1999, Huang, 2003). VRML was unstable as each vendor implemented their own versions of browser plug-ins. Further, geographic datasets imposed a significant load on bandwidth when transmitted over the web as they are generally large in size. These and other limitations of VRML were discussed in Chapter 3. This section discusses the approaches used for enabling database connectivity from web-based GIS.

4.1.1 Standardised DBMS Connectivity

In examining database connectivity as achieved by desktop GIS, we observe that desktop GIS are generally able to access distributed Database Management Systems (DBMS). Distributed DBMS can be classified into two groups, namely homogenous and heterogeneous DBMS (Worboys and Duckham, 2004). The former is composed of multiple storage units each using the same DBMS software and data model. Similarly, the latter is composed of multiple storage units, however each using different DBMS software or different data models. To ensure interoperability the computing community developed the Open Database Connectivity (ODBC) interface which allows miscellaneous client applications to access any conformant DBMS. Figure 4.1 shows the basic architectures for accessing database servers through the ODBC, these include the two-tier and the three-tier model. The rest of this subsection discusses these architectures.

In a two-tier model the database processing task is divided into two distinct parts – the database application (such as a GIS) on the client side and the database management on the server side. In this architecture, the database connectivity component can be installed on either the client side or the server side. Chao (2006) observes that installing the database connectivity component on the client side reduces the load on the server and also the network traffic. However, installing the database connectivity component on the server side ensures that it is available to all client applications all the time. Chao (2006) suggests that a limitation of the two-tier model is that it keeps a connection to a DBMS alive, which can overwhelm a database server. We highlight that this is a limitation that depends on the software design, as connectivity API such as the Java Database Connectivity (JDBC) API allow for connections to DBMS to be opened and closed at will (Horstmann and Cornell, 2004), thereby freeing resources on the database server.

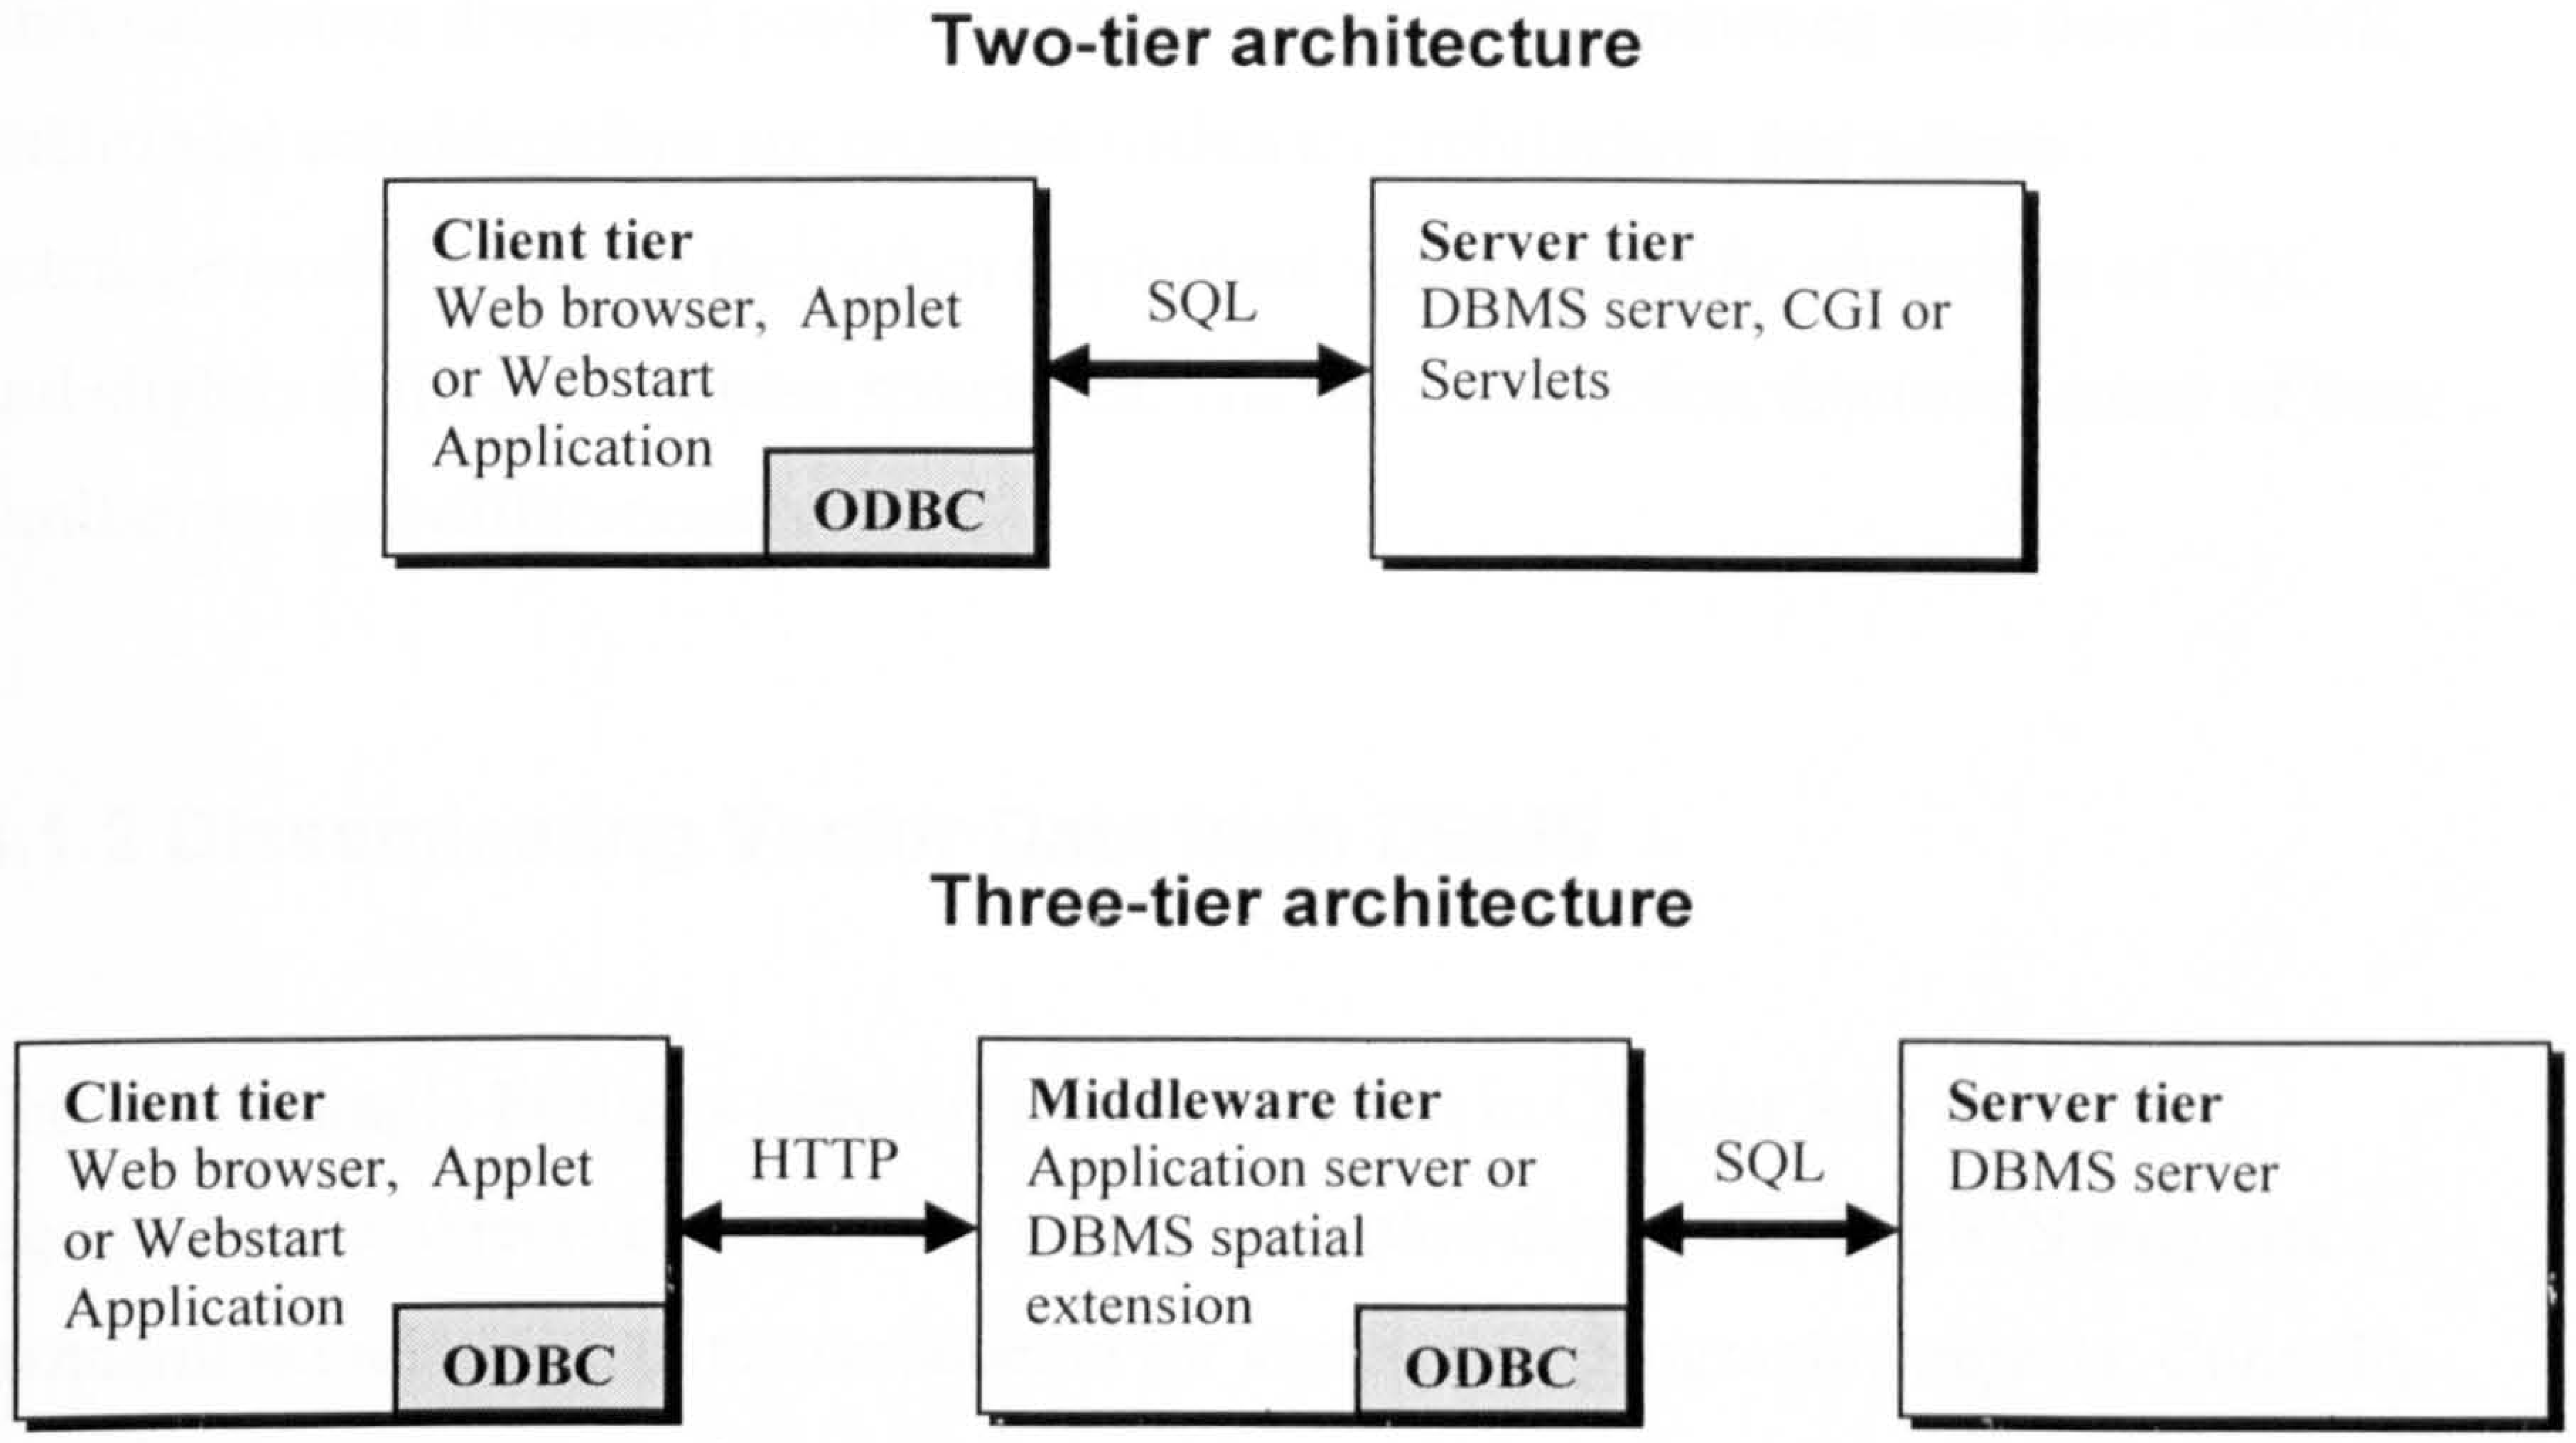

Figure 4.1 The architectures popularly adopted for database connectivity (adapted from Ghanem and Aref, 2004)

The three-tier model introduces an additional tier of middleware between the client tier and the server tier. However, it is the most popular and mature architecture for GIS development (Longley et al., 2001). We attribute this popularity to the following reason; as not all DBMS are spatially-enabled, the middleware tier allows GIS developers to implement spatial processing functionality that supports a wide range of DBMS. This is evidenced by ESRI ArcSDE, which offers uniform connectivity to Microsoft SQL Server, Oracle and several other DBMS. Another example is PostGIS, which offers spatial capabilities to the PostgreSQL database. An additional benefit of the three-tier model, is that it overcomes database limitations imposed on unsigned Java applets. The applet security model allows unsigned applets to only access data sources from the same hosting server the applet is deployed from—that is, to run from a 'sandbox'. Although this limitation can be overcome by signing an applet (Horstmann and Cornell, 2001), users might be reluctant to run signed applets as this also grants the applets complete access to the client's local files. The middleware tier, therefore, allows Java applets to access distributed database servers whilst also restricting unnecessary access to the rest of the client machine.

This subsection discussed possible architectures for disseminating data from DBMS. Additional considerations are required within an architecture that adopts heterogeneous DBMS as they often implement vendor-specific variations of SQL and slightly different database structures. The next subsection discusses some of their similarities and differences.

4.1.2 Disseminating Vector Data from DBMS

The OGC Simple Features specification, referred to in Chapter 3 as the SFS, recognises the object-relational nature of geographic databases. The SFS suggests a standardised relational database schema for abstracting geographic objects. Connolly et al., (1999) define a relational schema as the attributes of a relation and their associated domains. The previous subsection discussed architectures for disseminating data from DBMS. In this subsection we discuss how the SFS is implemented inside different DBMS. Two example relational schemas from popular spatially-enabled DBMS are discussed. The SFS schema describes four tables—a table being a physical representation of a relation. The schema, which is illustrated in Figure 4.2, is made up of the 'Geometry Columns' table that lists all geographic datasets held in the database, the 'Spatial Reference Systems' table that lists all available coordinate systems, the 'feature table' that holds the attributes of all geographic datasets and the 'geometry table' that lists the geometries of the features listed in the feature table.

In relational algebra, a relation that contains only those tuples (records) of another relation that satisfy a particular predicate is referred to as a 'selection'. For an example relation called *R*, Connolly et al., (1999) represent a selection in the following way $\sigma_{predicate}(R)$. A projection, represented as $\Pi_{field1,field2,...,fieldn}(R)$, is a vertical subset of relation *R*. That is, it contains all records from *R* but only a few of the fields. A join between two relations *R* and *S*, is an integration of records from *R* and *S* over a common field. One occurrence of the field that is common to both relations is then eliminated from the resulting relation. Connolly et al, (1999) represent it as $R \bowtie S$. The Cartesian product of two relations is a concatenation of every record in relation *R* with every record in relation *S*. It can be represented as $R \times S$. Other relational operations can be found in the literature; however the aforementioned are adequate for our discussion.

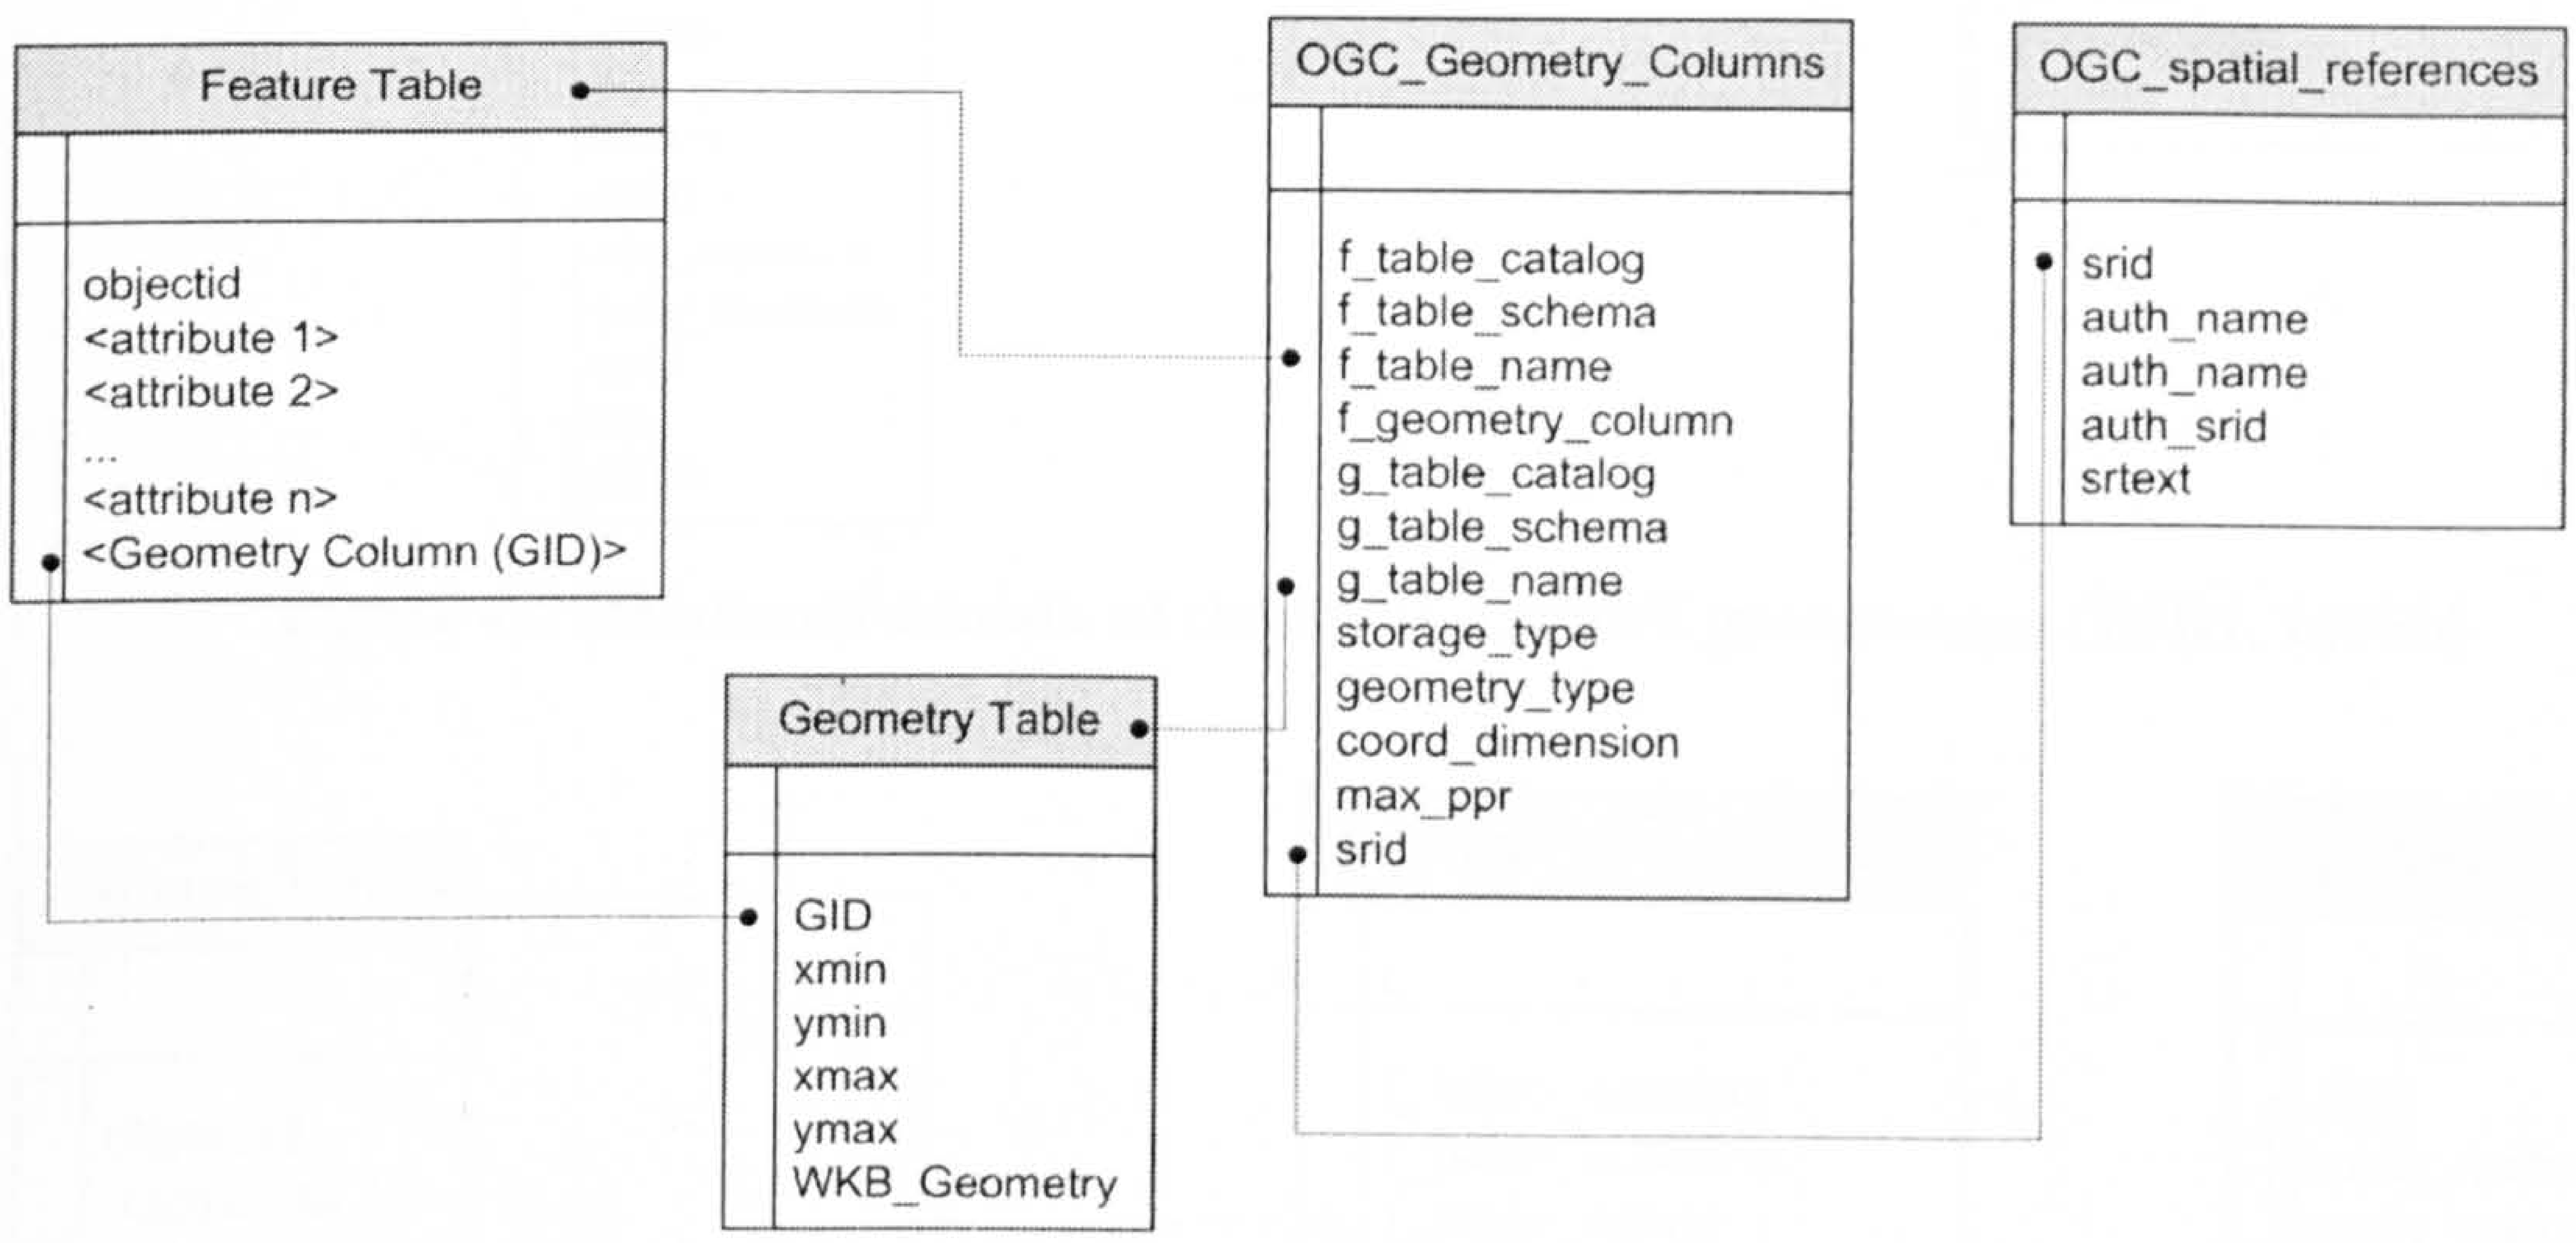

Figure 4.2 The OGC SFS Relational Schema (derived from OGC, 2005)

Where a geometry table is relation *G*, a feature table is relation *F*, a table of spatial referencing systems is relation *Z* and the geometry columns table is *C*, we observe that the collection of all geographic features (containing both attributes and geometry) in the SFS schema is represented by the join $G \bowtie F$, where the unique identifier GID is the field common to both relations. This relation can also be represented using the selection operation $\sigma_{G.GID = F.GID}(G \times F)$. Similarly, the ArcSDE database schema, illustrated in Figure 4.3, also represents a collection of all geographic features through the same natural join $G \bowtie F$. In contrast, each geographic dataset on PostGIS stores both geometry and attributes in a single relational table, as illustrated in Figure 4.4. We also observe that even though ArcSDE and PostGIS implement more fields in relation *Z* than the SFS, the relation of spatial referencing systems in the SFS can be obtained through the projection $\Pi_{srid,auth\_name,auth\_srid,srtext}(Z)$.

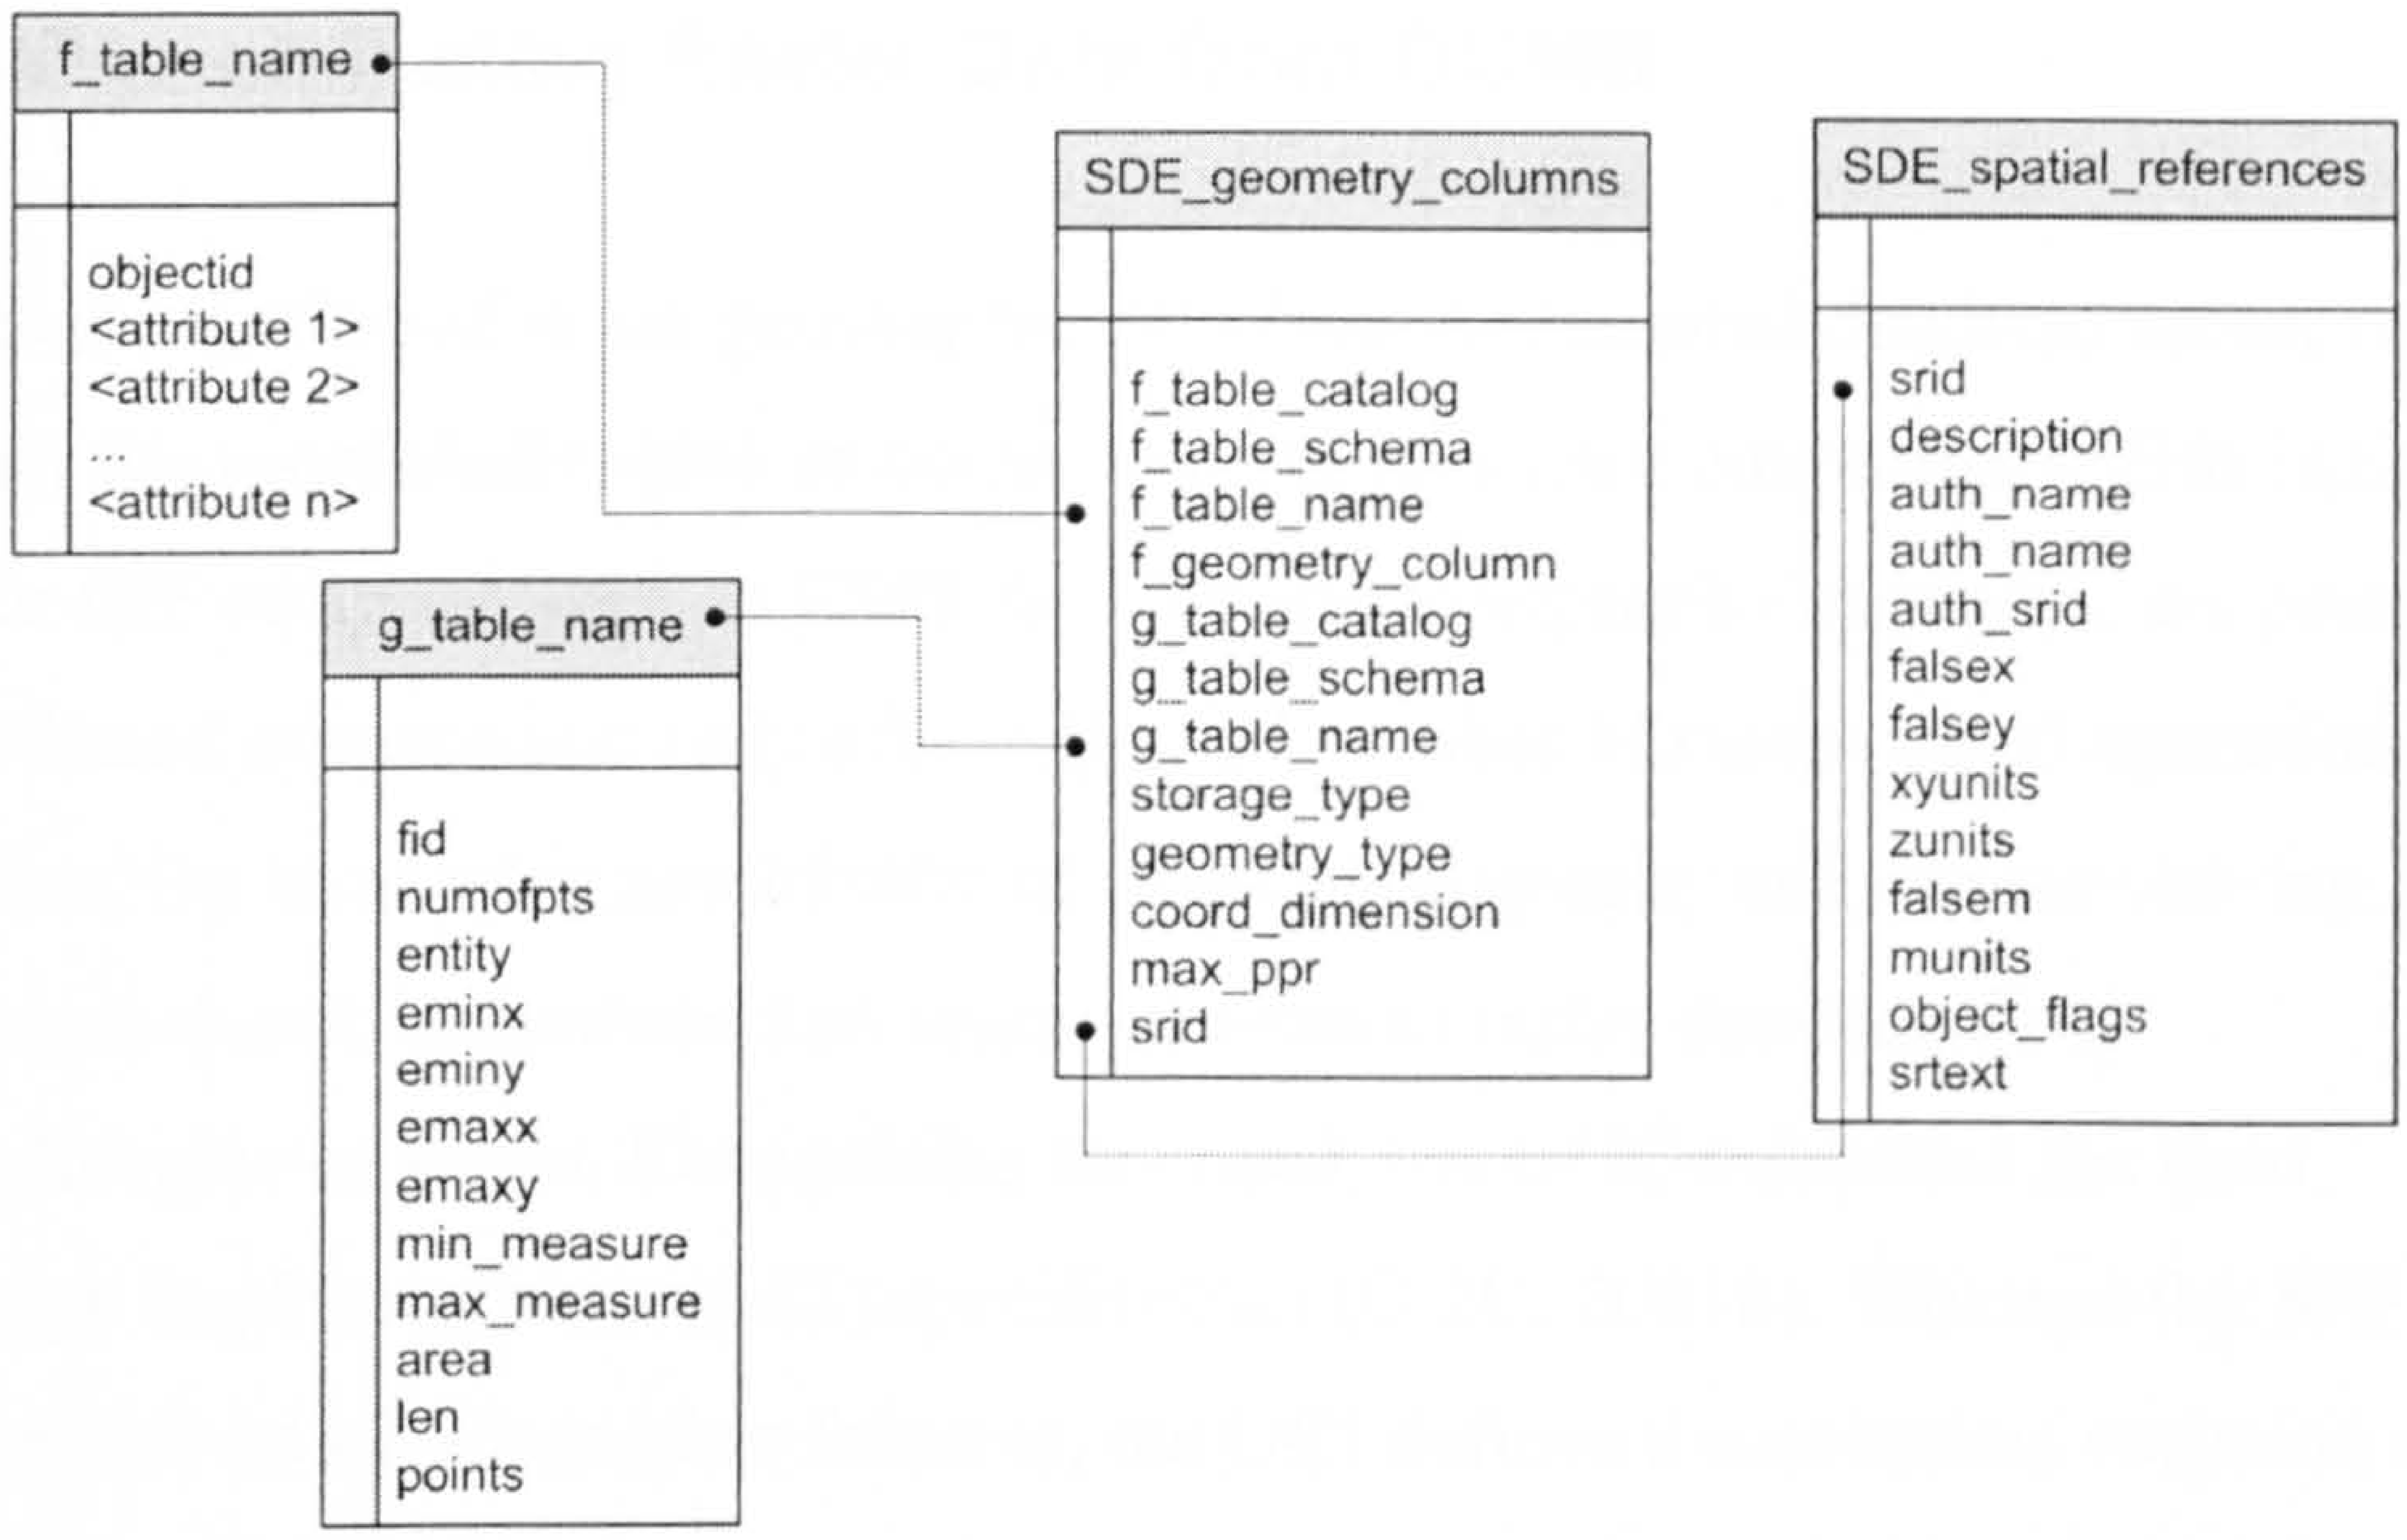

Figure 4.3 Relational schema of the ESRI ArcSDE geodatabase (ESRI, 2005b)

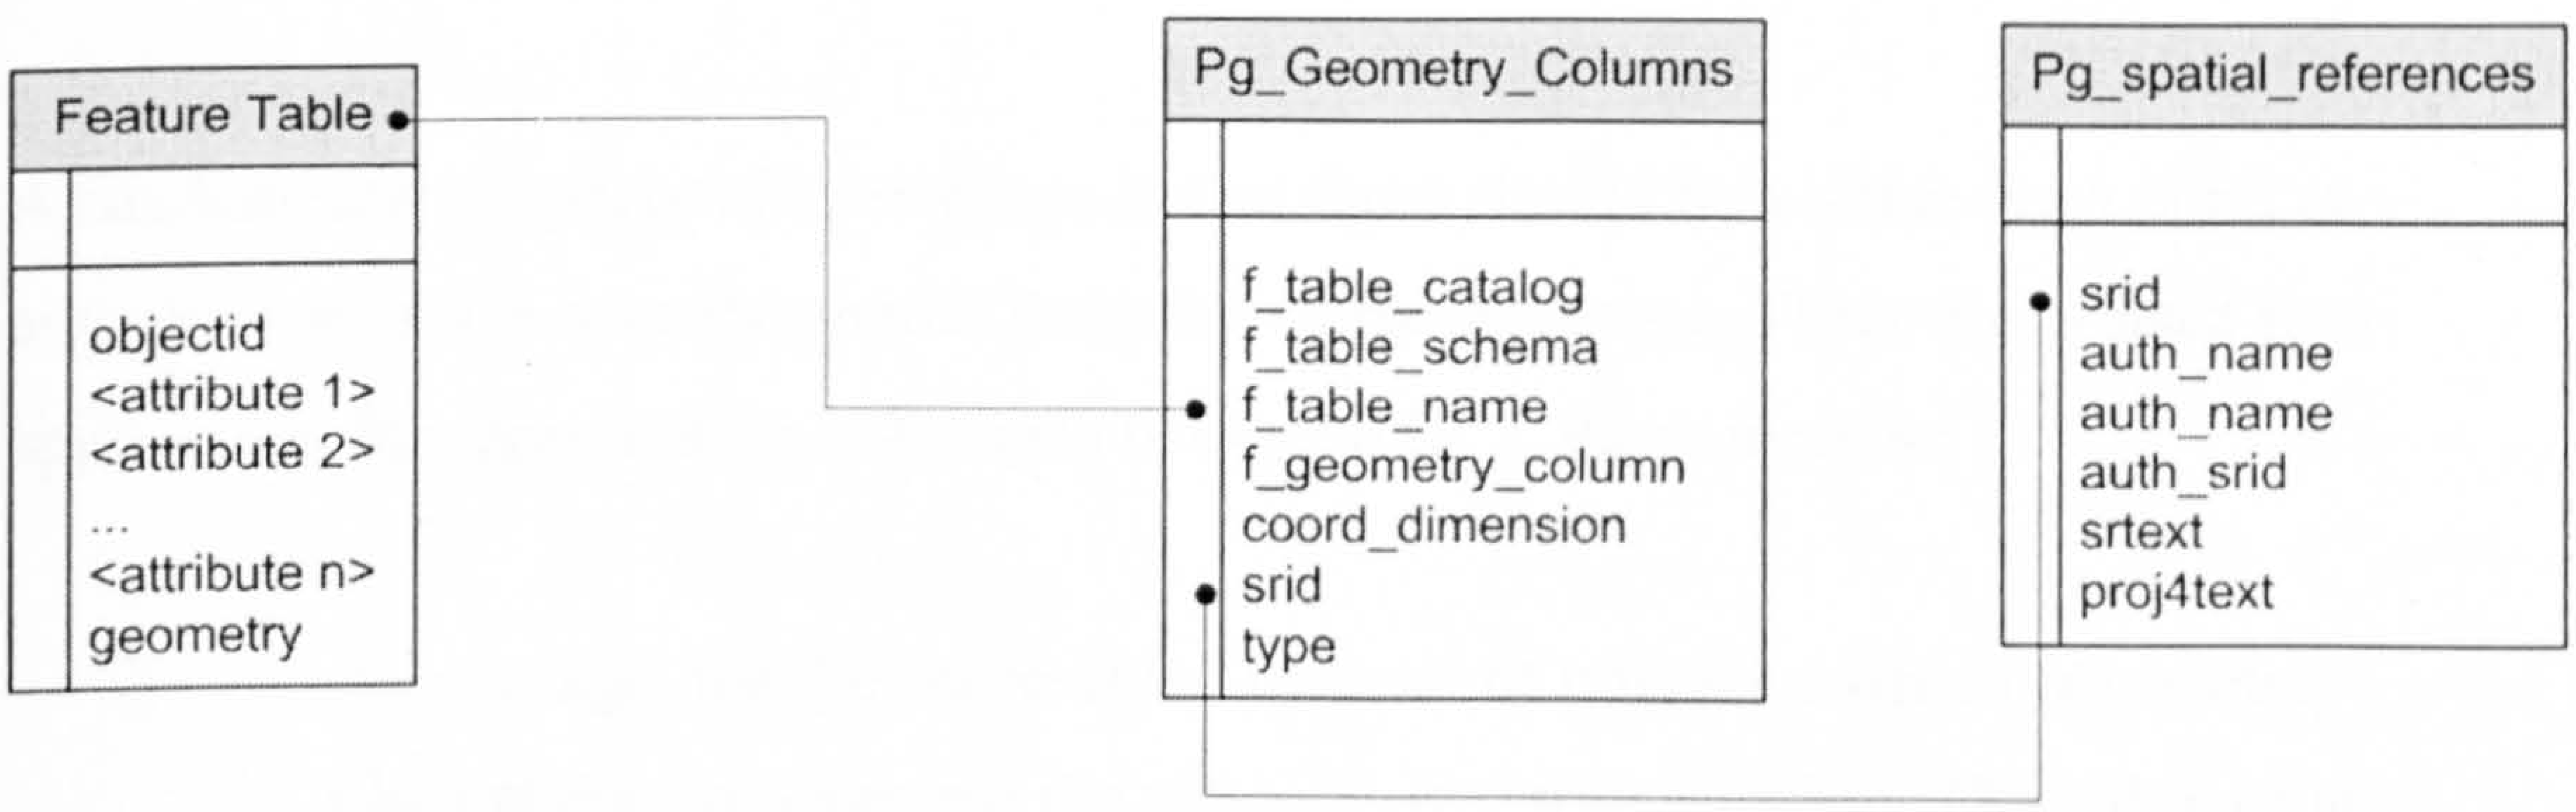

Figure 4.4 Relational schema of a PostGIS database (PostgreSQL Global Development Group, 2005)

Both ArcSDE and PostGIS implement several other tables within their schema. This subsection has only discussed those that relate to the SFS. Our discussion of these relational schemas illustrates the variation in database schema between heterogeneous spatially-enabled DBMS. These variations necessitate a middleware tier to formulate appropriate queries for each DBMS and provide an interface for client-side applications. The middleware applications therefore, have to be aware of the names and types of fields and the tables in which specific fields can be found. This subsection therefore has discussed the dissemination of vector geographic data from DBMS. The next section discusses dissemination of raster geographic data also from DBMS.

4.1.3 Disseminating Raster Data from DBMS

The dissemination of raster geographic data from a relational database server is possibly the most challenging in comparison to its vector counterpart. This is because popular GIS software such as ESRI ArcGIS and Intergraph Geomedia, are primarily vector-based systems and only offer a limited number of raster-based operations. However, the increasing availability of aerial photography and remote sensing imagery makes it imperative that improvements to raster data handling methodologies be made. Recognising this need, the OGC published the Grid Coverage Implementation (GCI) specification (OGC, 2001b). Whereas the SFS defined individual geographic features, the GCI defines a continuous variation of geographical phenomenon over a given space. The SFS and GCI therefore, derive from the vector and raster data models respectively.

A fundamental property of a coverage is that there should be a value for a point at any given location over the spatial footprint of the coverage. There are at least four approaches for representing coverages using the GCI, these include:

i) A coverage may be represented by a set of polygons which completely tile a plane, for example the triangles in a triangulation. The value of an attribute at any point on the raster is obtained from the polygon that contains that point.

ii) A coverage may also be represented by a grid of values such that the value of returned by a grid for a point is the grid value that is closest to that point. This is the conventional approach adopted by several GIS.

iii) Alternatively, a coverage may be represented by a mathematical function such that the value returned by the coverage for a location is the range of output values for that function.

iv) A coverage can also be represented by any combination of the aforementioned approaches; for example, where the attribute value of the polygons in the first approach is a mathematical function.

Due to these highly variable means of encoding raster data, the GCI does not attempt to formalize a relational model for storing raster data. However, other studies have suggested possible relational schemas and approaches for storing geographic raster data. For example, ArcSDE cuts a raster dataset into smaller subsets, referred to as "blocks", and stores them in individual rows in a separate block table (ESRI, 2005a). The blocks are stored as several small binary large objects (BLOBs). Each block is referenced in the *SDE_raster_columns* table which is also used for referencing the different bands that may comprise a raster. The aggregating of raster datasets into smaller blocks also allows the creation and storage of lower resolution versions of the raster dataset, referred to as pyramids. The pyramids are created based on a resampling algorithm of the user's choice for example nearest neighbour, convolution filters or bilinear interpolation. This means that rasters of the same pyramid have the same geographic extent but differ in resolution. Figure 4.5 illustrates the creation of multi-resolution rasters using the pyramid concept. Figure 4.6 shows the relational schema for ArcSDE raster datasets.

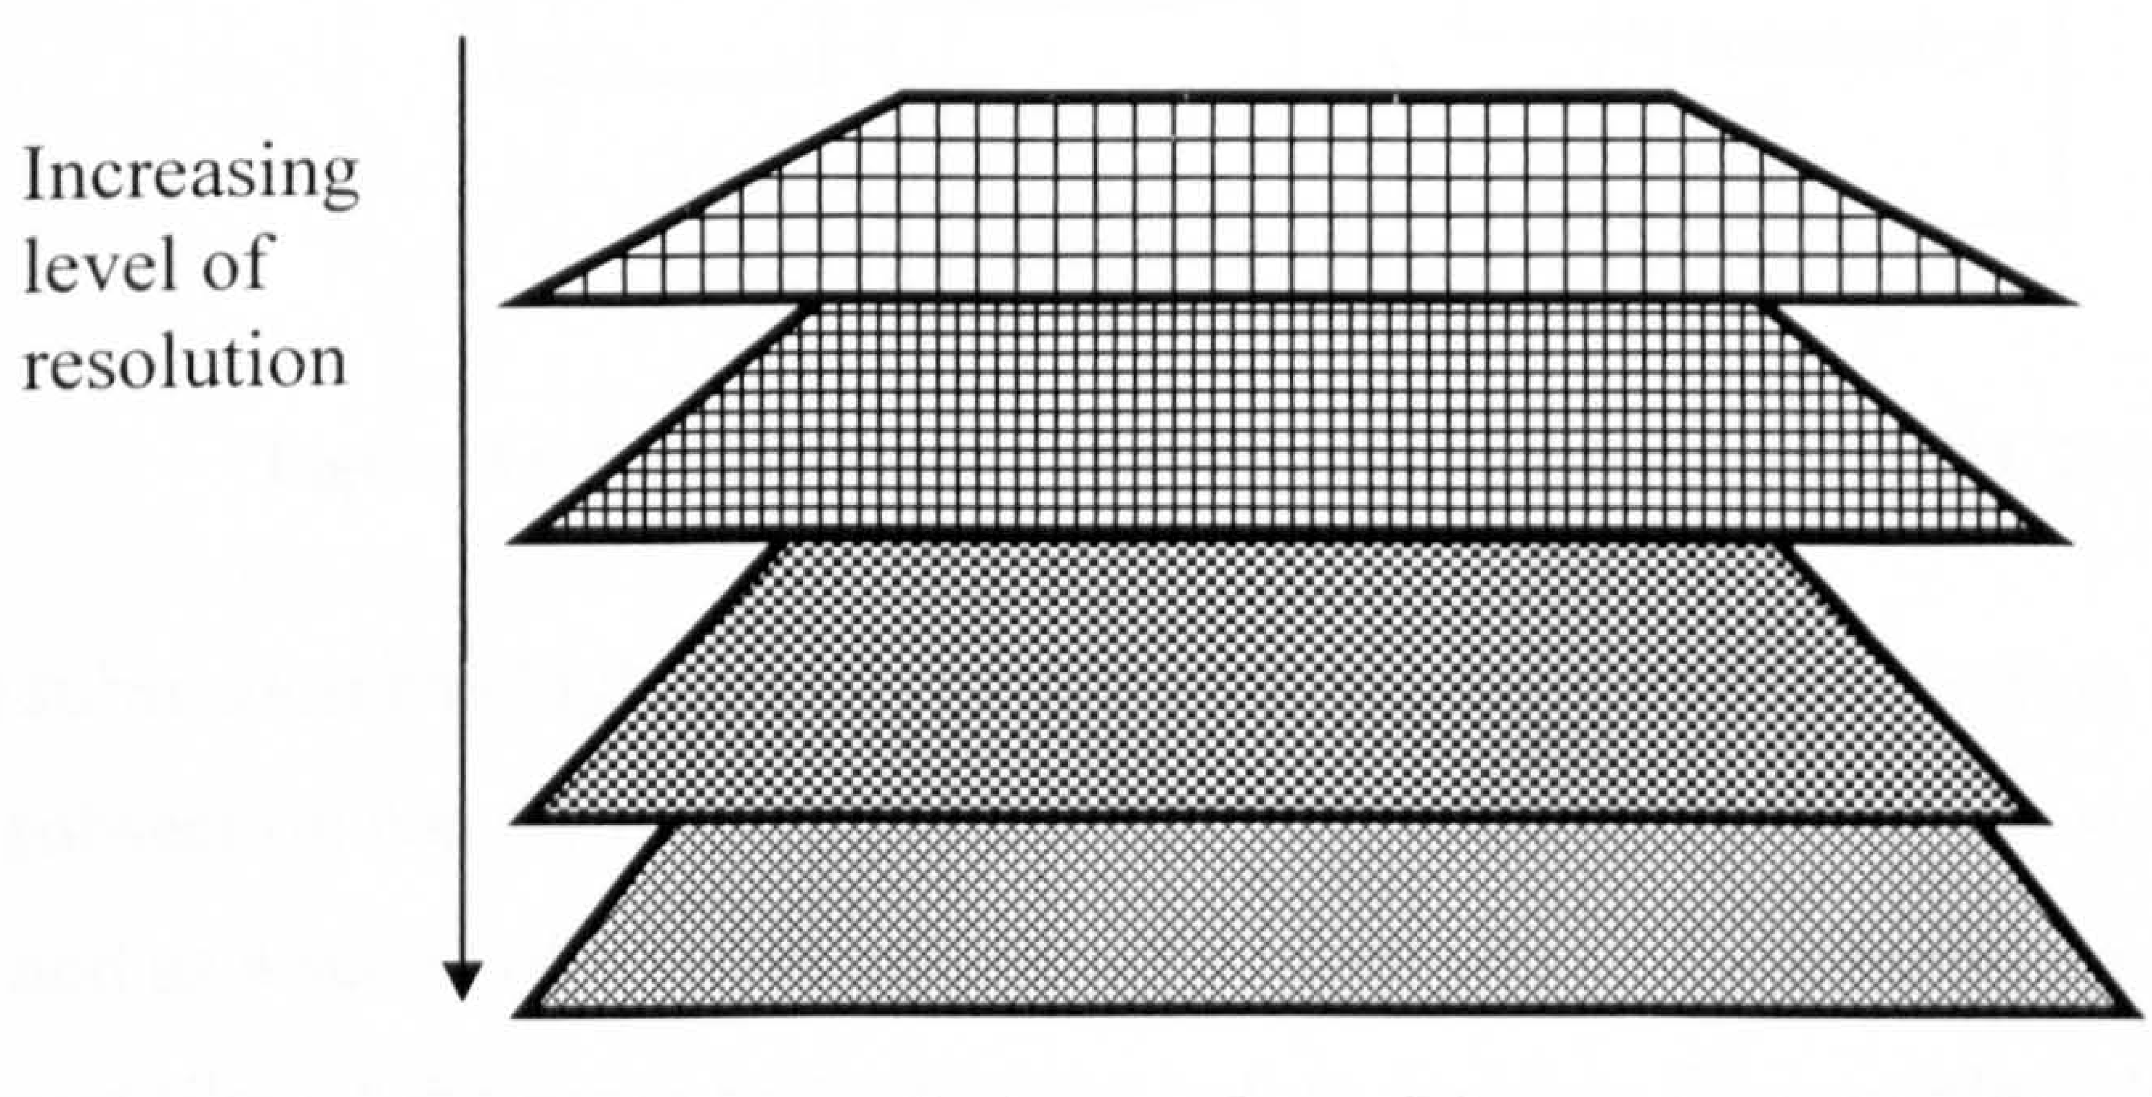

Figure 4.5 A pyramid of smaller resolution versions of the same raster

Of particular importance is the *SDE_blk_n* table which stores the actual pixel values as BLOBs. From the schema, it can be observed that the table holds information of the location of the block within each level of the multi-resolution pyramid (*row_nbr, col_nbr*). The table also identifies the band to which the block belongs (*rasterband_id*). The table therefore defines a four dimensional matrix with the band, pyramid level, row and columns as dimensions. Consequently, each block obtains its spatial referencing from its position within the matrix of blocks. The matrix which represents the complete raster dataset is spatially referenced by a polygon feature stored in a geometry field with other thematic (business) fields that a user may wish to implement. The raster dataset is delivered over the internet as a single raster after the ArcSDE engine merges all the aggregate blocks into a single dataset.

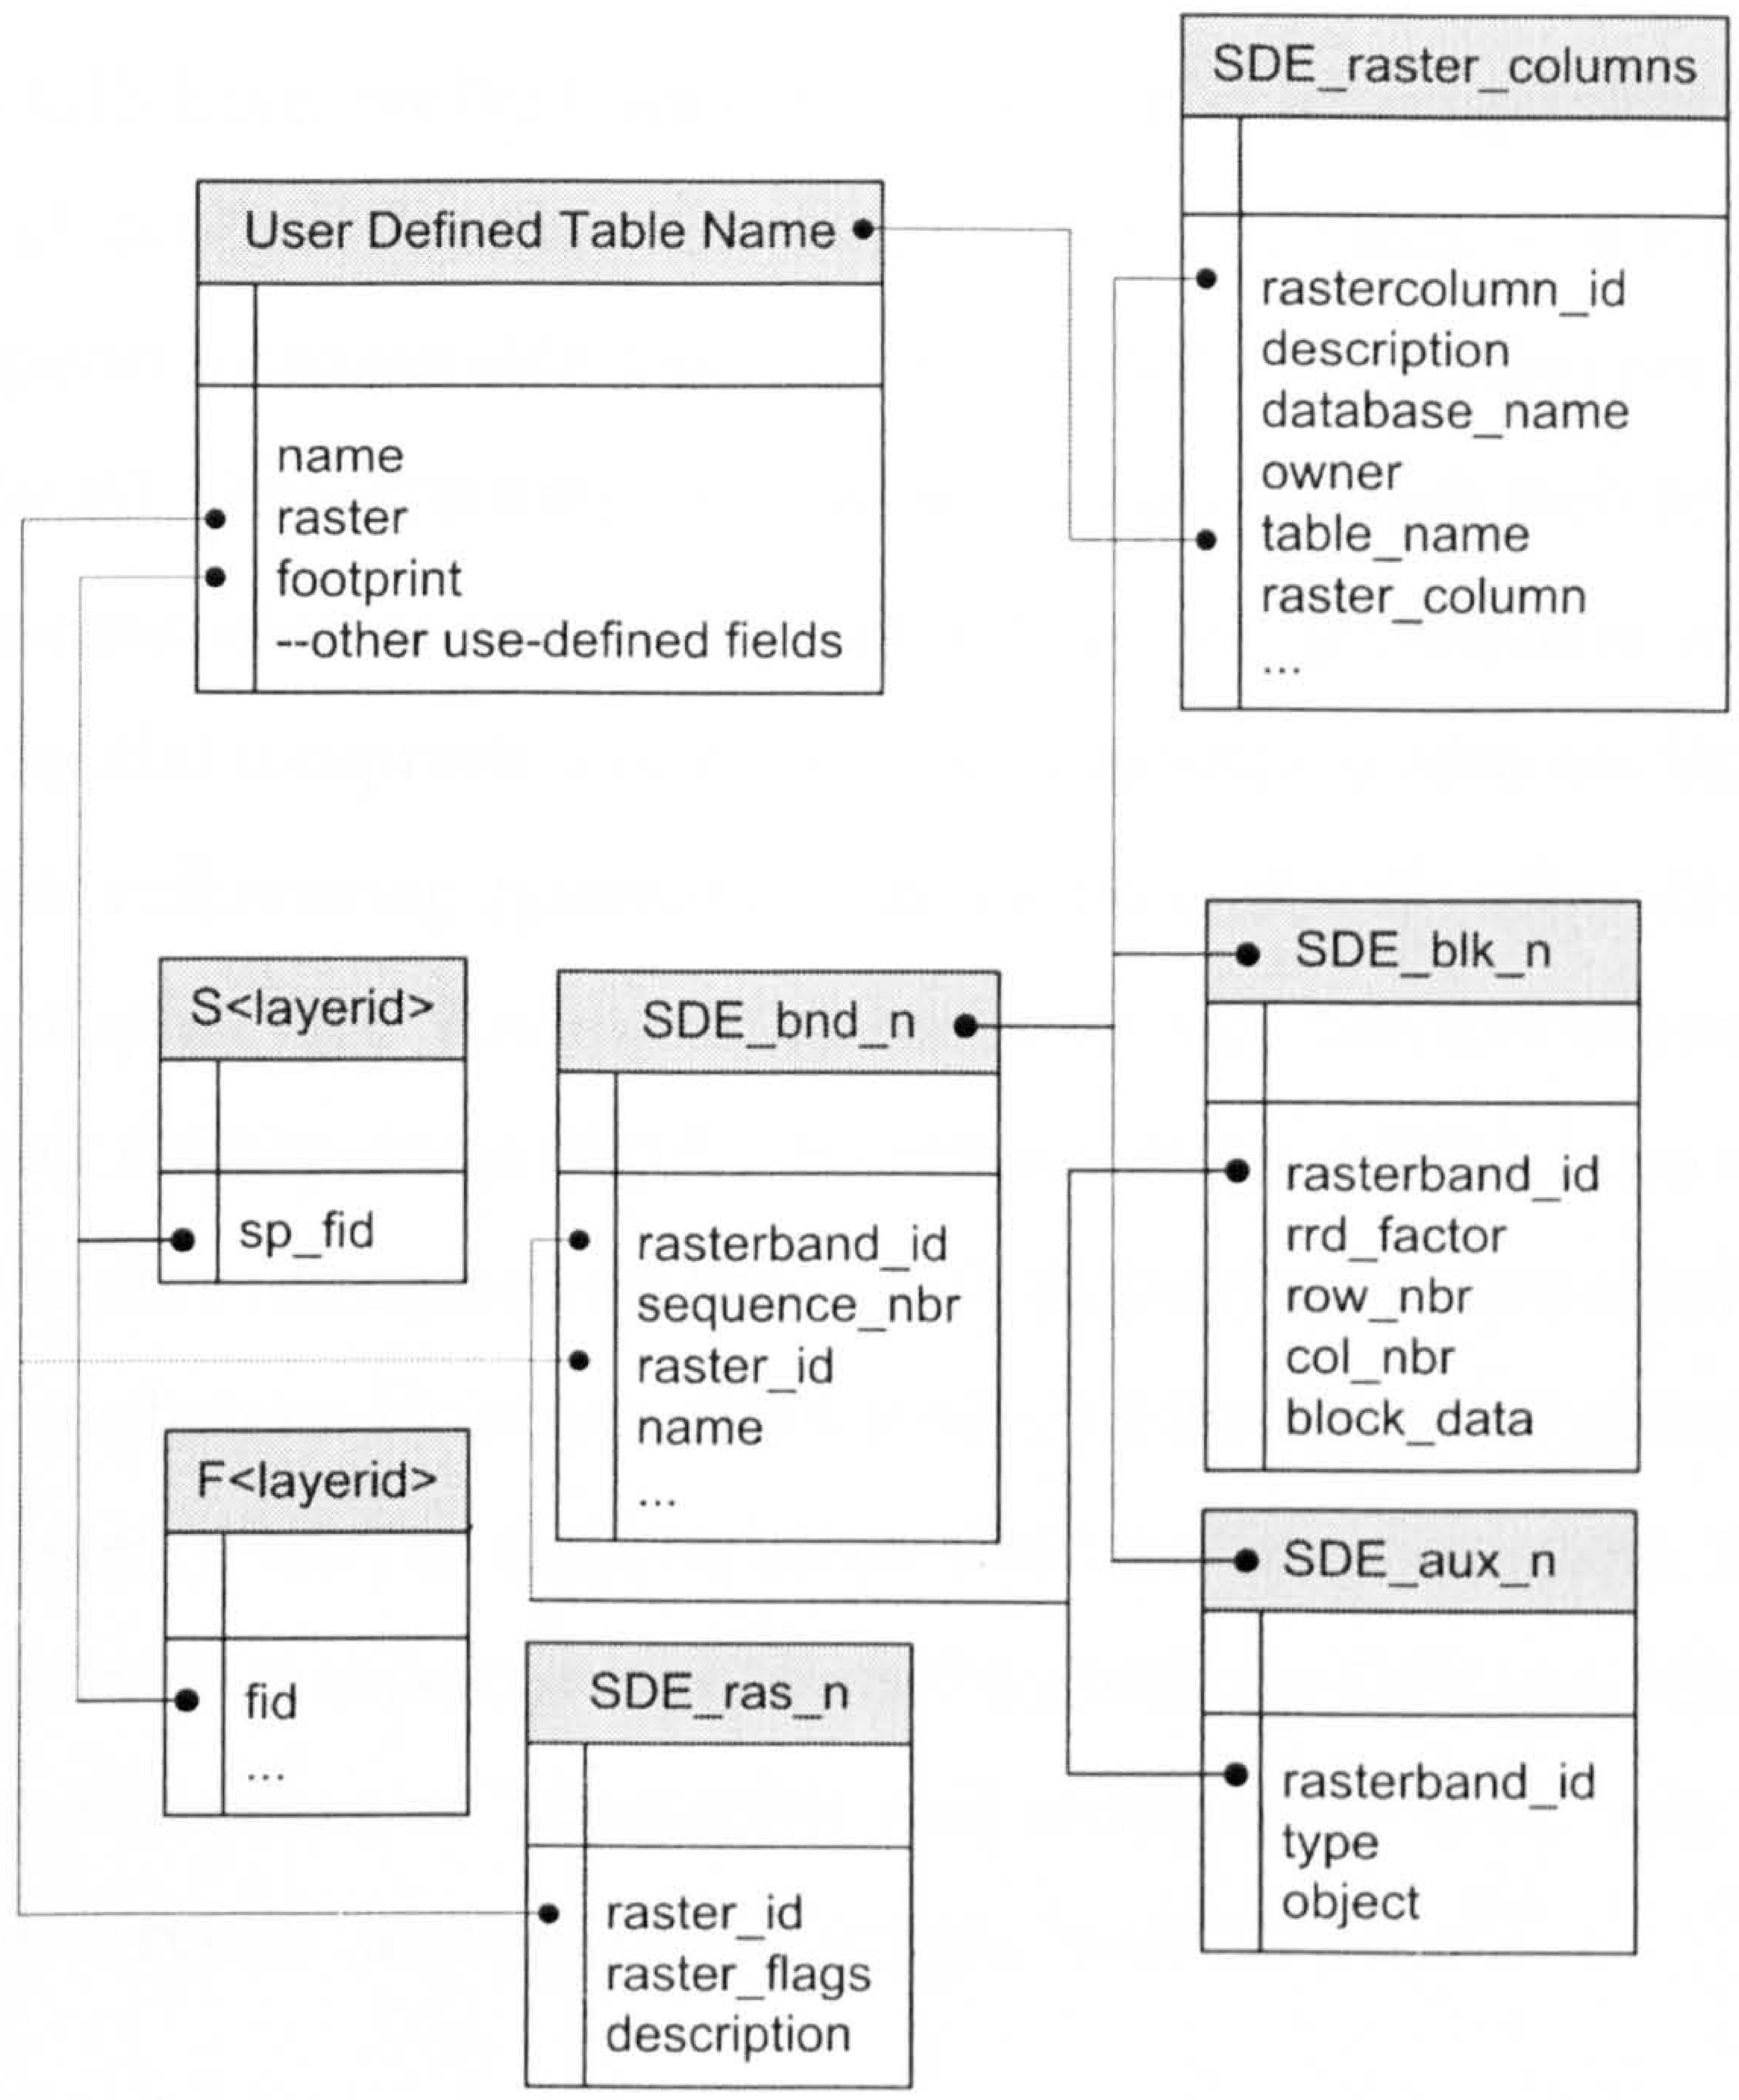

Figure 4.6 ArcSDE raster table schema (adapted from ESRI, 2005a)

This subsection has highlighted the different approaches for encoding raster data. The subsection has also presented the ArcSDE relational model for storing raster data and as a series of tiles. A middleware layer allows for a mediator to integrate different tiles of the same raster dataset before sending the complete dataset to the client. Where a raster dataset is to be derived from a mathematical model, the raster dataset could be created by the mediator before sending to the client. The three-tier model is therefore a possibly appropriate architecture for providing database access to a web-based geovisualisation application. This section, in whole, has discussed the implementation of database servers and the schemas of relational tables adopted by the GIS community for disseminating geographic data. The next section discusses an alternative approach for disseminating geographic data using web services, which may also be regarded as mediators.

4.2 Dissemination through Web Services

Contemporary GIS have evolved against a backdrop of developments in the dissemination of data and functionality through Web Services — a software system designed to support interoperable machine-to-machine interaction over a network (W3C, 1999-2006). Dissemination of geographic data through web services is different to dissemination of static geographic files as i) geographic web services have dynamic spatial footprints ii) can be served in various formats iii) can be served in various spatial referencing systems iv) can be merged with other datasets or layers from the same map service. Architectures based on web services are referred to as Service-Oriented Architectures (SOA). In its basic form, a SOA is an example of a three-tier architecture (discussed earlier in this chapter) with web services acting as middleware or mediators. However, SOA extended the three-tier paradigm into a multiple tier model (*n*-tier) as shall be illustrated later in this section. This section discusses web service technologies published by both the OGC and the W3C, we refer to OGC web services as 'geographic web services', to W3C web services as 'traditional web services' and to both as simply 'web services'.

4.2.1 Traditional Web Services

Traditional web services are built on technologies that can be divided into three main groups: communication protocols, service descriptions and service discovery. Communication protocols allow for the exchange of information between applications. An example of these is the Simple Object Access Protocol (SOAP). Service descriptions provide information on how to use a web service, that is, what functions a web service performs and the types of messages it receives and sends. The most popular method of describing a service is through the Web Service Description Language (WSDL). Service discovery technologies offer a catalogue of available services and their publishers. The current standard for service discovery is the Universal Description, Discovery and Integration (UDDI). Applications implementing these web service standards assume three types of roles, namely service requester, service provider and service registry. The applications interoperate to publish, find and bind clients to available relevant services, as illustrated in Figure 4.7. The rest of this section discusses the aforementioned web service technologies; an extensive discussion of other web service technologies can be found in Weerawarana et al.(2005).

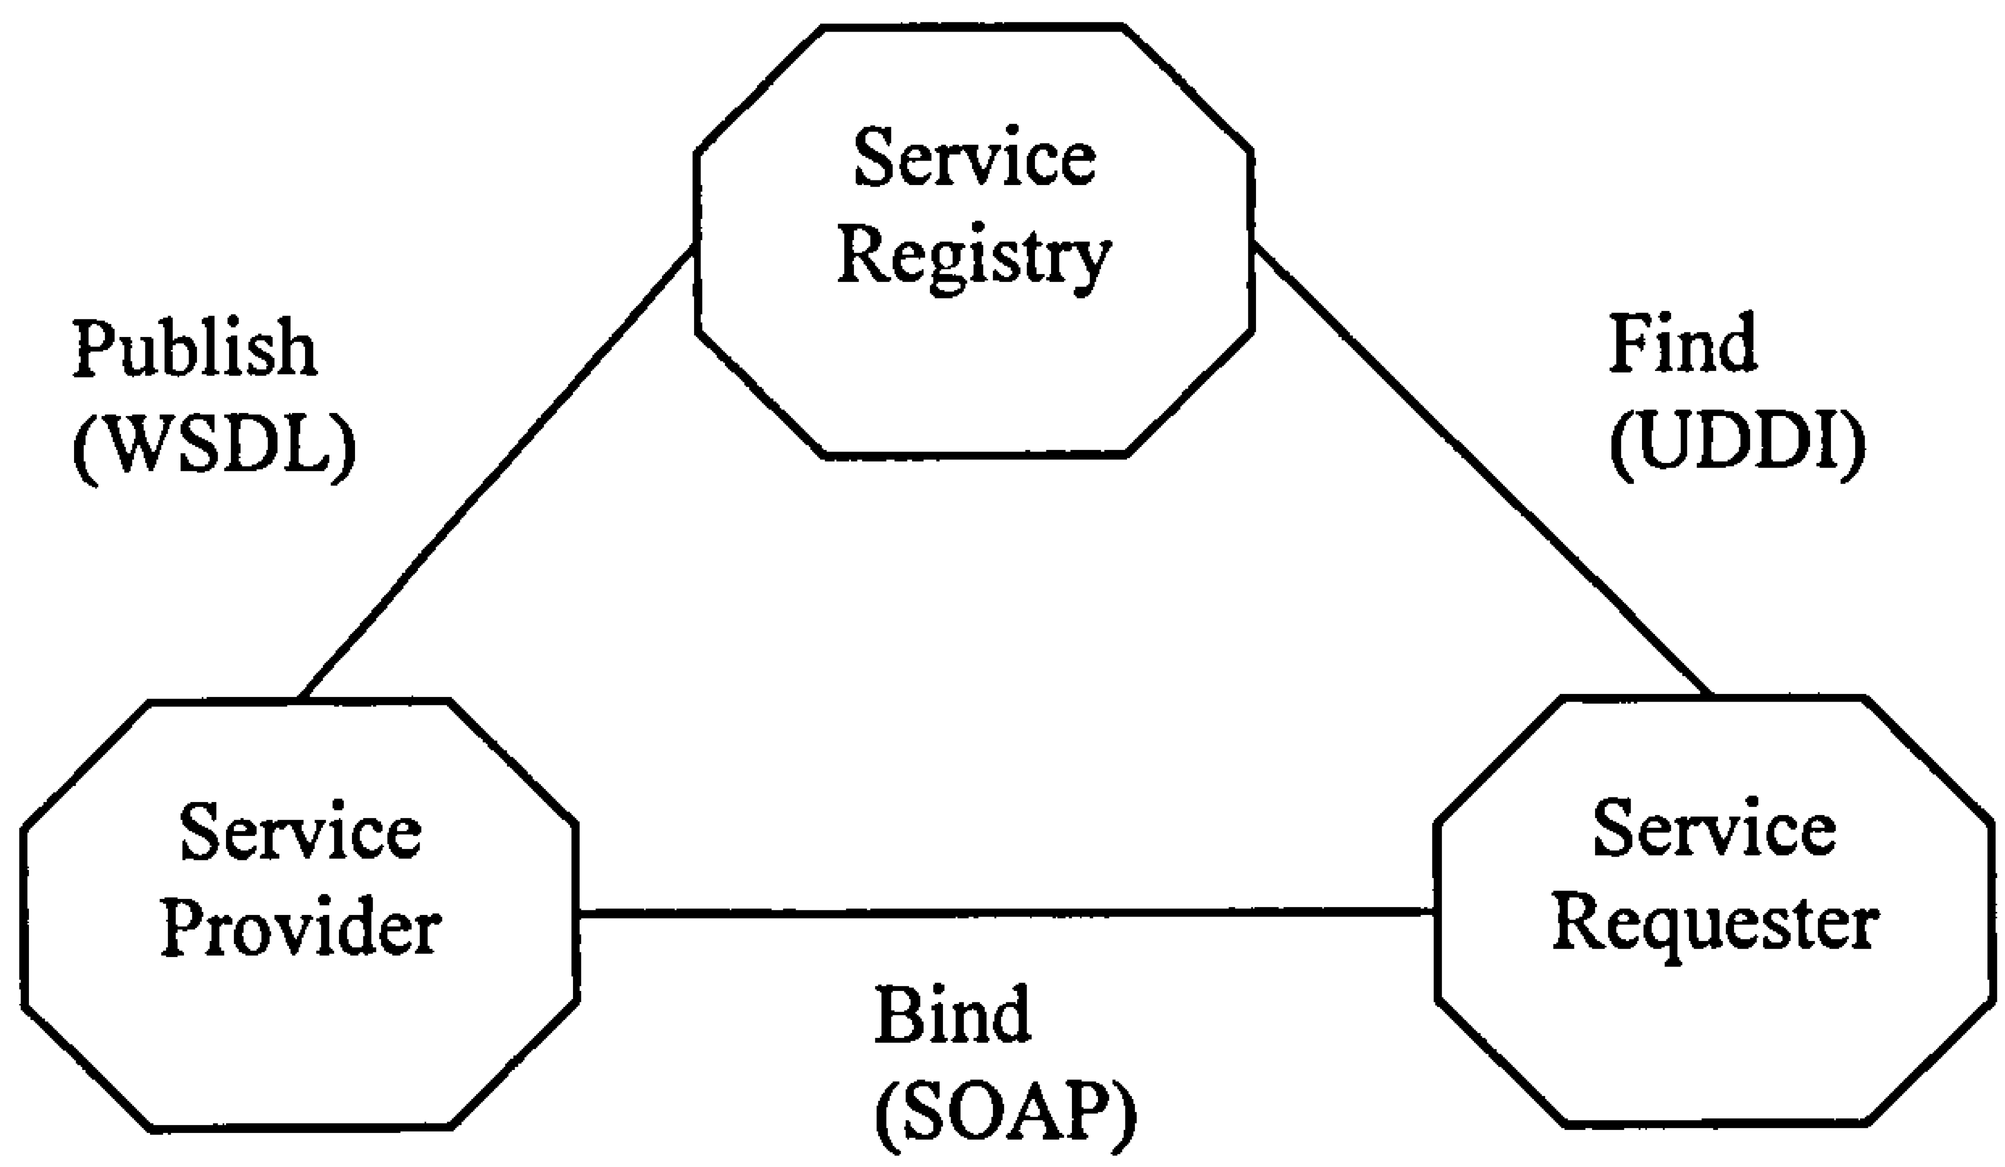

Figure 4.7 Web Services architectural model (adapted from Huhns and Singh, 2005)

SOAP is a platform-independent protocol for exchanging structured information in a decentralised and distributed environment (W3C, 1999-2006). SOAP messages are encoded independently of the transmission protocol and therefore can be transmitted via HTTP, Secure HTTP (HTTPS) or any other proprietary transmission protocol. In its simplest form, a SOAP message is made up of three parts, a header, an envelope and a body. From Figure 4.8, which shows a simple SOAP message, the header is shown by lines 5 to 14 and the body by lines 15 to 22. The envelope, which is primarily meant for conformance with the XML specification, is illustrated by lines 2 to 4 and line 23 in Figure 4.8. It should be noted, however, that the XML elements are not restricted to the ones illustrated in the Figure. They can adopt any name and therefore make it possible to represent most structured data models including those of object-oriented GIS.

The principal method of describing a web service is through its WSDL specification. If we consider a web service to be a process which accepts input parameters and outputs some form of response, then a WSDL is an XML-based language for describing the name of a service, the operations offered and the data types of the input and output parameters. The specification supports any data type system regardless of the platform used by the service requestor or service provider. This is achieved by referencing data types to namespaces (qualifiers) encoded as Uniform Resource Identifiers (URI). For example, the integer and string data types defined within the W3C XML schema (W3C, 1999-2006) are supported by most web service engines. Application-centric data types can also be supported by WSDL through the referencing of unique namespaces.

```
<?xml version="1.0" ?>
<env:Envelope
xmlns:env="http://www.w3.org/2003/05/soap-envelope"
xmlns:xml="http://www.w3.org/XML/1998/namespace">
<env:Header>
<env:Upgrade>
<env:SupportedEnvelope qname="ns1:Envelope"
xmlns:ns1="http://www.w3.org/2003/05/soap-
envelope"/>
<env:SupportedEnvelope qname="ns2:Envelope"
xmlns:ns2="http://schemas.xmlsoap.org/soap/
envelope/"/>
</env:Upgrade>
</env:Header>
<env:Body>
<env:Fault>
<env:Code><env:Value>env:VersionMismatch
</env:Value></env:Code>
<env:Reason><env:Text xml:lang="en">Version
 Mismatch</env:Text></env:Reason>
</env:Fault>
</env:Body>
</env:Envelope>
```

Figure 4.8 An example of a SOAP message (W3C, 1999-2006)

Development of the UDDI specification is spearheaded by the Organization for the Advancement of Structured Information Standards (OASIS). Although OASIS is a separate entity to the W3C, the two organisations recognise each other's work and therefore build on each other's standards; for example, UDDI Application Programmers' Interfaces use SOAP (a W3C standard) for messaging. UDDI defines a set of services for the description and discovery of web services providers (such as businesses or governmental organisations), the web services they provide and the interfaces used for accessing the services. Although the current version of UDDI can represent Spatial Referencing Systems and degrees of latitude and longitude through specification of technical models (tModels), a major obstacle to the adoption of UDDI within the GI community has been that UDDI registries do not support spatial queries, which are a fundamental aspect of geographic data handling (Alameh, 2003).

4.2.2 Geographic Web Services

Some of the specifications by the OGC include the Catalogue Services Specification (CAT), Web Map Service (WMS), Web Feature Service (WFS) and the Web Coverage Service (WCS). They all fall under an umbrella specification referred to as the OGC Web Service specification (OWS). The OWS defines features that are common amongst all OGC web service specifications. Much of the consensus on OWS has been a direct result of earlier interoperability work on geographic data models. One of the key products of the geographic interoperability initiatives of the OGC has been the development of the Geography Markup Language (GML) as a data exchange format. This section starts with a discussion of GML because it is a key component of requests to and from geographic web services.

```
<needs:geofeature uniqueid="123456">
<needs:owner>John Smith</needs:owner>
<needs:polygon>
<gml:Polygon srsName="osgb:BNG">
<gml:outerBoundaryIs>
<gml:LinearRing>
<gml:coordinates>424745.71875,564970.75
424743.9375,564972.625
424746.0625,564974.875
424745.71875,564970.75</gml:coordinates>
</gml:LinearRing>
</gml:outerBoundaryIs>
</gml:Polygon>
</needs:polygon>
</needs:geofeature>
```

Figure 4.9 An extract from a GML file

GML is a platform-independent language (protocol) for exchanging geographic information (Lake et al., 2004); therefore it recognises spatial referencing systems and geometry. Similar to SOAP, GML is also serialised in XML. The standard formalises the elements describing geometry into, for example *<gml:point>*,

<gml:LineString> and *<gml:polygon>* as illustrated in Figure 4.9. Unsurprisingly, the geometric elements of a GML file are mandatory, however other fields may vary depending on the particular profile of GML. As a result, different organisations have implemented their own profiles of GML for example the Ordnance Survey offers OS MasterMap and the Dutch National Mapping Agency (Topografische Dienst Nederland) offers Top10NL. Software vendors have similarly developed vendor-specific profiles of GML for example Safe Software's FME Workbench. The main advantage of GML over SOAP for geographic data dissemination is that across all GML profiles the geometry and spatial referencing systems are encoded using exactly the same notation and convention (schema). This significantly improves interoperability between different systems and helps ensure data integrity. In contrast, SOAP is not bound to any object structure and hence structures (classes) encoded in SOAP may differ from vendor to vendor.

The WMS specification defines a web service that renders and returns dynamic maps as digital images. The maps are transmitted, as streams of bytes, in image formats such as PNG, GIF or JPEG. This allows the receiving browser to render them as conventional images. The service therefore does not provide geographic data itself; instead it provides a pictorial view of the data layers. The input parameters for WMS include a reference to a style descriptor (for colours and symbols), the geographic extent, the spatial referencing system, the output format and the list of layers to render onto the produced image. If the layers defined in the request have differing SRS, the server projects the layers into the SRS defined in the request. The specification adopts at least three basic services, the `getCapabililites`, `getMap` and `getFeatureInfo`. The `getCapabilities` operation returns a list of all available layers including their geographic extents and spatial reference systems. The default style for each layer is also mentioned. Additional metadata is included for describing the WMS including the list of all image file types that can be returned by the service. The return from the operation is serialised in XML, making it human readable. The `getMap` operation returns a digital image of the requested layers using rendering styles defined in the `getMap` request. Last, the `getFeatureInfo` operation allows attributes to be retrieved from layers that allow querying.

The WFS specification defines an HTTP-based distributed computing platform for the *insert, update, delete, lock, query* and *discovery* of geographic objects. Unlike the WMS, WFS return the actual content of the geographic data. It thus returns a collection of geographic objects serialised in GML, as a minimum. Alternatively, legacy geospatial formats such as the ESRI shapefile or Autodesk DXF may also be retrieved from a WFS. Description of geographic objects within the interface is however, restricted to GML. This restriction is further enforced by the rule that the interfaces must be in eXtensible Markup Language (XML). The WFS adopts a `getCapabilities` operation for describing the layers and operations available on the service. The `getCapabilities` operation of a WFS describes which feature types (layers) are available through the service, and which operations are supported on each feature type. The `getCapabilities` operation, therefore, offers similar metadata to that offered by the WSDL descriptions. However, it also formalizes specification of geographic extents and spatial operations such as 'intersection' and 'containment' that can be executed on available datasets. Formalizing geographic operations offers a higher-level of interoperability than offered by traditional web services, as they do not offer an equivalent facility.

The OGC defines the WCS as a specification for supporting electronic interchange of geospatial datasets representing space-varying phenomena (coverage). Some of the geographic data models regarded as being of coverage type include raster models and Digital Elevation Models (DEM). These are generally encoded as a georeferenced regular grid of pixel values; however, coverages can be returned as image files. The specification defines three operations `getCapabilities`, `getCoverage` and `describeCoverage`. Similar to both the WMS and the WFS, the `getCapabilities` method provides metadata about the available layers and their spatial reference systems. An important distinction between WMS and WCS is that the latter can return either an image or the raw data itself. Images are generally restricted to Red-Green-Blue (RGB) values ranging from 0 to 255 whereas raw data may include pixel values greater than 255 or lower than zero. However, images offer the possibility of returning several image bands stored in a single image file. A portrayal service, as introduced in Chapter 1, is then required to render coverages from a WCS before presentation on a client. Portrayal services are therefore, normally embedded within WCS, WMS or both.

Within a geographic data discovery framework, the challenge of discovering a suitable web service is a major area of study. To address this challenge, the OGC published the Catalog Services specification (CAT) for disseminating geographic metadata encoded in the ISO19115. The specification allows for implementing different interfaces for communication between client and server for example through HTTP, CORBA or Z39.50 protocols. Web services adopting the HTTP binding are often referred to as Catalogue Services for the Web (CSW). Nouegras-Iso et al., (2005: pp. 202) observes that "despite the relevance of catalog interface specifications, the implementations are not many in comparison with other OGC specifications". They attribute the slow uptake of CAT to the higher attention afforded to web mapping by the OGC. In Chapter 1, we examined metadata dissemination through the Z39.50 and presented results on a survey we carried out to determine the popularity of different Z39.50 servers. We can therefore add to the observation by Nouegras-Iso et al., (2005) by highlighting that it is apparent that geographic metadata catalogue servers have been implemented, however there is no indication that they have been updated to conform to the CAT or CSW specifications. The specification enforces interoperability by offering an extensible Common Catalogue Query Language (CQL) that is modelled on the SQL WHERE clause or relational algebra's predicates of selection operations. The CQL supports both tight and loose queries: where a tight query is defined as a query that returns a null set if a requested field is not supported; and a loose query is where any undefined field is considered to match the queries predicates. We acknowledge CAT as a possible geographic metadata delivery mechanism for the STORM browser that was proposed in Chapter 3.

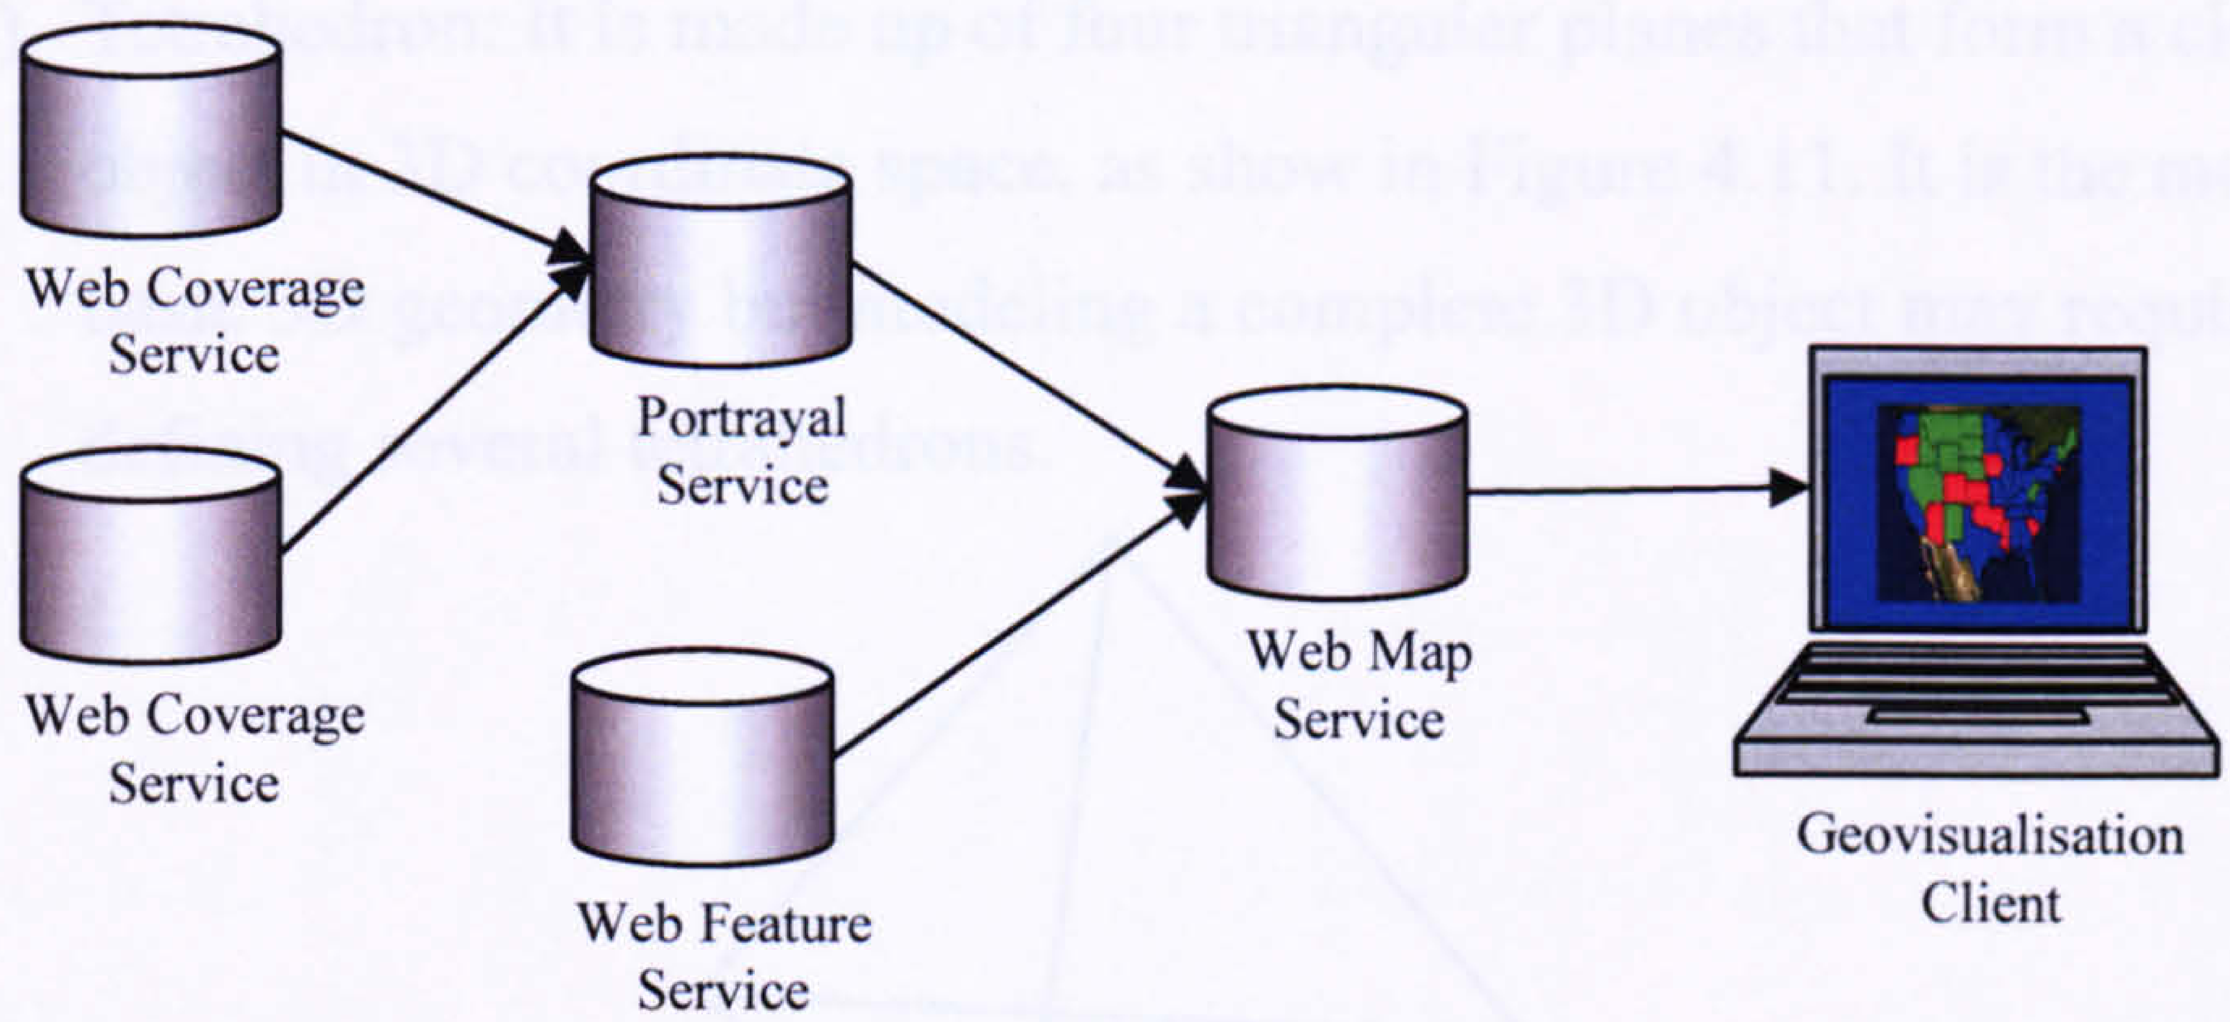

Figure 4.10 Chaining of geographic web services (adapted from Alameh, 2003)

Figure 4.10 illustrates how these geographic web services interoperate within a geo-centric SOA. It can be observed from the illustration that a geo-centric SOA can comprise of several tiers. The section has shown that geographic web services offer a high level of interoperability by formalizing parameters for requests, for example the bounding box and spatial reference system parameters for WMS requests. Although this is an advantage over the more traditional web service specifications, this means that traditional web services offer more dynamism through the flexible structure of SOAP objects. Consequently, this thesis suggests that it is advantageous for web-based geovisualisation applications to support both traditional and geographic web services. This section has examined both traditional and geographic web services technologies; the next section examines the dissemination of multidimensional geographic objects.

4.3 Delivering Multidimensional Objects

The previous chapter discussed different approaches in the visualisation of multidimensional geographic information. This section discusses the storage and dissemination of multidimensional geographic objects using the data delivery methods discussed earlier in this chapter. In particular, the section discusses the encoding of 3D geographic objects within an RDBMS and the serving of 3D geographic objects from a SOA. First, we examine 3D geometric primitives suggested by Arens et al, (2005) for modeling geographic objects:

a) Tetrahedron: It is made up of four triangular planes that form a closed object in 3D coordinate space, as show in Figure 4.11. It is the most basic 3D geometry but modeling a complete 3D object may require defining several tetrahedrons.

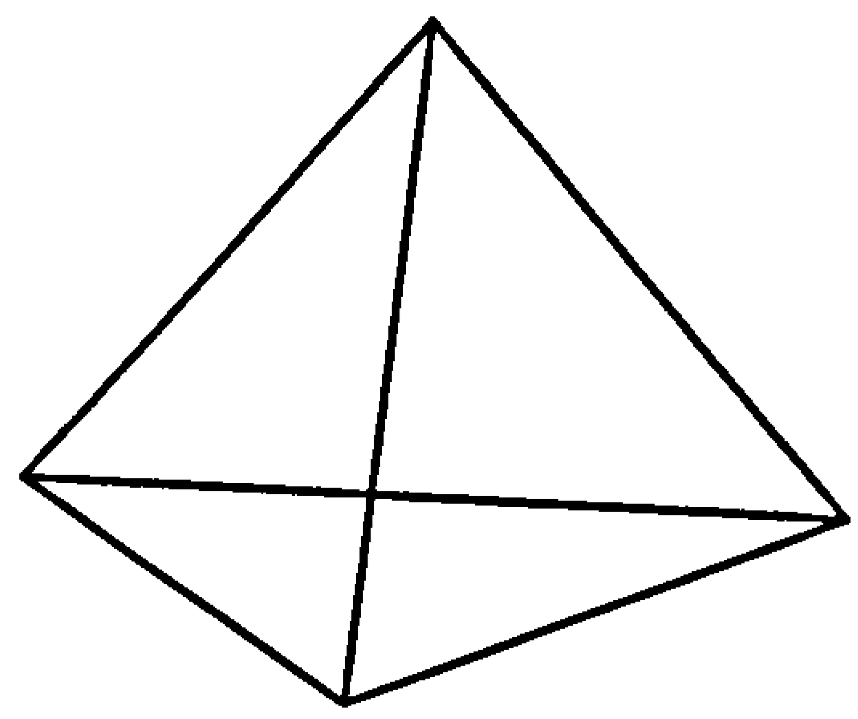

Figure 4.11 A tetrahedron

b) Polyhedron: It is made up of multiple polygonal faces and is thus the 3D equivalent of a polygon. An example is shown in Figure 4.12

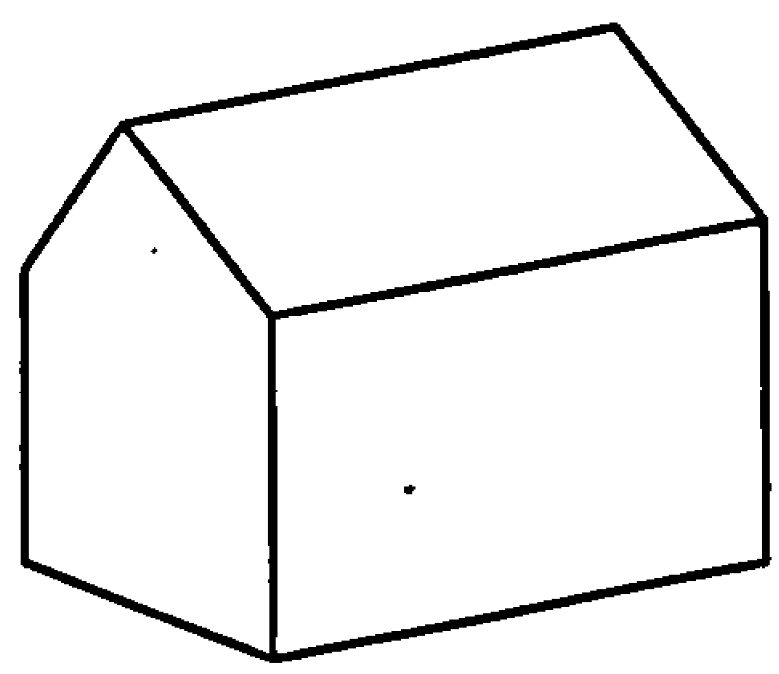

Figure 4.12 A polyhedron

c) Polyhedron combined with spherical and cylindrical patches: These three geometries are implemented in several real-world objects and therefore a combination of them can be used to model a single geographic entity. An example is shown in Figure 4.13

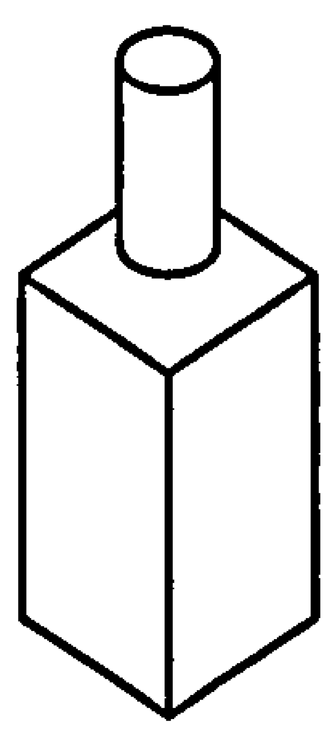

Figure 4.13 A polyhedron combined with cylinder

d) CAD objects: These models a generally beyond the scope of the current SFS model. For example, Constructive Solid Geometry (CSG) which allows a model to be created from a Boolean operation (such as union, intersection or difference) on other 3D models, as shown in Figure 4.14

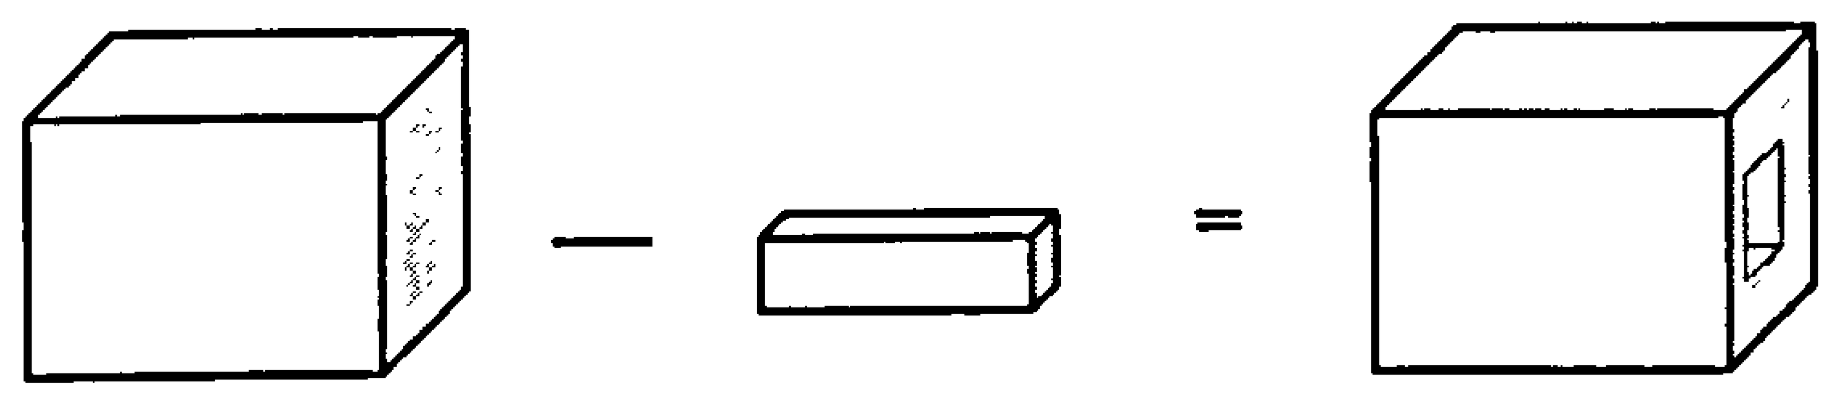

Figure 4.14 Illustrations of basic 3D geometric primitives (Arens et al., 2005)

To investigate the modeling of these primitives, our study examined both ArcSDE and PostGIS. ArcSDE uses CAD objects to model the surface of 3D geometries. The geometry type used is inherited from the ESRI Shapefile specification and is called a MultiPatch. ESRI define a MultiPatch object as an integration of different geometries, such as points, lines and polygons. As these geometries allow for 3D coordinates, the MultiPatch object can model almost any type of 3D shape. In contrast, PostGIS does not yet support MultiPatch objects. It however, supports the encoding of polygons with 3D coordinates, including triangular planes. As any surface can be approximated using a composite of triangular planes, this means PostGIS can support 3D geographic objects. A detailed discussion on the merits of any of the aforementioned approaches for modelling 3D objects would be beyond the scope of this chapter. However, Arens et al.,(2005) offer a detailed discussion. Instead, we examine two approaches for storing 3D geometries in an RDBMS and the serving the geometries through SOA.

As was established in Chapter 2, the basic geometric primitives include points, lines and polygons. The 3D coordinates of these geometries can be encoded as a sequence of X,Y,Z floating point numbers as implemented by the ESRI Multipatch geometries. For example, given four coordinates (x_1,y_1,z_1) (x_2,y_2,z_2) (x_3,y_3,z_3) and (x_4,y_4,z_4) that define a quadrilateral plane; a sequence of these coordinates could be $x_1,x_2,x_3,x_4,y_1,y_2,y_3,y_4,z_1,z_2,z_3,z_4$. It then becomes necessary to encode information about the number points encoded in the sequence and thereby enabling a visualisation application to parse the sequence and reconstruct the quadrilateral plane. The sequence of coordinates is particularly suited for storage in an RDBMS as it can be stored as a BLOB. A BLOB is technically a sequence of bytes. A group of eight bytes can represent a double precision floating point number. This approach to storing multidimensional coordinates as a sequence of bytes (BLOB) in a relational database is implemented in ArcSDE relational tables. The BLOB can then be served through a database server as discussed earlier in this Chapter.

An alternative approach is to disseminate the 3D object using a WFS or a SOAP-based traditional web service. The web service could encode the 3D geometry in a 3D scene description language such as VRML, X3D or another XML-encoded language. Scene description formats were discussed in the previous chapter. This approach requires that the visualisation application support the 3D scene description formats. The advantage of this approach is that high level 3D real-time APIs such as Java3D and OpenGL Performer support several of these 3D scene description formats, as was discussed in Chapter 3. However, the approach inherits the disadvantages of 3D scene description languages discussed in Chapter 3. By virtue of messaging through XML-encoded objects, traditional web services can also transmit instances of objects (for example, Java3D objects) directly to the visualisation application. The advantage of this approach is that parsing of the objects is automatically handled by the web service framework. However, the disadvantage is that the objects may be platform-specific, for example Java3D classes differ from OpenGL Performer classes.

This section discussed the storage and dissemination of 3D geometries using the two data dissemination approaches discussed earlier in this Chapter. The choice of which approach to adopt depends on whether a visualisation application is capable of parsing a BLOB to reconstruct the 3D geometry. If using traditional thin-client web applications, then delivering the 3D object using a standard 3D description language such as VRML or X3D is more appropriate. However, if the visualisation application is capable of parsing a BLOB and reconstructing a 3D object then the DBMS approach can be adopted. Alternatively, a web service that translates a BLOB into a standard 3D scene description format could mediate between the visualisation application and the DBMS.

4.4 Chapter Summary

This chapter has presented two approaches for the dissemination of multidimensional geographic data namely; dissemination through DBMS and web services. Two-tier and three-tier architectures for disseminating geographic data through DBMS were examined. Similarly, two approaches for dissemination through web services were also examined, that is through traditional (SOAP-based) web services and through geographic web services. The chapter also highlighted that catalogue services implementing the CAT specification can be used for supporting the STORM browser that was proposed in Chapter 3. A summary of the different web service technologies discussed in this chapter is presented in Table 4.1.

Task	Traditional Web Service	Geographic Web Service
Messaging	SOAP	GML,WFS, WMS, WCS
Description	WSDL	`getCapabilities`
Discovery	UDDI	Catalogue Service

Table 4.1 Tasks of the different traditional and geographic web services

The chapter also discussed the storage and dissemination of 3D objects using DBMS and web services. From the discussion in this chapter, we observe that all the aforementioned approaches for dissemination of data can implement a geometry model similar to the OGC SFS. We thus conclude that a multidimensional geovisualisation application should be capable of supporting DBMS and web services by resolving geometric structures from these delivery mechanisms into a representation of the SFS. These data delivery mechanisms could be used to support the web-based GeoDOVE application that was proposed in Chapter 3. The next chapter discusses the design and implementation of both the STORM and GeoDOVE prototypes.

Chapter 5 Design and Implementation

5.1 Introduction

One of the stated research questions was "could the suggested approaches be incorporated into a conventional geoportal?", this question will be addressed in this chapter. The architecture of a geoportal was presented in Figure 1.3. The figure illustrates the inter-play between catalogue, data, portal and portrayal services. Some of these services were examined in previous chapters, namely Web Map Services (WMS), Web Feature Services (WFS), Catalogue Web Services (CSW) and Web Coverage Services (WCS). In Chapter 3 we presented a variety of geovisualisation approaches and proposed novel methods for visualising geographic data and the results of a geographic search. In addition to the aforementioned webs services, database-oriented approaches for disseminating multidimensional geographic data were discussed in Chapter 4. In this chapter, we consider the development of web-based applications based on the visualisation approaches proposed in Chapter 3 and discuss how the applications are incorporated into a geoportal. As by definition, a geoportal is "a web site that presents an entry point to geographic content on the web" (Tait, 2005: pp. 34), the geovisualisation applications discussed in this chapter are served through some of the web services discussed in Chapter 4. It is envisaged that the geoportal design will support the incorporation of heterogeneous geovisualisation and geographic metadata visualisation applications.

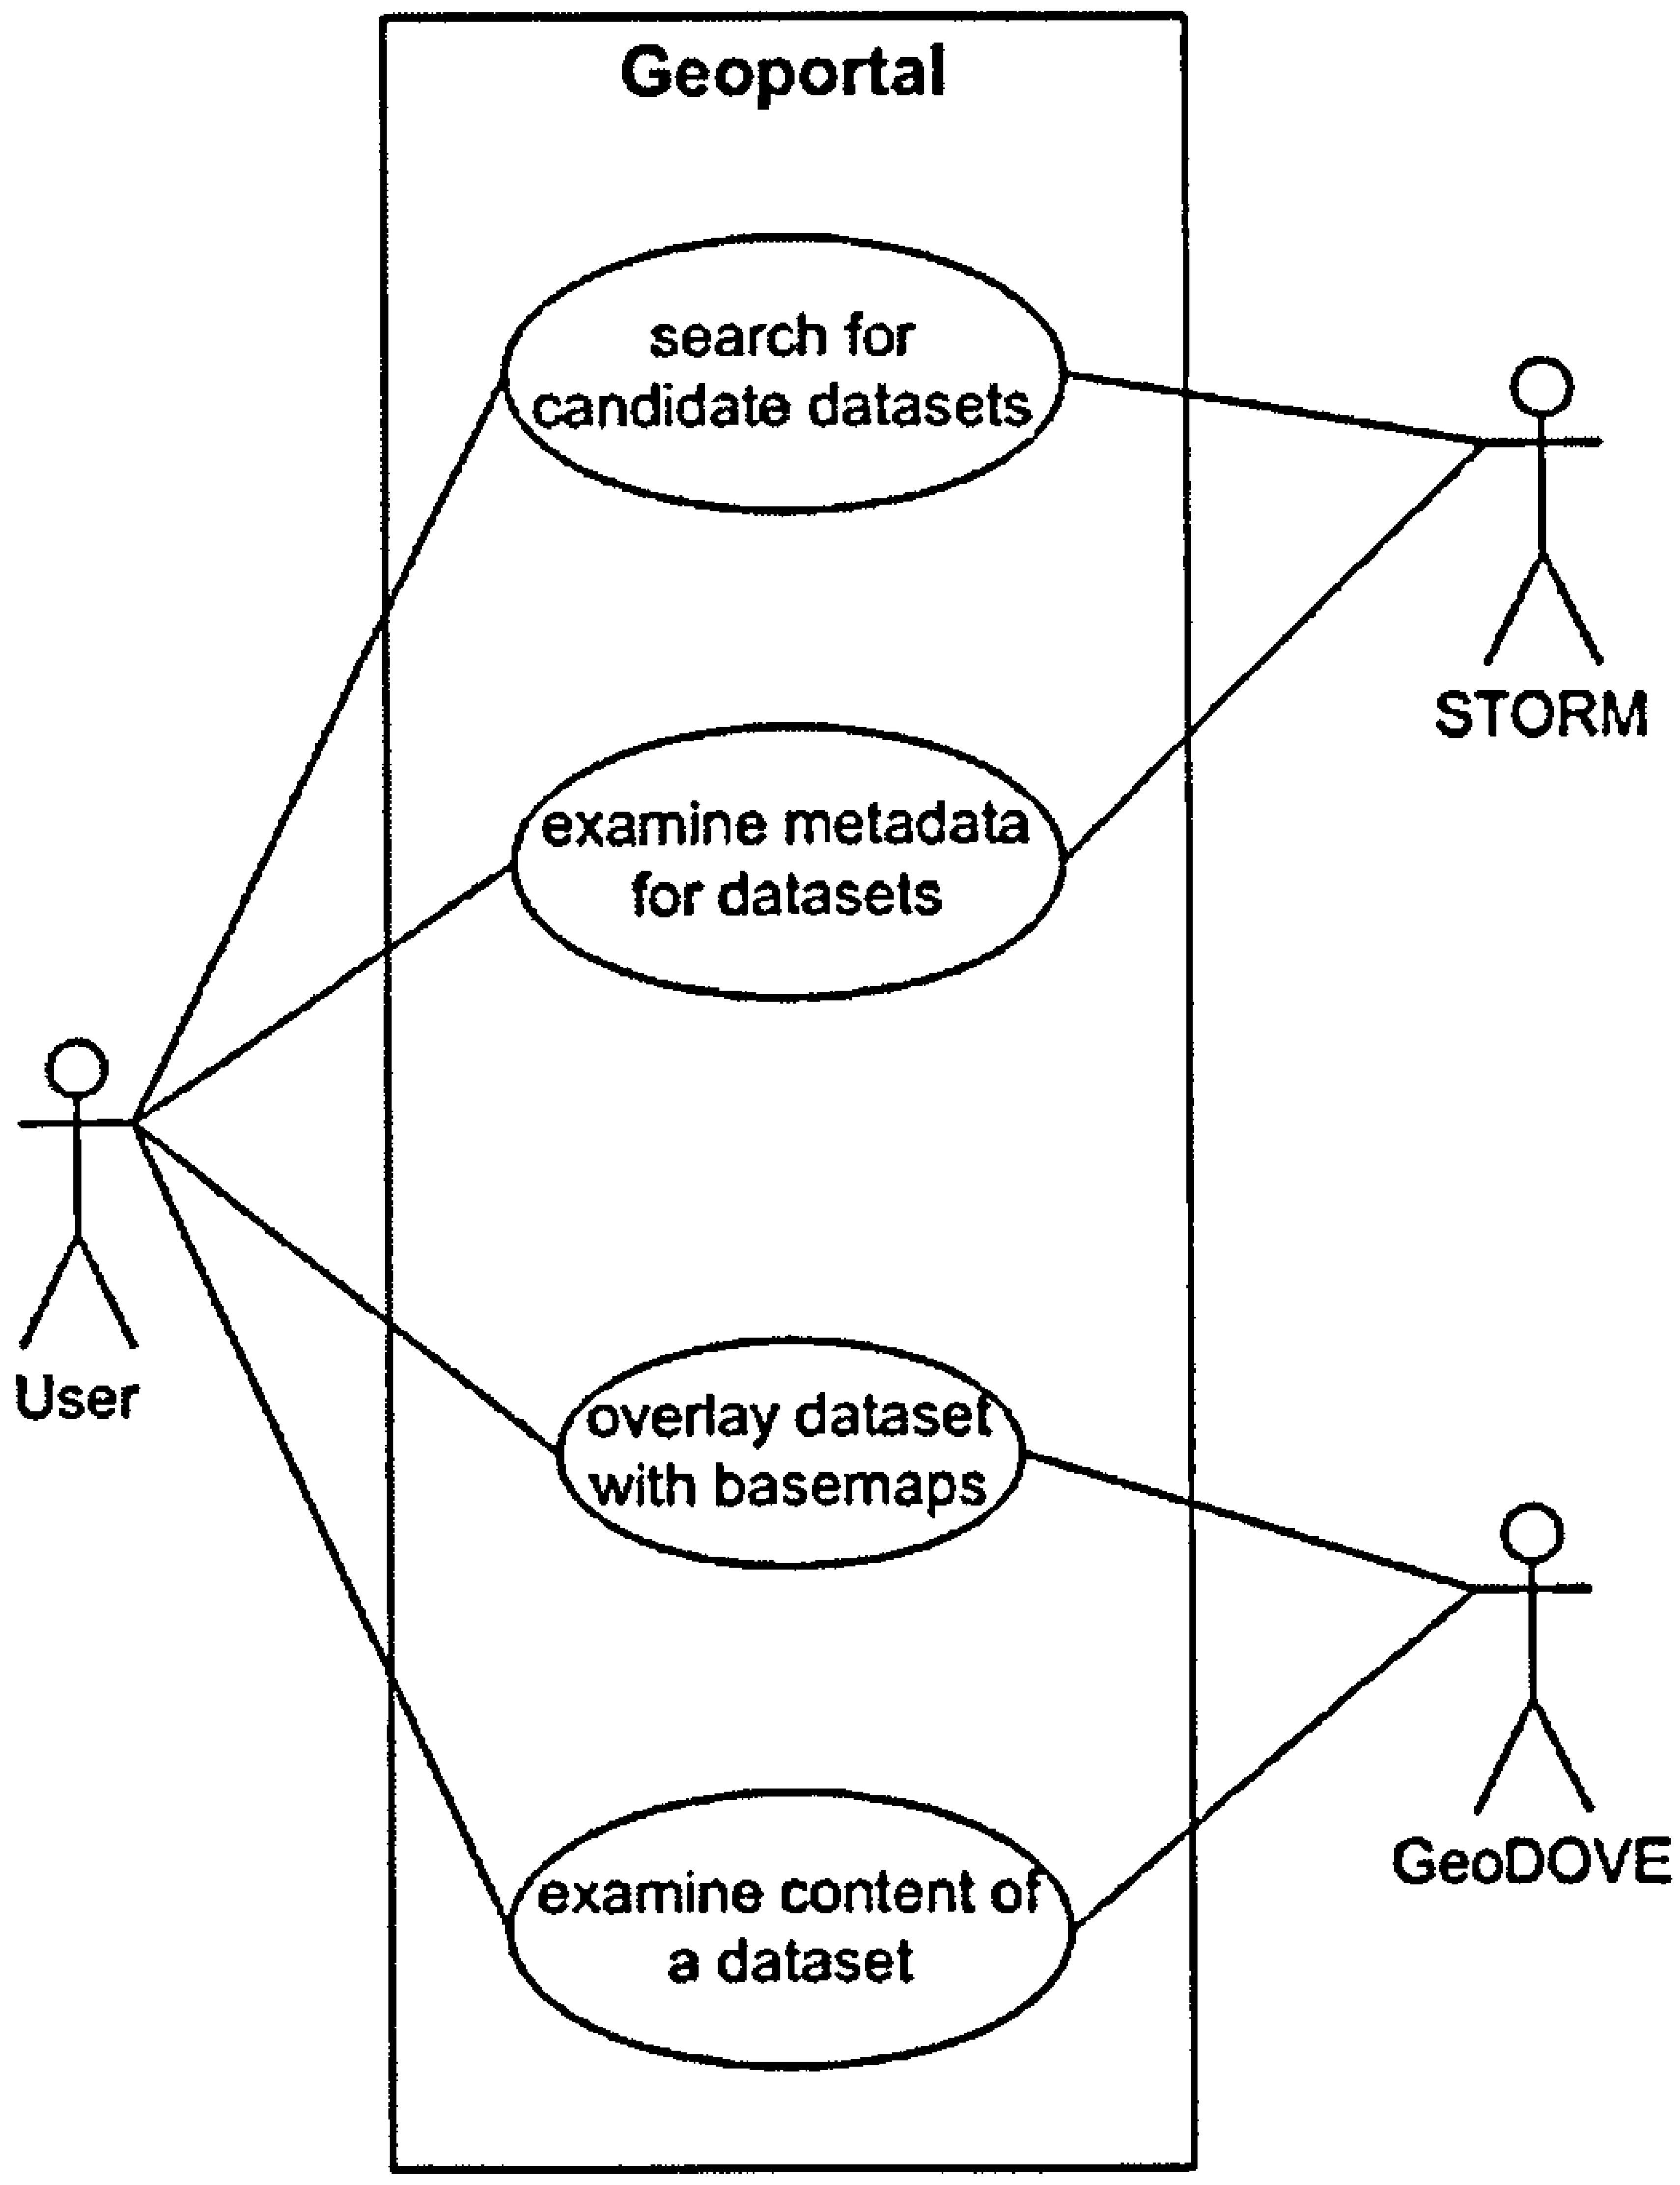

Figure 5.1 UML use case diagram showing requirements of STORM and GeoDOVE

The aim of a GIR system, including a geoportal, is to enable users to efficiently discover relevant geographical datasets. The use case presented in Figure 5.1, illustrated in UML use case notation, describes the requirements of the prototype applications in relation to the geoportal and the user. The prototypes and the user are illustrated as 'actors', with the geoportal as the platform from which the actors interact. The limitations of existing approaches for web-based geovisualisation and presenting the results of geographic search were discussed in detail in Chapters 2 and 3. Our rationale for choosing Java3D is that it is open source, web-deployable, platform independent and its applets can be embedded in a web page.

5.2 Developing the STORM System

5.2.1 Architecture

In Chapter 3 we proposed an approach for presenting ranked geographic metadata records in a 3D visualisation. The proposed approach ranks metadata records according to spatial, temporal and semantic relevance. This study developed a prototype called the STORM browser to investigate the design and implementation considerations for such an application. The system adopts a three tier client-middleware-server model. The meta-database is held and disseminated from metadata servers on the server-side. Requests and responses are processed by a metadata harvesting tool and ontology-based middleware. The resultant metadata records are then ranked and presented on the client-side by a Java3D-based client. Figure 5.2 illustrates the architecture adopted. This subsection presents a more detailed and technical discussion of the design and implementation.

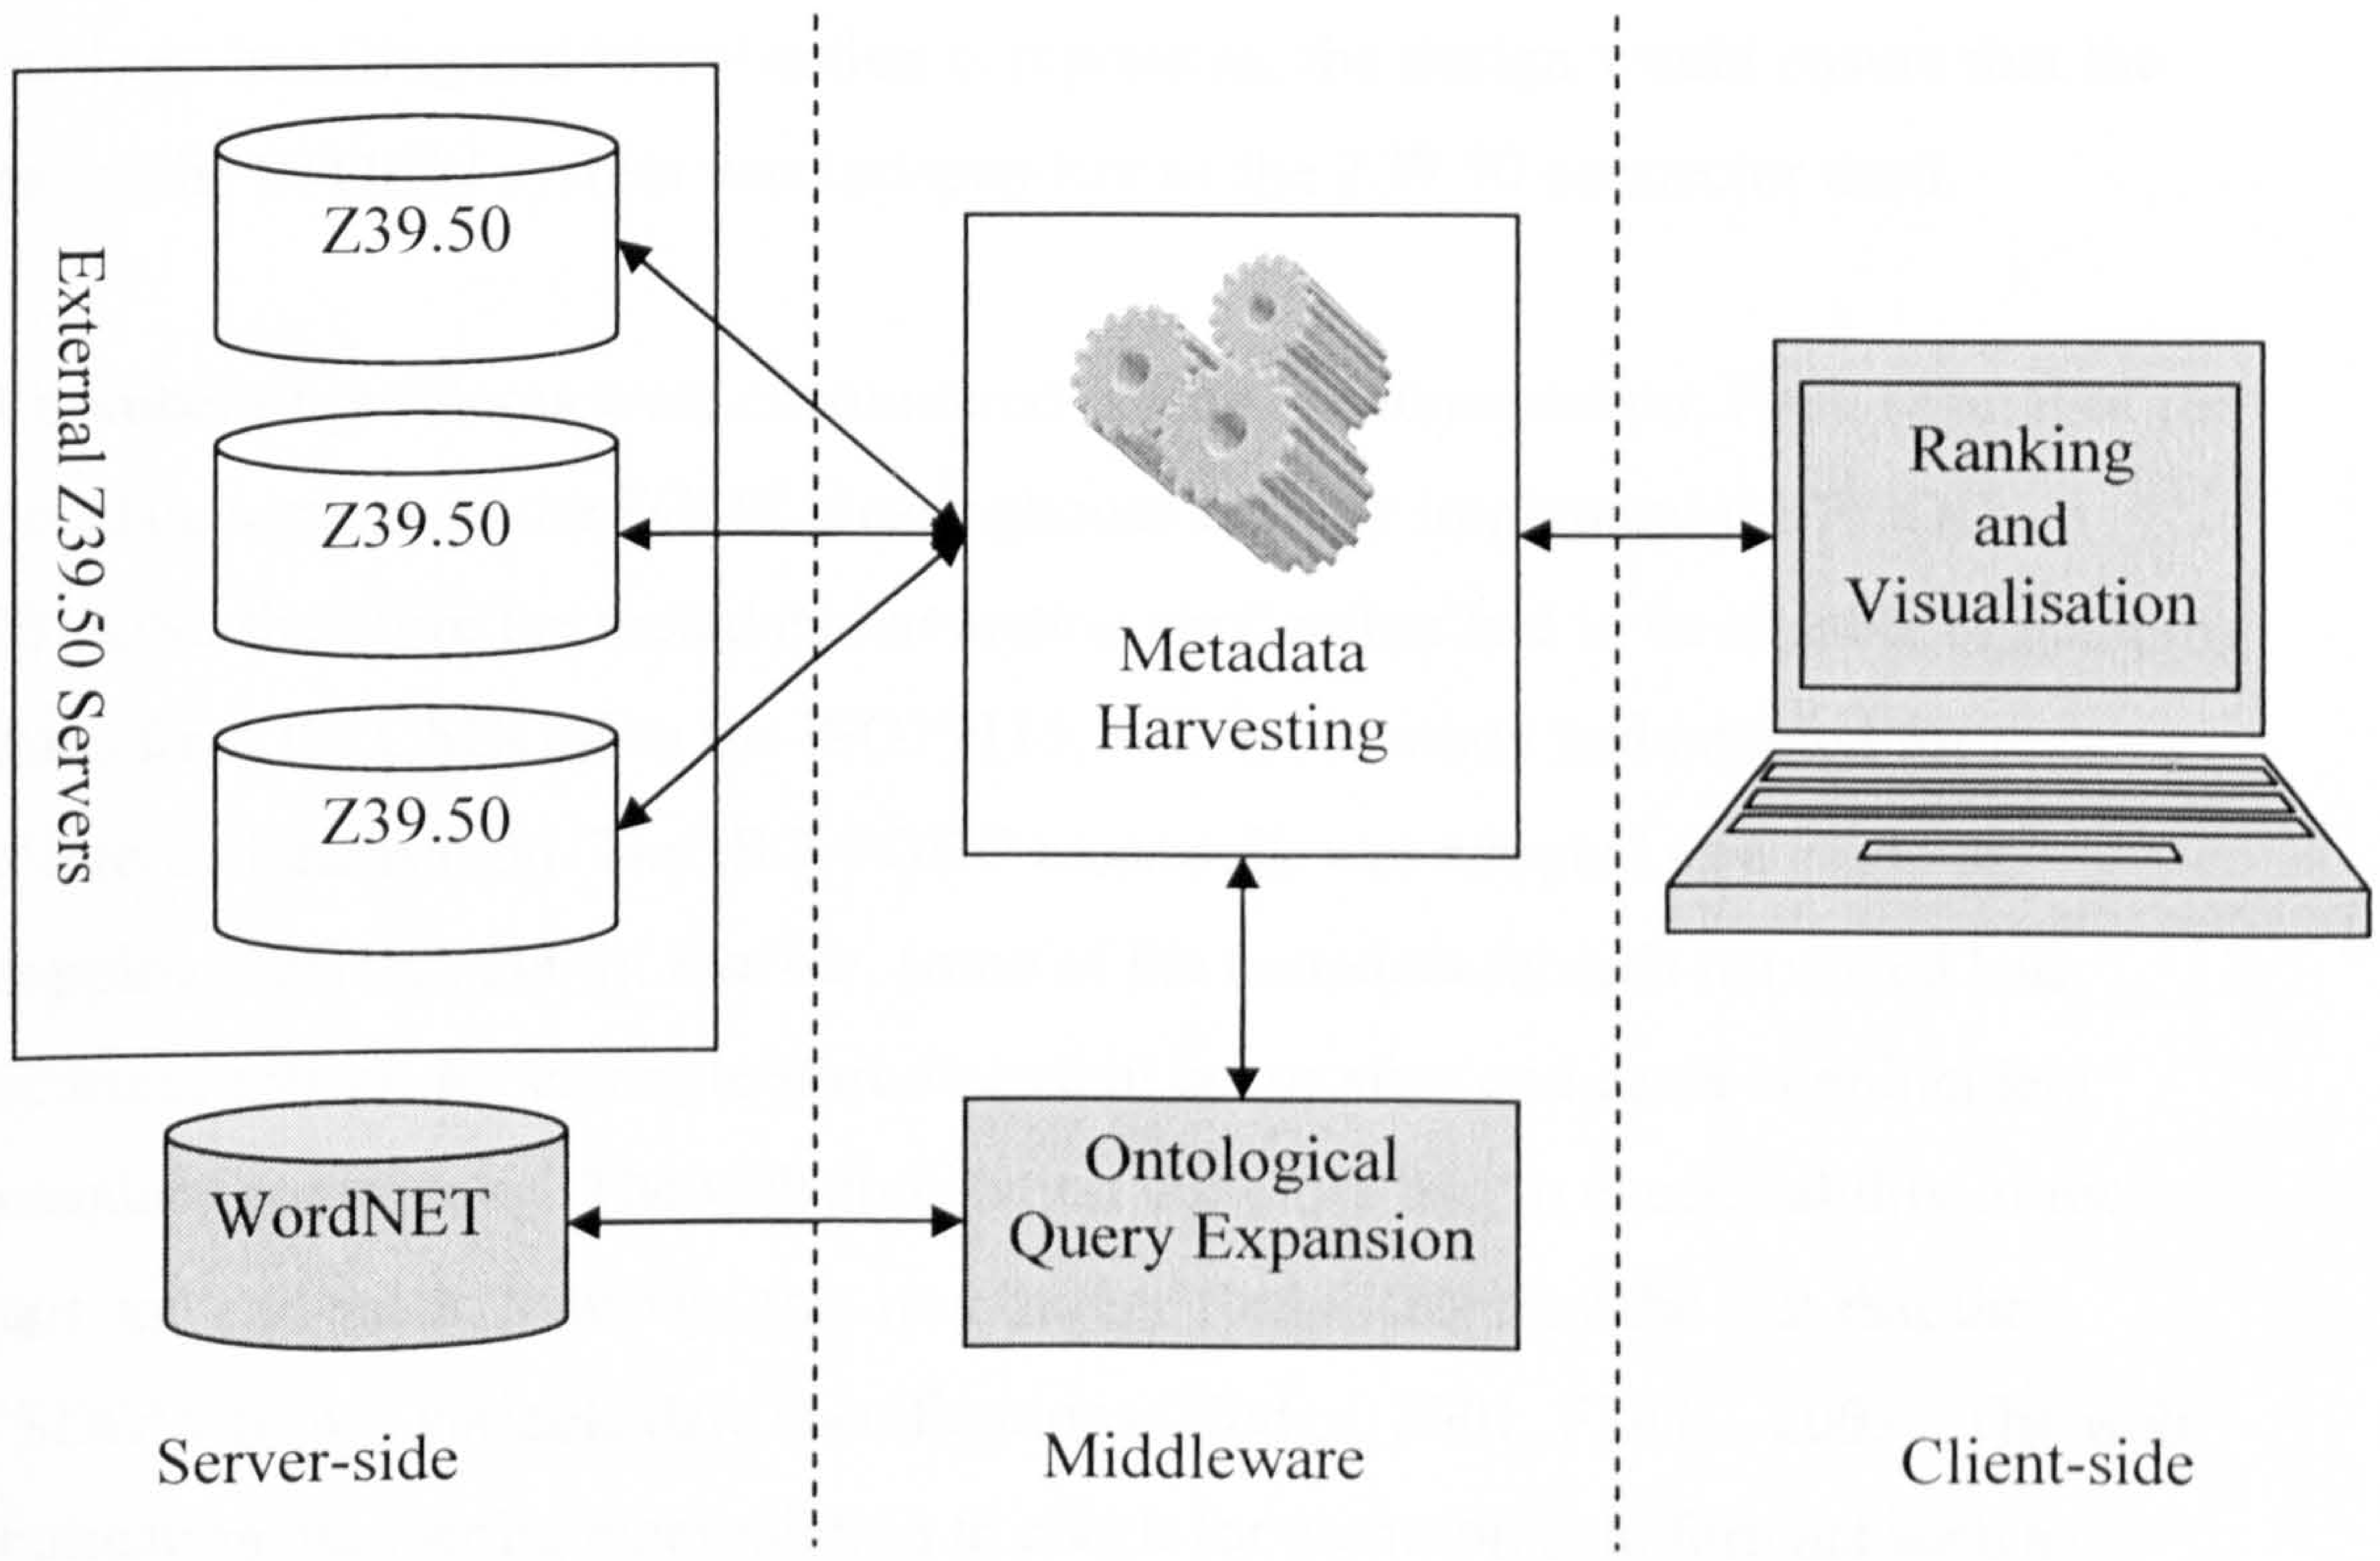

Figure 5.2 The STORM architecture

The first implementation issue considered was the harvesting of metadata. In Chapters 1 and 4 we discussed Catalogue web services (CSW). The FGDC Clearinghouse registry was identified as a major source of geographic metadata served through the Z39.50 protocol. Over 400 servers located worldwide are listed on the FGDC registry(FDGC, 2005). The Isite Z39.50 server, which is freely distributed by the FGDC, was selected for metadata dissemination. An initial version of the server-side component was implemented using VBZOOMC (Habing, 2002-

2003) open source libraries. During testing, it was observed that VBZOOMC was not adequately stable, it was replaced with the YAZ toolkit (Index Data, 2006) for interfacing with Z39.50 servers.

The web application accepted eight HTTP parameters; the server's Uniform Resource Locator (URL), the meta-database name, the search title, keyword, bounding east, bounding west, bounding north and bounding south coordinates. The web application then automatically constructed a query statement before sending the query to a targeted metadata server. Once the results are returned from the metadata servers, the web application serialises them in XML before forwarding them to the STORM browser. By separating the Z39.50-based metadata harvesting from the ontology-handling and visualisation components, the design would ensure that the rest of the STORM system was independent of the Z39.50 connector used.

A number of problems were encountered during implementation. First, several of the metadata servers on the FGDC clearinghouse registry implement the FGDC CSDGM, therefore the metadata harvesting application had to be capable of mapping fields from the CSDGM to the ISO19115, which our study had adopted for reasons of internationalisation. The ISO-FGDC crosswalk was adopted as a guide to mappings (FGDC, 2004). Further, some of the metadata records harvested had incorrect entries, for example start dates that occur after end dates or coincident boundary coordinates. The web application therefore had to check validity of the start and end dates, however, this was further complicated by the fact that the CSDGM allows variable date specifications (Hodge, 2001, FDGC, 2005). The web application uses string manipulation to check for common date formats such as DD/MM/YYYY, MM/DD/YYYY, YYYY, 'before YYYY' and 'after YYYY' – where the day of the month is DD, the month is MM and the year is YYYY. However, we cannot claim to handle all possible date specifications as the CSDGM allows for dates encoded as 'free text'. The difficulty in automated processing of free text metadata is echoed by Podolak and Demšar (2004).

5.2.2 Ontology Access

Chapter 2 discussed recent research on ontologies and their use in related studies. The STORM browser implements an Ontological Query Expansion tool for determining relationships between search terms and related concepts. As shown in Figure 5.2, the Ontological Query Expansion tool is an integral part of the STORM system. If the user's query includes a thematic keyword, the metadata harvesting tool sends a request to the Ontological Query Expansion tool which then returns a list of terms ontologically related to the search keyword. The additional keywords are then automatically added to the list of terms to query. Figure 5.3 presents a flowchart of the ranking algorithm designed for the system.

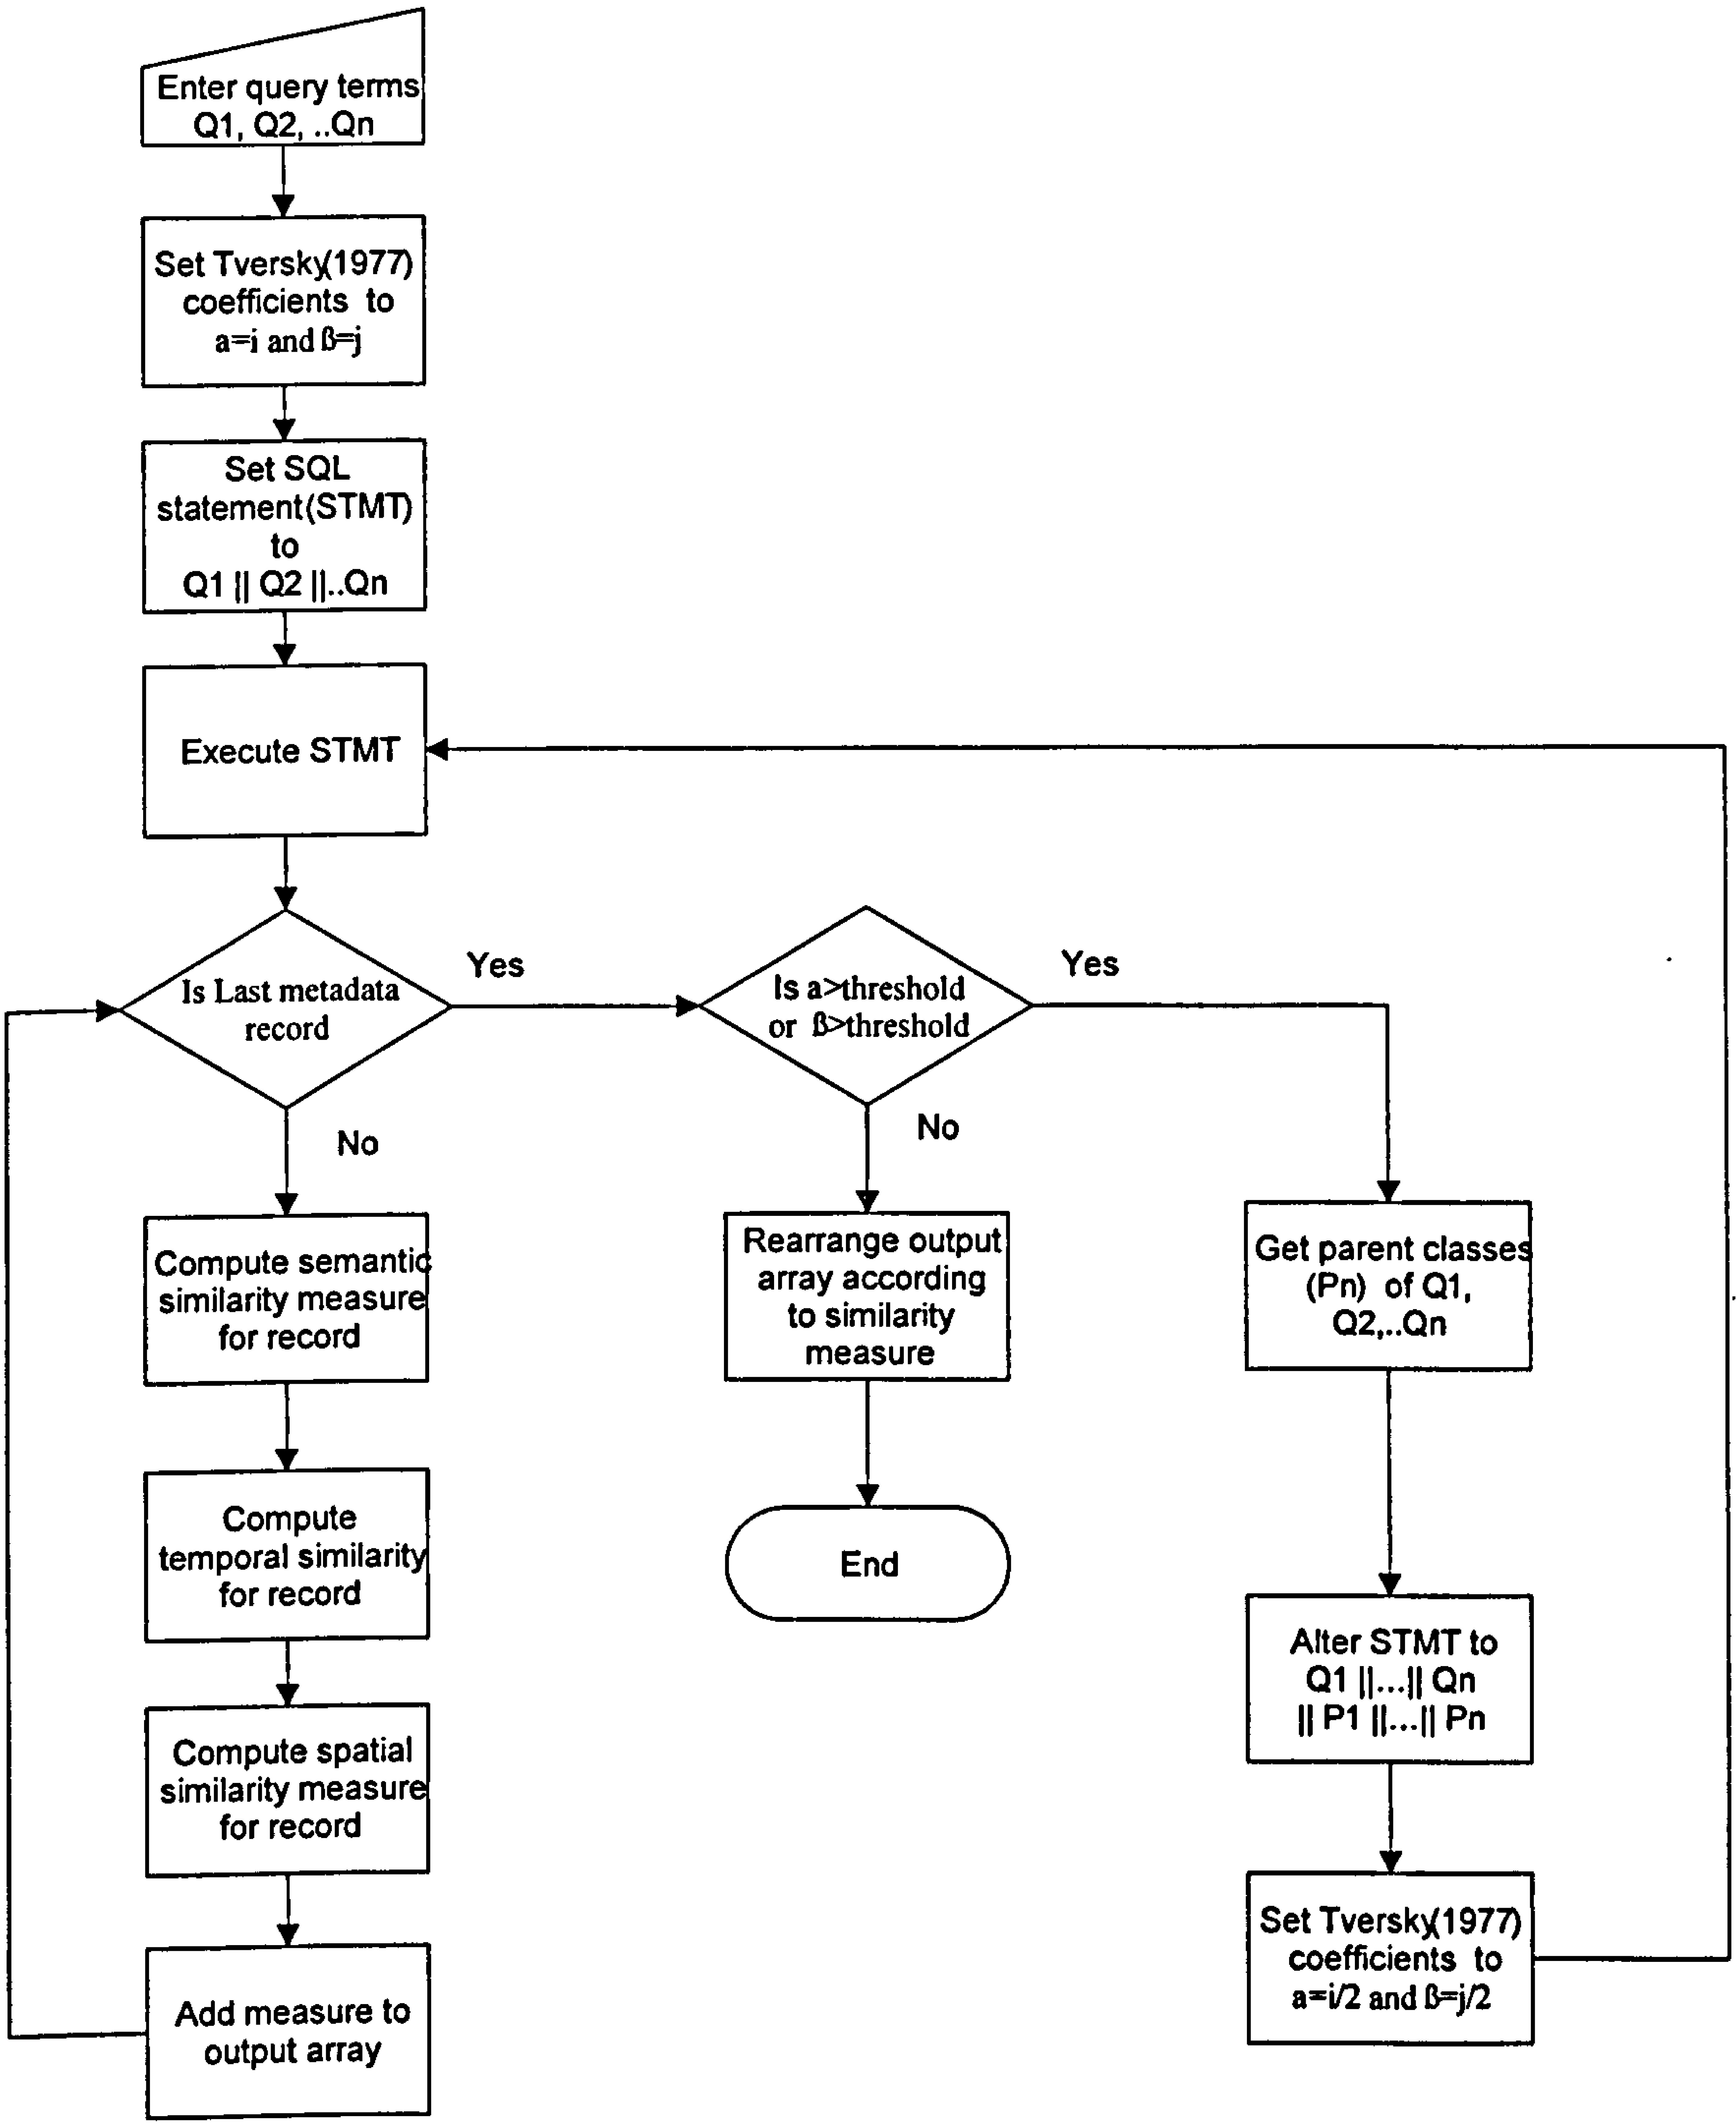

Figure 5.3 flowchart of the implemented ranking algorithm

During the ranking process, the related terms are assigned lower weighting values, as per the Tversky model presented in Chapter 2, to reflect the difference in semantics between the discovered terms and the user's original search term. For example, for a search term *skyscraper* with a weight of 1.0, a discovered term of *building* would be assigned a weight of approximately 0.7 to show that even though a skyscraper is a building, a building may not be a skyscraper (Soanes and Hawkes, 2005). The weighting values depend on the application, however they have an influence on the proximity of objects within the visualisation. It may, therefore, be necessary to select a large weight difference to reduce clatter on the semantic similarity axis. However, as our implementation is based on the model by Tversky(1977), we also consider those terms that are not related to the query term. Therefore a document A that has the original search term in its metadata may rank lower than a document B that contains a semantically-related term, if there are several other terms within document A that are not related to the original search term. Figures 5.3, 5.4 and the following equation indicate how Tversky's model has been implemented,

$$s = \frac{i}{i + |u \times (q - i)| + |v \times (o - i)| + |w \times (m - i)|}$$

where:

- i is the number of intersections between a query (or semantically related term) and a term in the metadata field
- u, v and w are weights assigned to the differences between query terms, metadata terms and intersections. The weights are fractions that sum up to 1 to depict asymmetric similarity as discussed in section 2.1.2.
- q is the number of terms in the query
- o is the number of semantically related terms
- m is the number of terms in the metadata fields (e.g. title, abstract or others)
- s is the similarity measure

From the equation, similarity decreases with an increase in the number of terms that are not shared between the query terms and metadata fields. The process illustrated in Figure 5.4 increases the intersection index (i) each time a match between a query term and a term in a metadata field is found. After the query terms have been compared to the metadata, the process is repeated for terms semantically related to the query terms. This feature-based approach differs to edge-counting methods of semantic similarity as we assume equal distance between concepts. As discussed in section 2.4.4, conceptual distance between concepts is not always uniform leading to irregular densities of links between some concepts. Assumption of equal distance between concepts is, therefore, a limitation of our approach. Our approach differs to information content-based methods, as well, as we ignore the frequency of occurrence of a term within a corpus. This, however, protects our similarity measure from terms with multiple meanings, for example, the term 'bank' could refer to a river bank or a commercial bank.

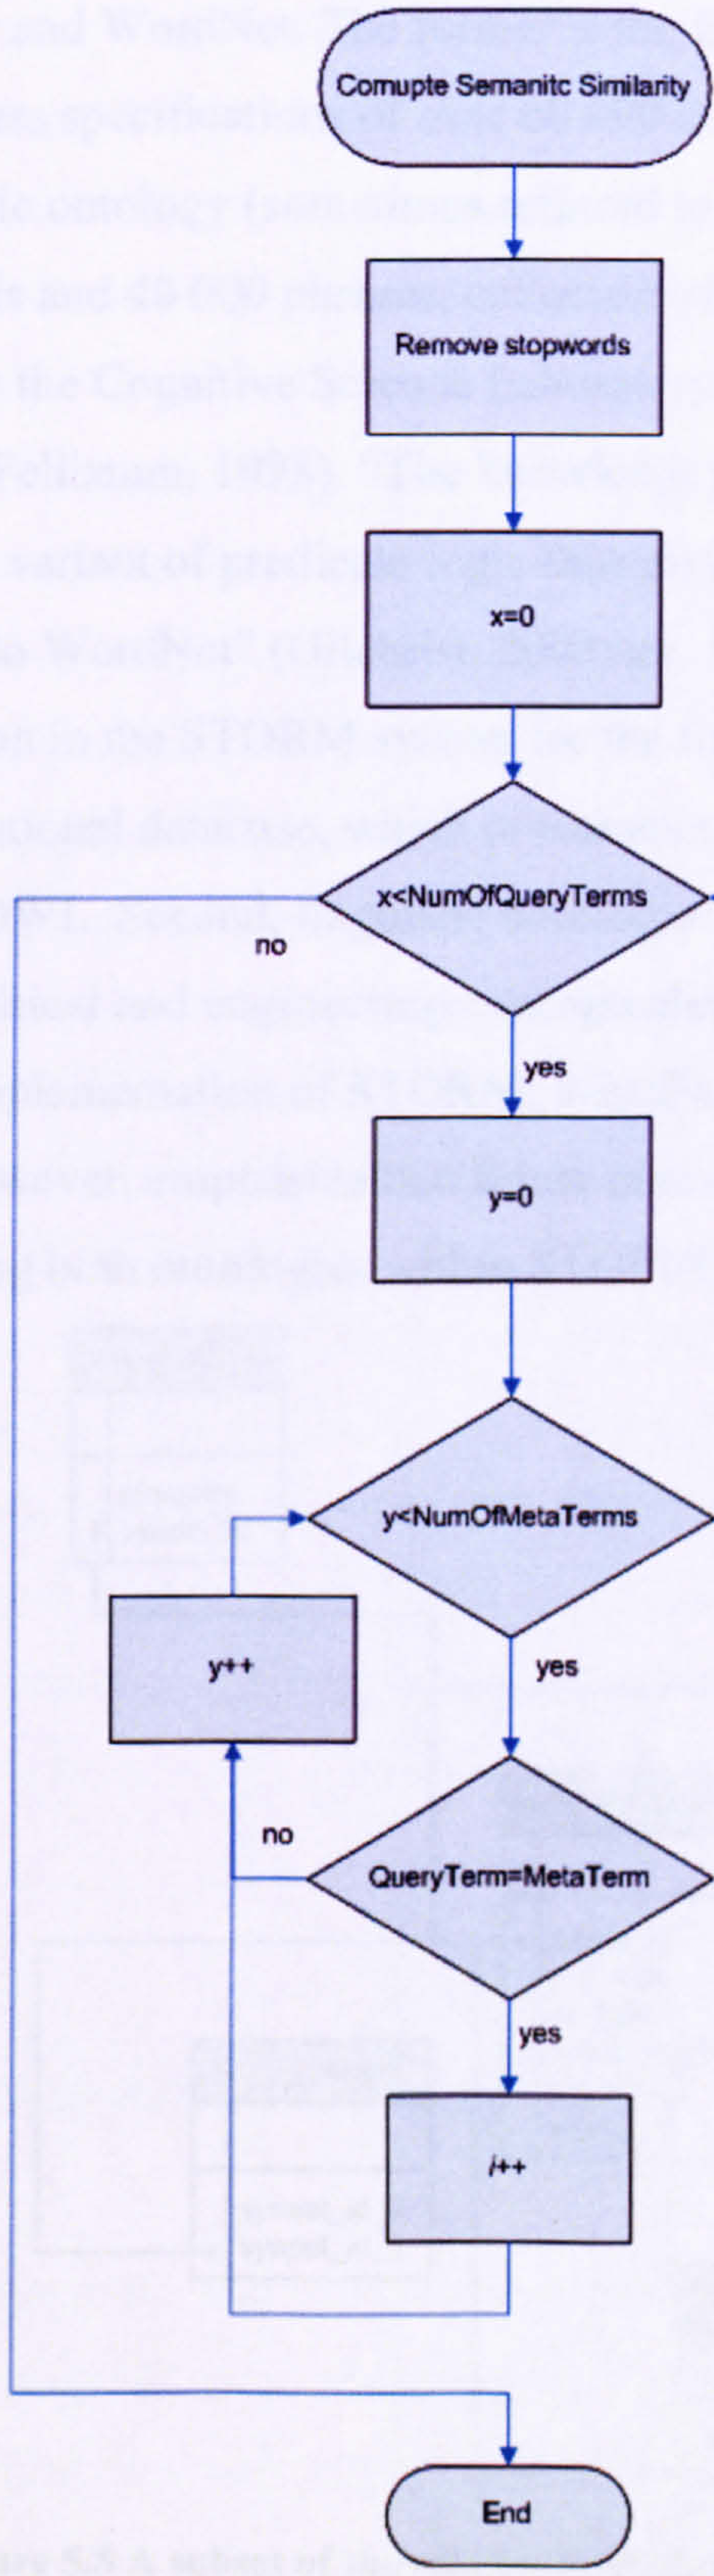

Figure 5.4 Flowchart of semantic comparison of query to terms in metadata descriptions

For implementing the ontological components, two ontologies were examined, namely OpenCyc and WordNet. The former is the free subset of the Cyc ontology and currently offers specifications of over 60 000 concepts (CycCorp, 2003). The latter is a linguistic ontology (sometimes referred to as a lexical database) containing over 50 000 words and 40 000 phrases, collected into over 70 000 sense meanings and developed by the Cognitive Science Laboratory at Princeton University (Agarwal, 2005, Fellbaum, 1998). "The knowledge in Cyc is represented declaratively in a variant of predicate logic thus giving CYC more reasoning power than is available to WordNet" (Gilchrist, 2003: pp. 13). Our design adopted WordNet for implementation in the STORM system for the following reasons: First, it is available as a relational database, which processes queries faster than ASCII-based formats such as OWL. Second, linguistic ontologies, such as WordNet, offer a bridge between philosophical and engineering conceptualisations (Agarwal, 2005). Third, at the time of the implementation of STORM, WordNet offered more concepts than OpenCyc. We however, emphasize that future research should examine the possibility of using both ontologies within STORM.

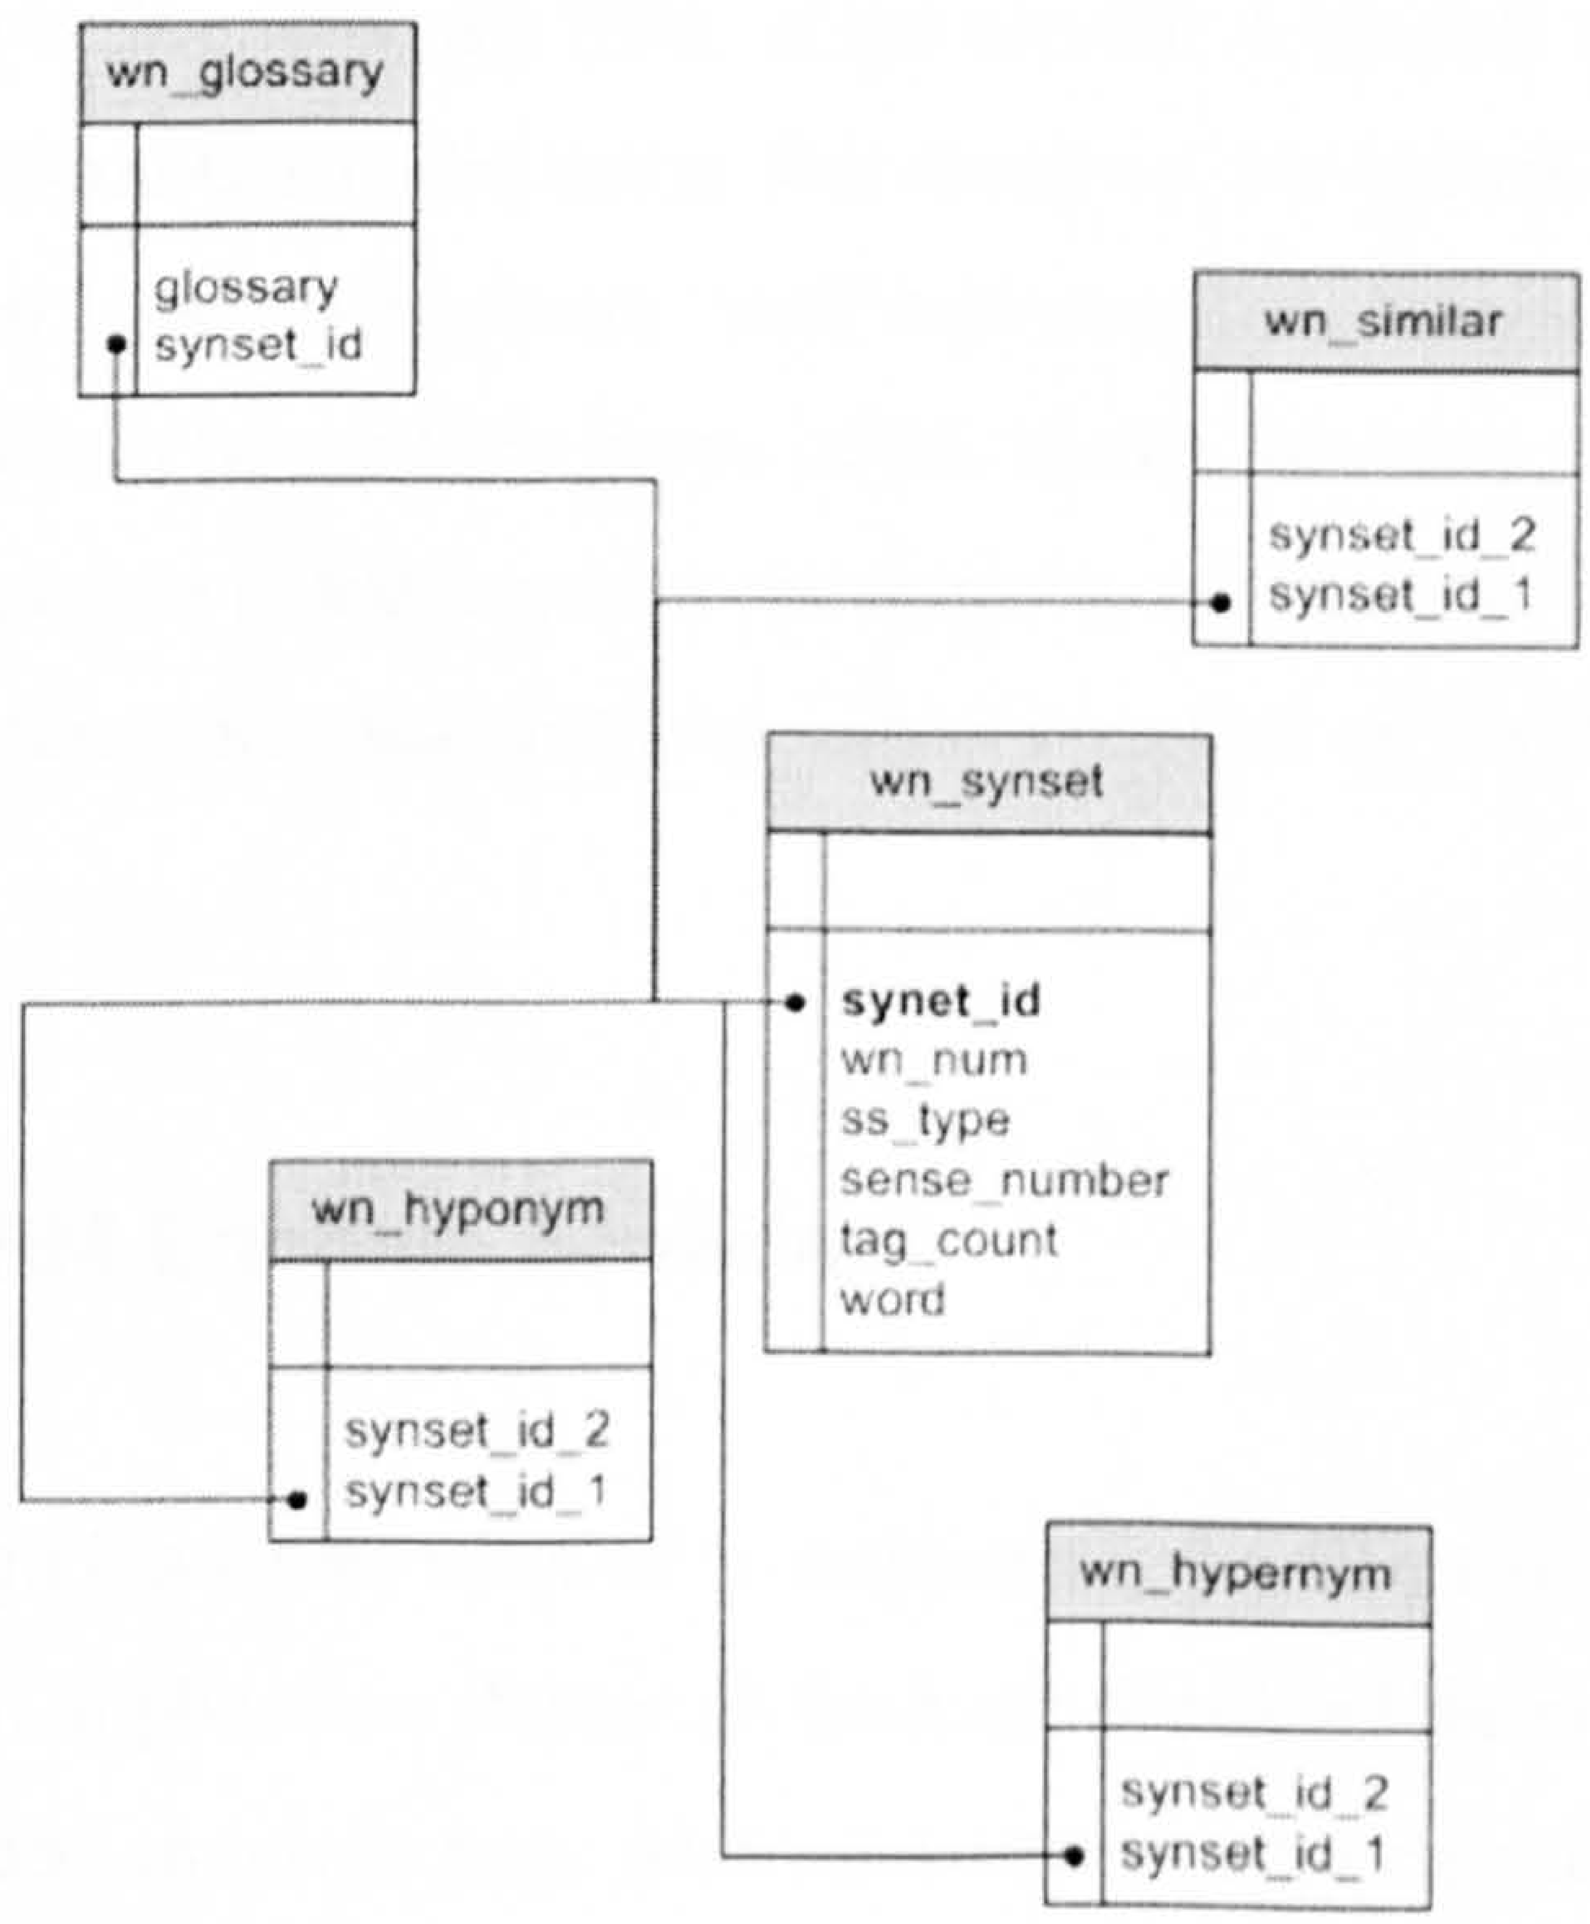

Figure 5.5 A subset of the relational implementation of Wordnet

A relational implementation of Wordnet was adopted for providing a lexical database (Android Technologies Inc., 2006). This database contains a relational table for similar terms, hypernyms, hyponyms, meronyms, antonyms and other relations. To reduce the possibility of retrieving irrelevant terms, our implementation only adopted the tables of similar terms (wn_similar), hypernyms (wn_hypernym), hyponyms (wn_hyponym) and synsets (wn_synset). An illustration of these tables is presented in Figure 5.5. The process of retrieving terms from the relational database was to,

- create a view (called wn_tempo) from the join of the wn_synset and wn_similar tables, with both the synset_id and word fields as constraints
- select all from the join of wn_tempo and wn_synset, with the synset_id_2 field as the constraint.
- repeat the first two steps for other semantic relations (in our case, hypernyms and hyponyms).

The associations created by the relational joins, are illustrated as circular-headed lines in Figure 5.5. Once the associations have been retrieved, the extracted terms are exported as XML through a servlet. The STORM-enabled browser then retrieves the XML document by referencing the URL of the servlet. A limitation of our approach is that we do not disambiguate the terms. For example, our system does not distinguish between senses of the term 'well' referring to a 'borehole' and 'feeling'. Term disambiguation is beyond the scope of this thesis, however, we acknowledge that term disambiguation could improve the retrieval of terms for the STORM system. The next subsection describes the implementation of the STORM-enabled browser.

5.2.3 The STORM-Enabled Browser

An important design issue was whether to implement the client application as an applet or a webstart application. Although both are deployed through the web, applets are embedded on a web page whereas webstart applications run independently of the web browser. Nielsen (2000: pp. 258) proposes that applets that involve accessing real-world data existing external to the hosting web page should be displayed in a new non-browser window (that is, as webstart applications). The STORM browser was therefore designed to operate as a webstart application composed of three main Graphical User Interface (GUI) components, i) a five column table for the ranked list showing the one-dimensional rank, the title, spatial, semantic and temporal score of each dataset, ii) a 3D graphics panel for the STORM visualisation and iii) an HTML panel for presenting geospatial metadata for a selected dataset. A panel is a section of a GUI window. This design satisfies recommendations by Ware and Plumlee (2005) that all relevant information should be placed within the same field of view to aid navigation; in our case, the 3D visualisation, the metadata panel and the ranked list are presented in the same window. Another recommendation addressed includes the use of hyperlinks through the selection of datasets by clicking on records on the ranked list or clicking on thumbnails in the 3D visualisation.

An additional design consideration was how to assist the user in detecting distances through the provision of calibrated axes. The size and shape of the thumbnails is kept constant so as to enable the user to detect depth through perspective, that is, objects that appear smaller are further away than those that appear bigger. Different graphics are rendered for vector and raster datasets on the thumbnails. A vector dataset is represented with a cartographic map of a globe, whereas a raster dataset is represented with a satellite view of the earth. Both of these thumbnails are illustrated in Figure 5.7. This type of symbolism allows for immediate identification of the type of dataset. With three panels on the main window, another design consideration was how the user interacts with the different GUI components. We followed a conventional GIS model for interaction; a desktop GIS allows clicking of an object in a visualisation to highlight the representation of that object on an attribute table and vice versa (Longley et al., 2001). In our case, however, there were three panels (i.e. including the HTML panel for presenting metadata) hence selection of an object in the visualisation or a record on the table would have to update the metadata HTML panel as well.

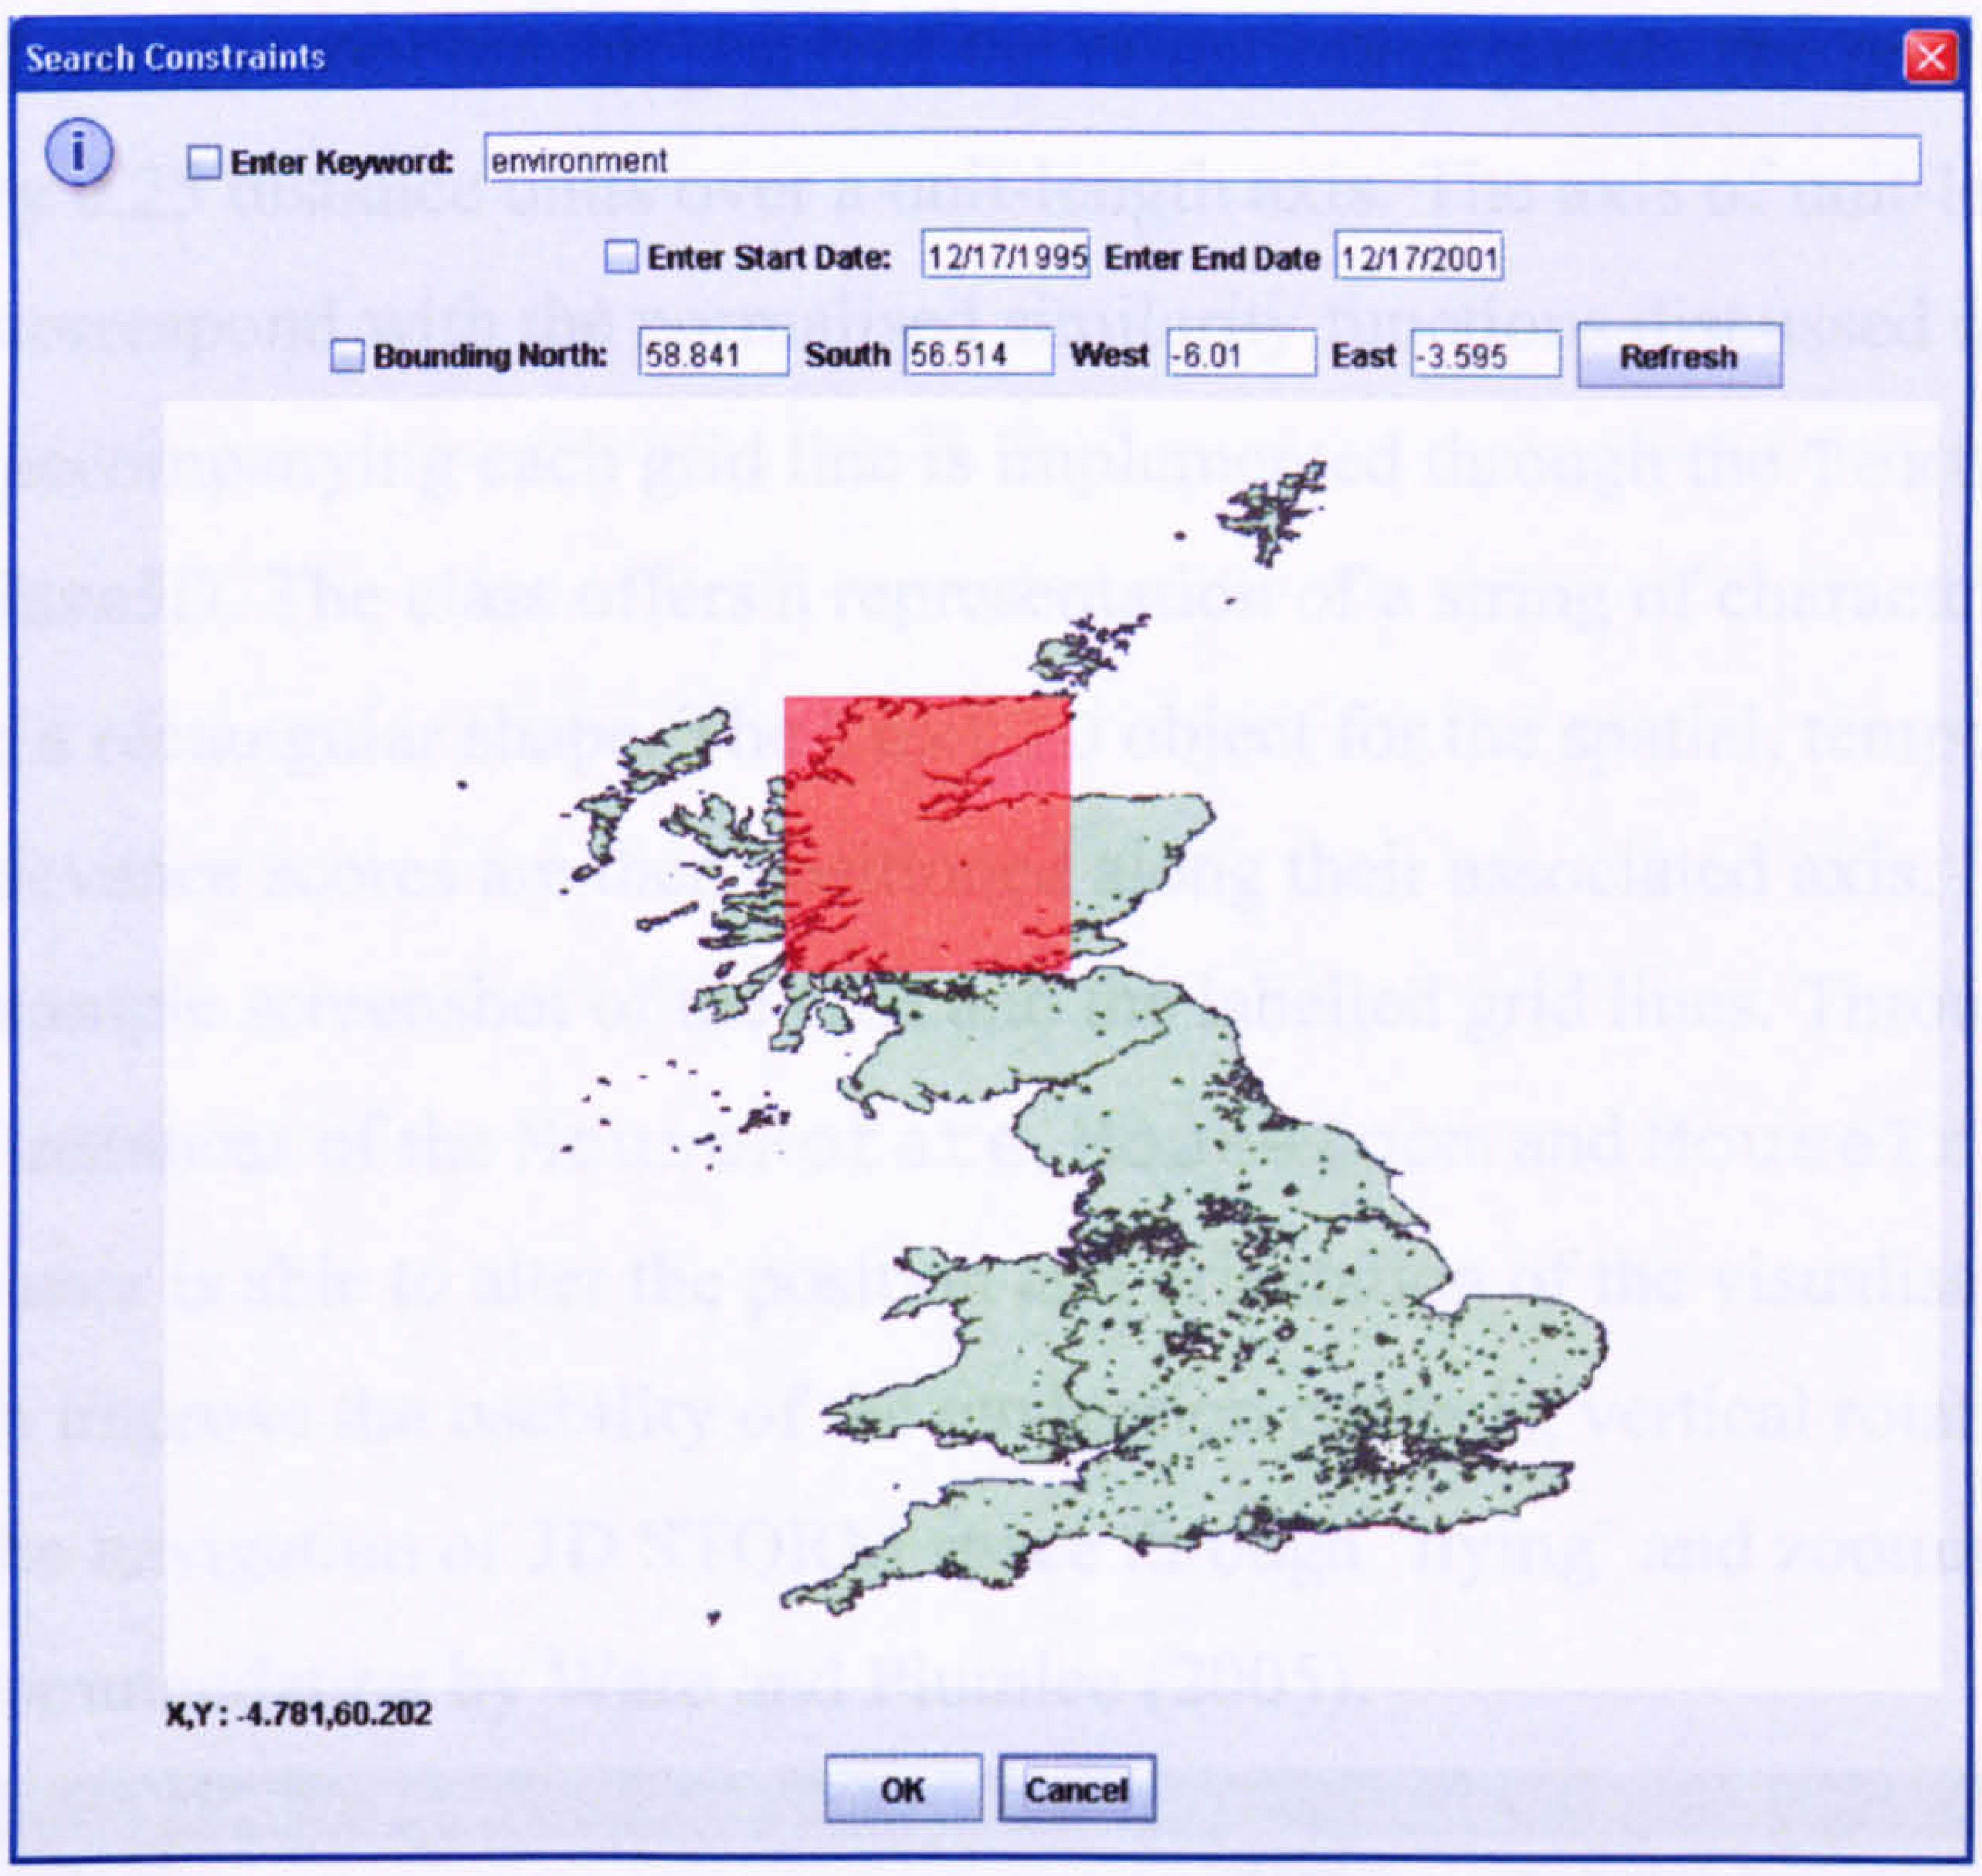

Figure 5.6 Dialog for entering search constraints. Ordnance Survey © Crown Copyright

Another design consideration was how to provide a query interface. To avoid overwhelming the user with an excessive number of controls, a separate query window only appears upon clicking on a menu button, as illustrated in Figure 5.6. The query window included textboxes for entering the keywords, temporal and spatial constraints to search. The spatial constraints are entered either as bounding coordinates or through an interactive map that automatically alters the coordinates in the textboxes. The metadata for each dataset is compared with the keyword entered by the user to determine semantic similarity. The start and end dates entered are compared with the dates of creation retrieved from the metadata. The bounding coordinates entered by the user are tested against the bounding coordinates of each dataset to determine if the query area contains or overlaps the geographic extent of the dataset. Although not implemented in the current version of the browser, an alternative method for specifying spatial constraints could have been the use of an ontology-driven gazetteer, as implemented by the SPIRIT project (Jones et al., 2001).

The visualisation panel extends the `Canvas3D` class which is the main drawing GUI component for rendering in Java3D. A grid of `LineArray` objects is implemented and inserted into the root branch graph of the visualisation panel. The grid lines are spaced every 0.25 distance units over a unit-length axis. The axis of unit-length is adopted to correspond with the normalised similarity functions discussed in Chapter 2. The text accompanying each grid line is implemented through the `Text2D` class offered by Java3D. The class offers a representation of a string of characters rendered on a rectangular shape. The `Text2D` object for the spatial, temporal and semantic relevance scores are then positioned along their associated axis. Figure 5.7 shows an example screenshot of the GUI and the labelled grid lines. Through the addition of instances of the `MouseRotate`, `MouseZoom` and `MouseTranslate` classes, the user is able to alter the position and orientation of the visualisation. However, to improve the usability of the navigation controls, vertical rotation was disabled. The navigation of 3D STORM space through 'flying' and zooming satisfies another recommendation by Ware and Plumlee (2005).

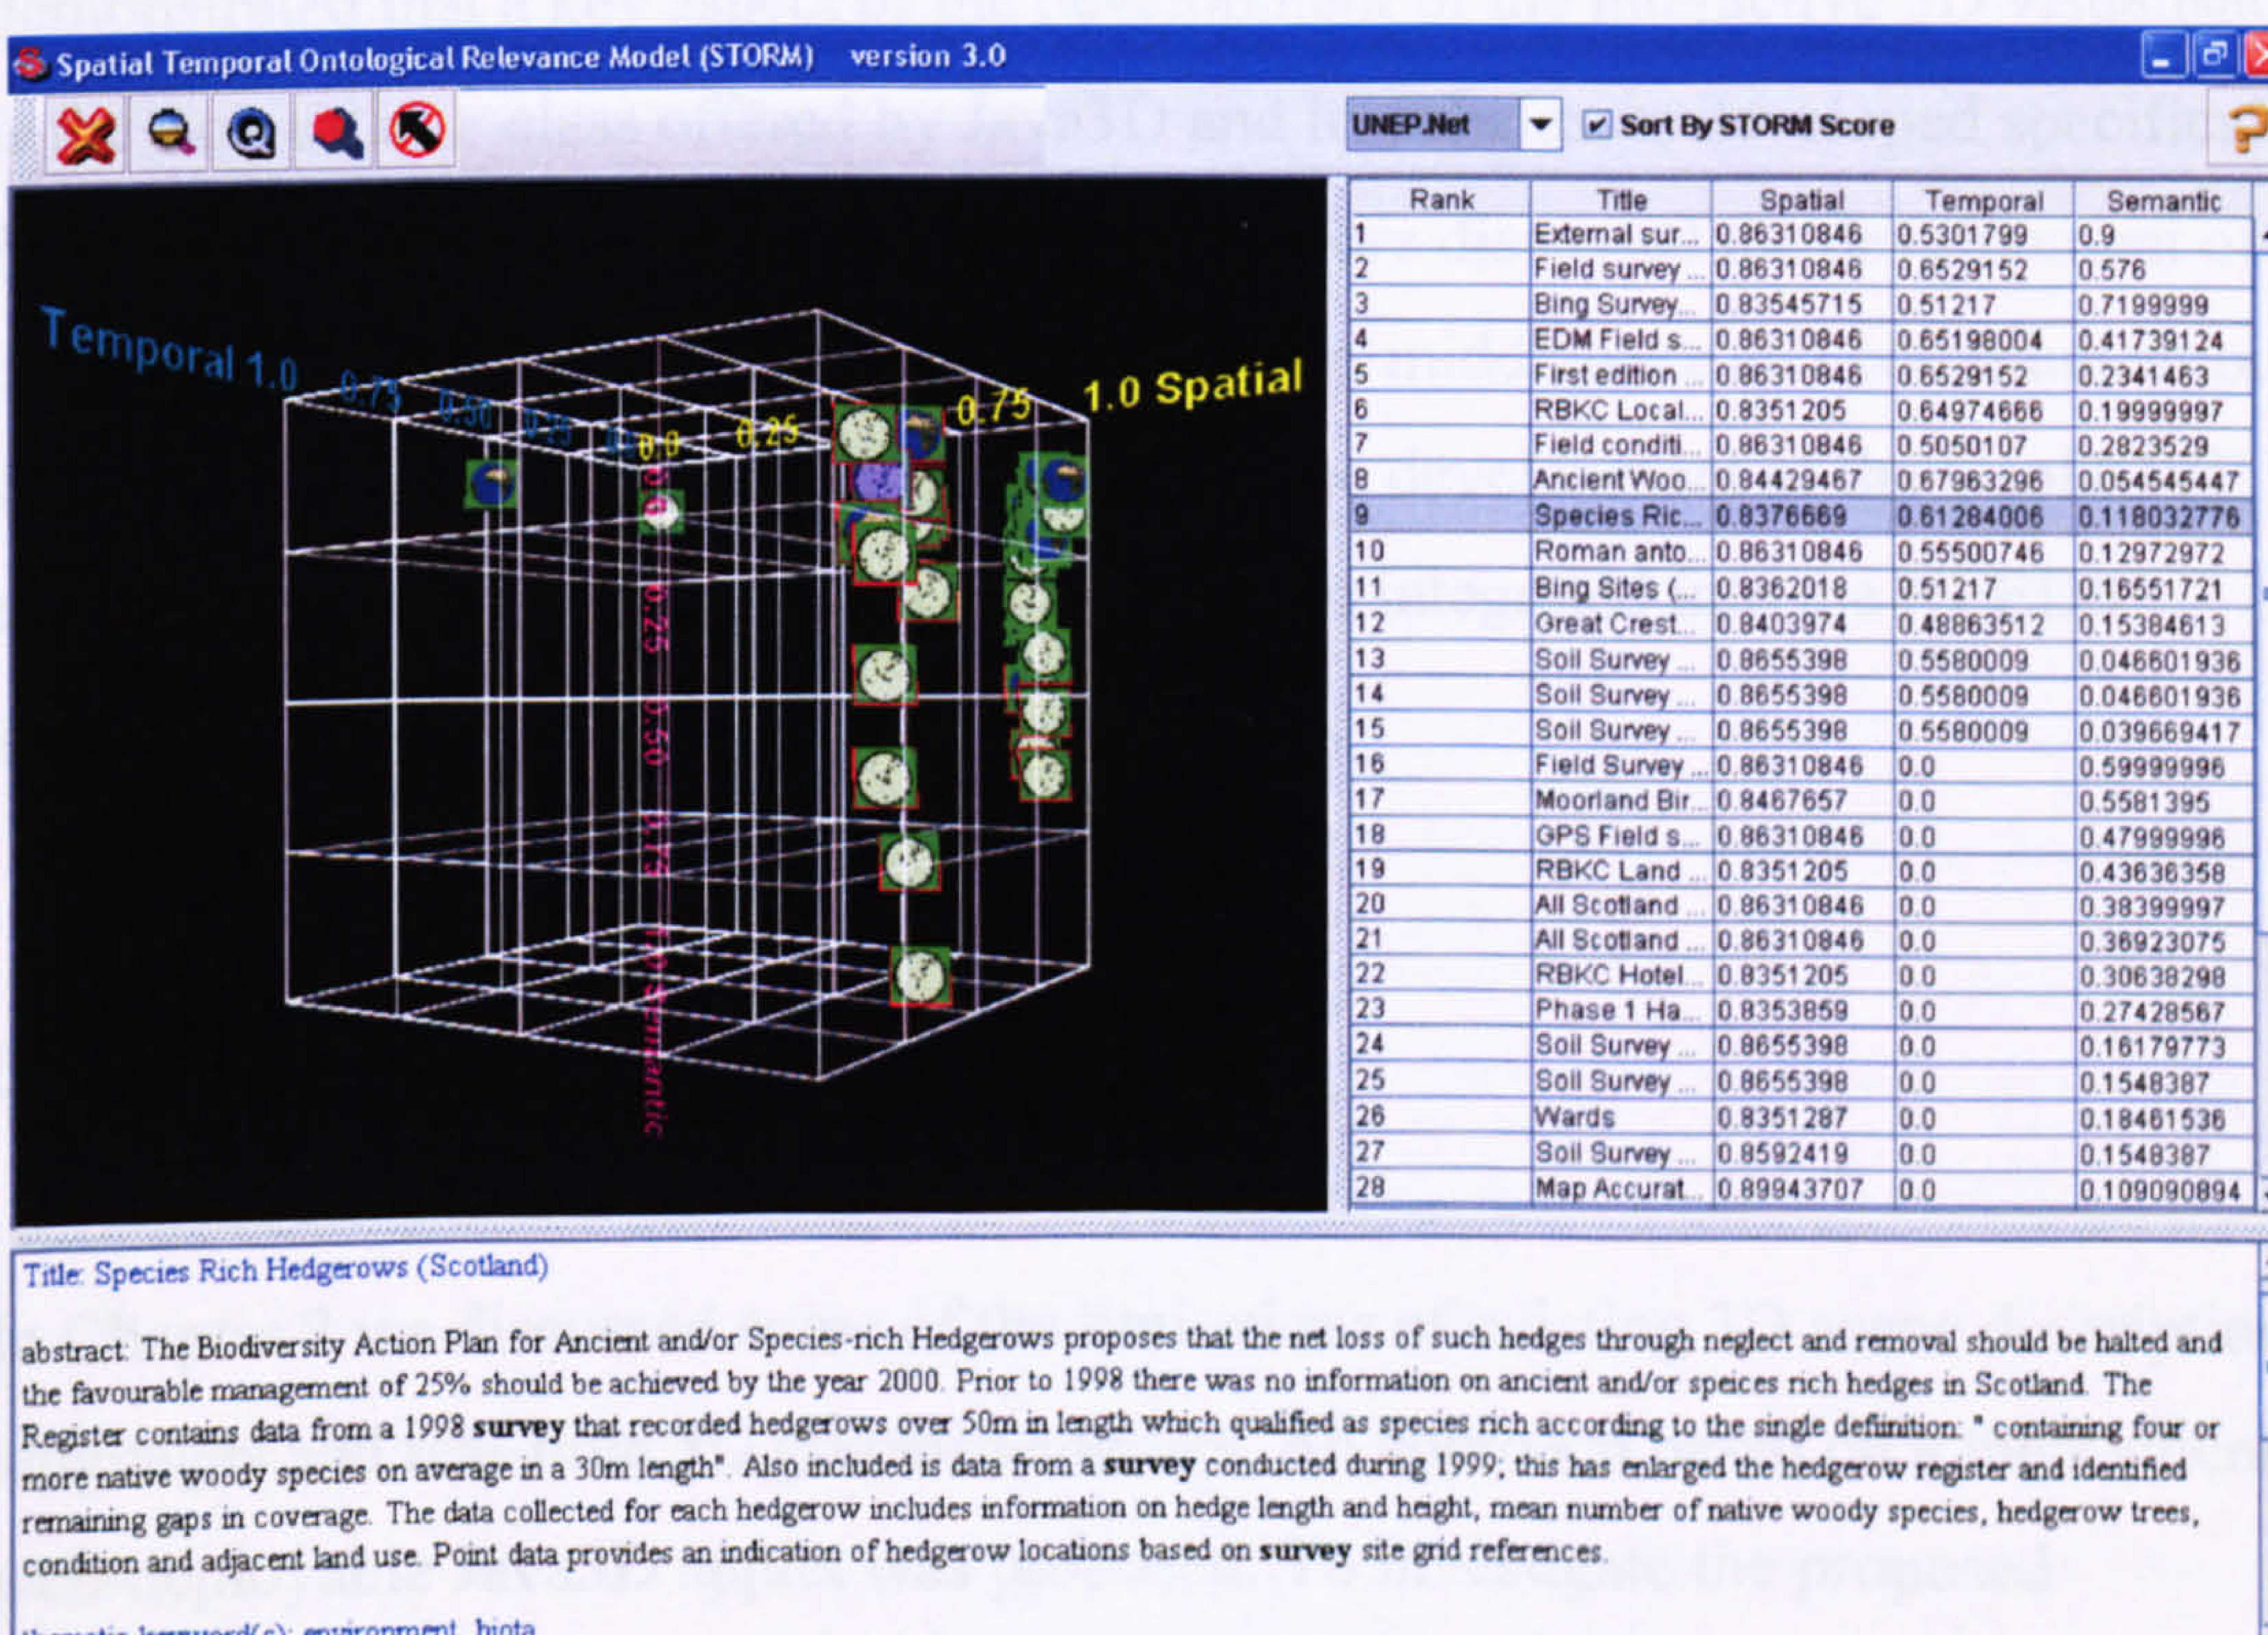

Figure 5.7 Viewing from an arbitrary angle (axis)

As aforementioned, an HTML panel was implemented for presenting the geographic metadata. The applet stores metadata for each dataset in an instance of the class `Dataset`. Within each `Dataset` object the values and names of metadata fields are stored in an encapsulated Java `Hashtable` object. The `Hashtable` class offers temporary storage of key-value pairs (Horstmann and Cornell, 2001). A program can extract each entry by specifying a key (field) to retrieve. For multiple occurrences of metadata fields such as two place names in the same metadata record (for example United Kingdom and England), STORM concatenates the values with commas separating each value. When a dataset is selected either in the visualisation panel or the ranked list, an HTML script is automatically composed from the key-value pairs stored in the `Hashtable`. The HTML script is then added to an instance of the `JTextPane` class, a Java GUI component that renders HTML markup.

As was described in this section, the cross-operation of the browser, middleware and external metadata servers was made possible through communication based on standard approaches such as the Z39.50 protocol and XML. The section demonstrated that a key aspect of the development of the interactive 3D visualisation is the `Behavior` class offered by Java3D and its subclasses developed specifically for the STORM browser. This section has therefore discussed the development of the complete STORM system, including the browser, middleware and communication with servers. Subsequent sections will discuss the development of the GeoDOVE applet and how both STORM and GeoDOVE are integrated into the NEEDS geoportal.

5.3 Development of GeoDOVE

In Chapter 3 we discussed some of the limitations of existing 3D scene description languages and real-time 3D graphics engines. An approach based on a thick-client web-deployable Java3D applet was proposed. To investigate the proposed geovisualisation approaches, a prototype called GeoDOVE was developed as part of this study. The three main facilities that were identified for implementing GeoDOVE included the data upload, visualisation and exploration facilities. The user engages the aforementioned facilities in the following order:

1. From the NEEDS metadata page, a user clicks on a link to a dataset and is forwarded to the GeoDOVE-embedded webpage.
2. The applet's data upload facility automatically uploads the dataset that was described by the metadata page in the previous step.
3. As we cannot, at present, predict the schematic and semantic definitions of the contents of each dataset, the applet first asks the user if she wishes to extrude the dataset, and if so by how many distance units (generally metres).
4. The applet's visualisation facility initialises the 3D scene and automatically renders the dataset with the user's presentation settings already implemented.
5. Once the data has been rendered, the data exploration facility enables the user to access the associated table of attributes and assess the contents of the dataset.
6. Upon gaining a deeper understanding of the data, the user may wish to change the appearance of the presented data. The applet should support runtime modifications to the appearance of objects in the 3D scene.

5.3.1 Data Upload Facility

The data upload facility is responsible for the retrieval of geographic data and instantiation of Java classes that abstract the SFS. This means that it is the interface between heterogeneous DBMS, web services and the applet. The SFS model that is produced by the data upload facility is implemented using primitive Java data types, and hence does not include any Java3D classes. This was necessary to ensure that feature collections could be serialised and transmitted over the web, resulting in a three-tier client-middleware-server architecture. Important Java3D classes such as the `Geometry` class do not implement the `Serializable` interface and hence instances of these classes cannot be transmitted as Java objects over a network connection (Horstmann and Cornell, 2001). In a two-tier architecture where the client connects directly through the ODBC or JDBC, serialisation is not necessary as the feature collection is created on the client side. An initial design of GeoDOVE was based on the two-tier approach. Unfortunately, security restrictions from outside our University Intranet were restrictive to connections through any ports other than port

80 (the HTTP web server) and consequently, the architecture was re-designed to adopt a three-tier model; with the upload facilities implemented at middle-level. Both architectures are presented in Figure 5.8. The dashed line depicts the cordon between client and server. The illustration shows that the difference between these two implementations of GeoDOVE was characterised mainly by the position of the JDBC/Web Service connector which was part of the upload facility.

One of the most important design considerations was how to offer temporary storage for uploaded attributes. Although feature attributes could be easily stored in memory at runtime, it was necessary to offer the ability to query those attributes using standard methods of querying such as SQL. The applet was therefore designed to include an internal RDBMS. The internal RDBMS would offer SQL support for feature collections retrieved from data sources that may not support SQL-based querying. The RDBMS uses the popular HSQLDB, an open source database engine completely developed in Java (HSQLDB Development Group, 2005). The HSQLDB database is deployed with the rest of the GeoDOVE Java libraries when the applet web page is visited. A relational table is then created for each feature collection that is uploaded in the applet. A later section discusses the querying of attribute tables. Although the current implementation does not offer the creation of relational joins between uploaded feature collections, the HSQDB database supports relational joins and these may be included in future versions of GeoDOVE.

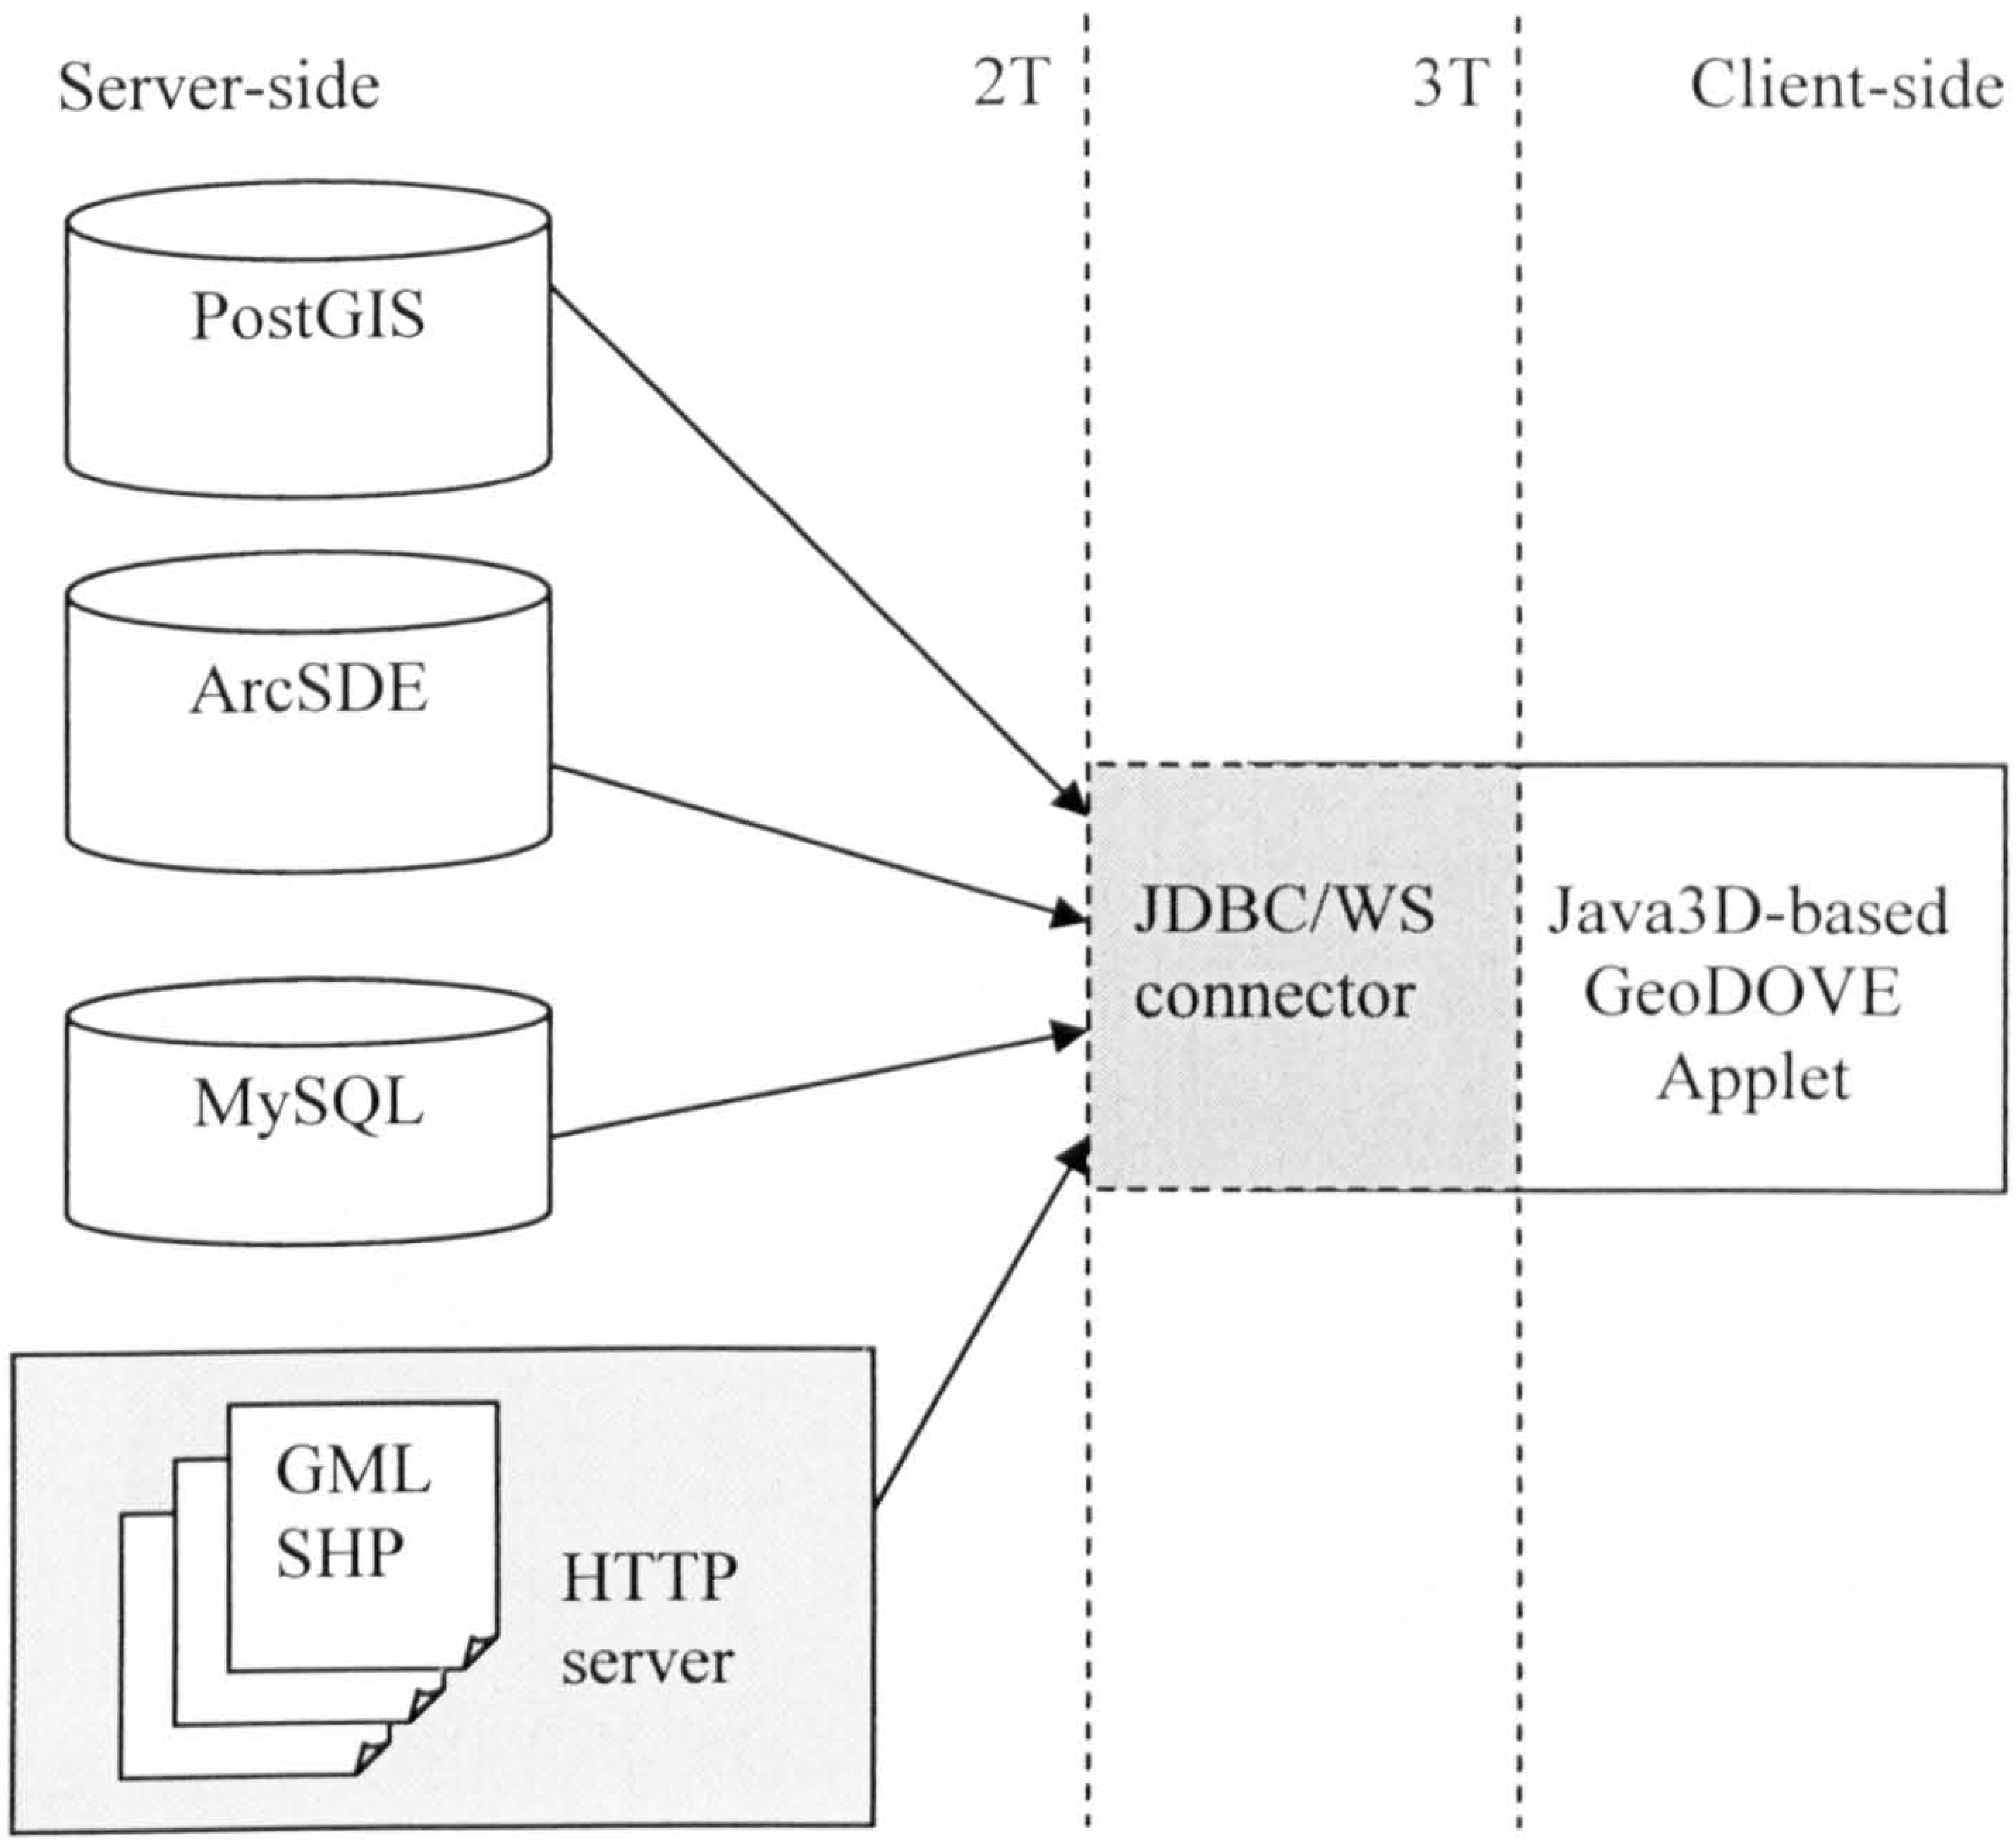

Figure 5.8 The GeoDOVE architecture

5.3.2 Visualisation Facility

In Chapter 3 we proposed an approach for abstracting OGC simple features using Java3D geometries. In this section we describe how the mechanisms for adding those Java3D geometries were implemented. The section is concerned with how objects are retrieved from the upload facility and added to the visualisation facility. Further, this section is concerned with how the scene is initialised to prepare it for user interaction. A screenshot of a geovisualisation from GeoDOVE is presented in Figure 5.9. The illustration shows buildings with textures attached to walls and an additional texture draped over the DEM as well.

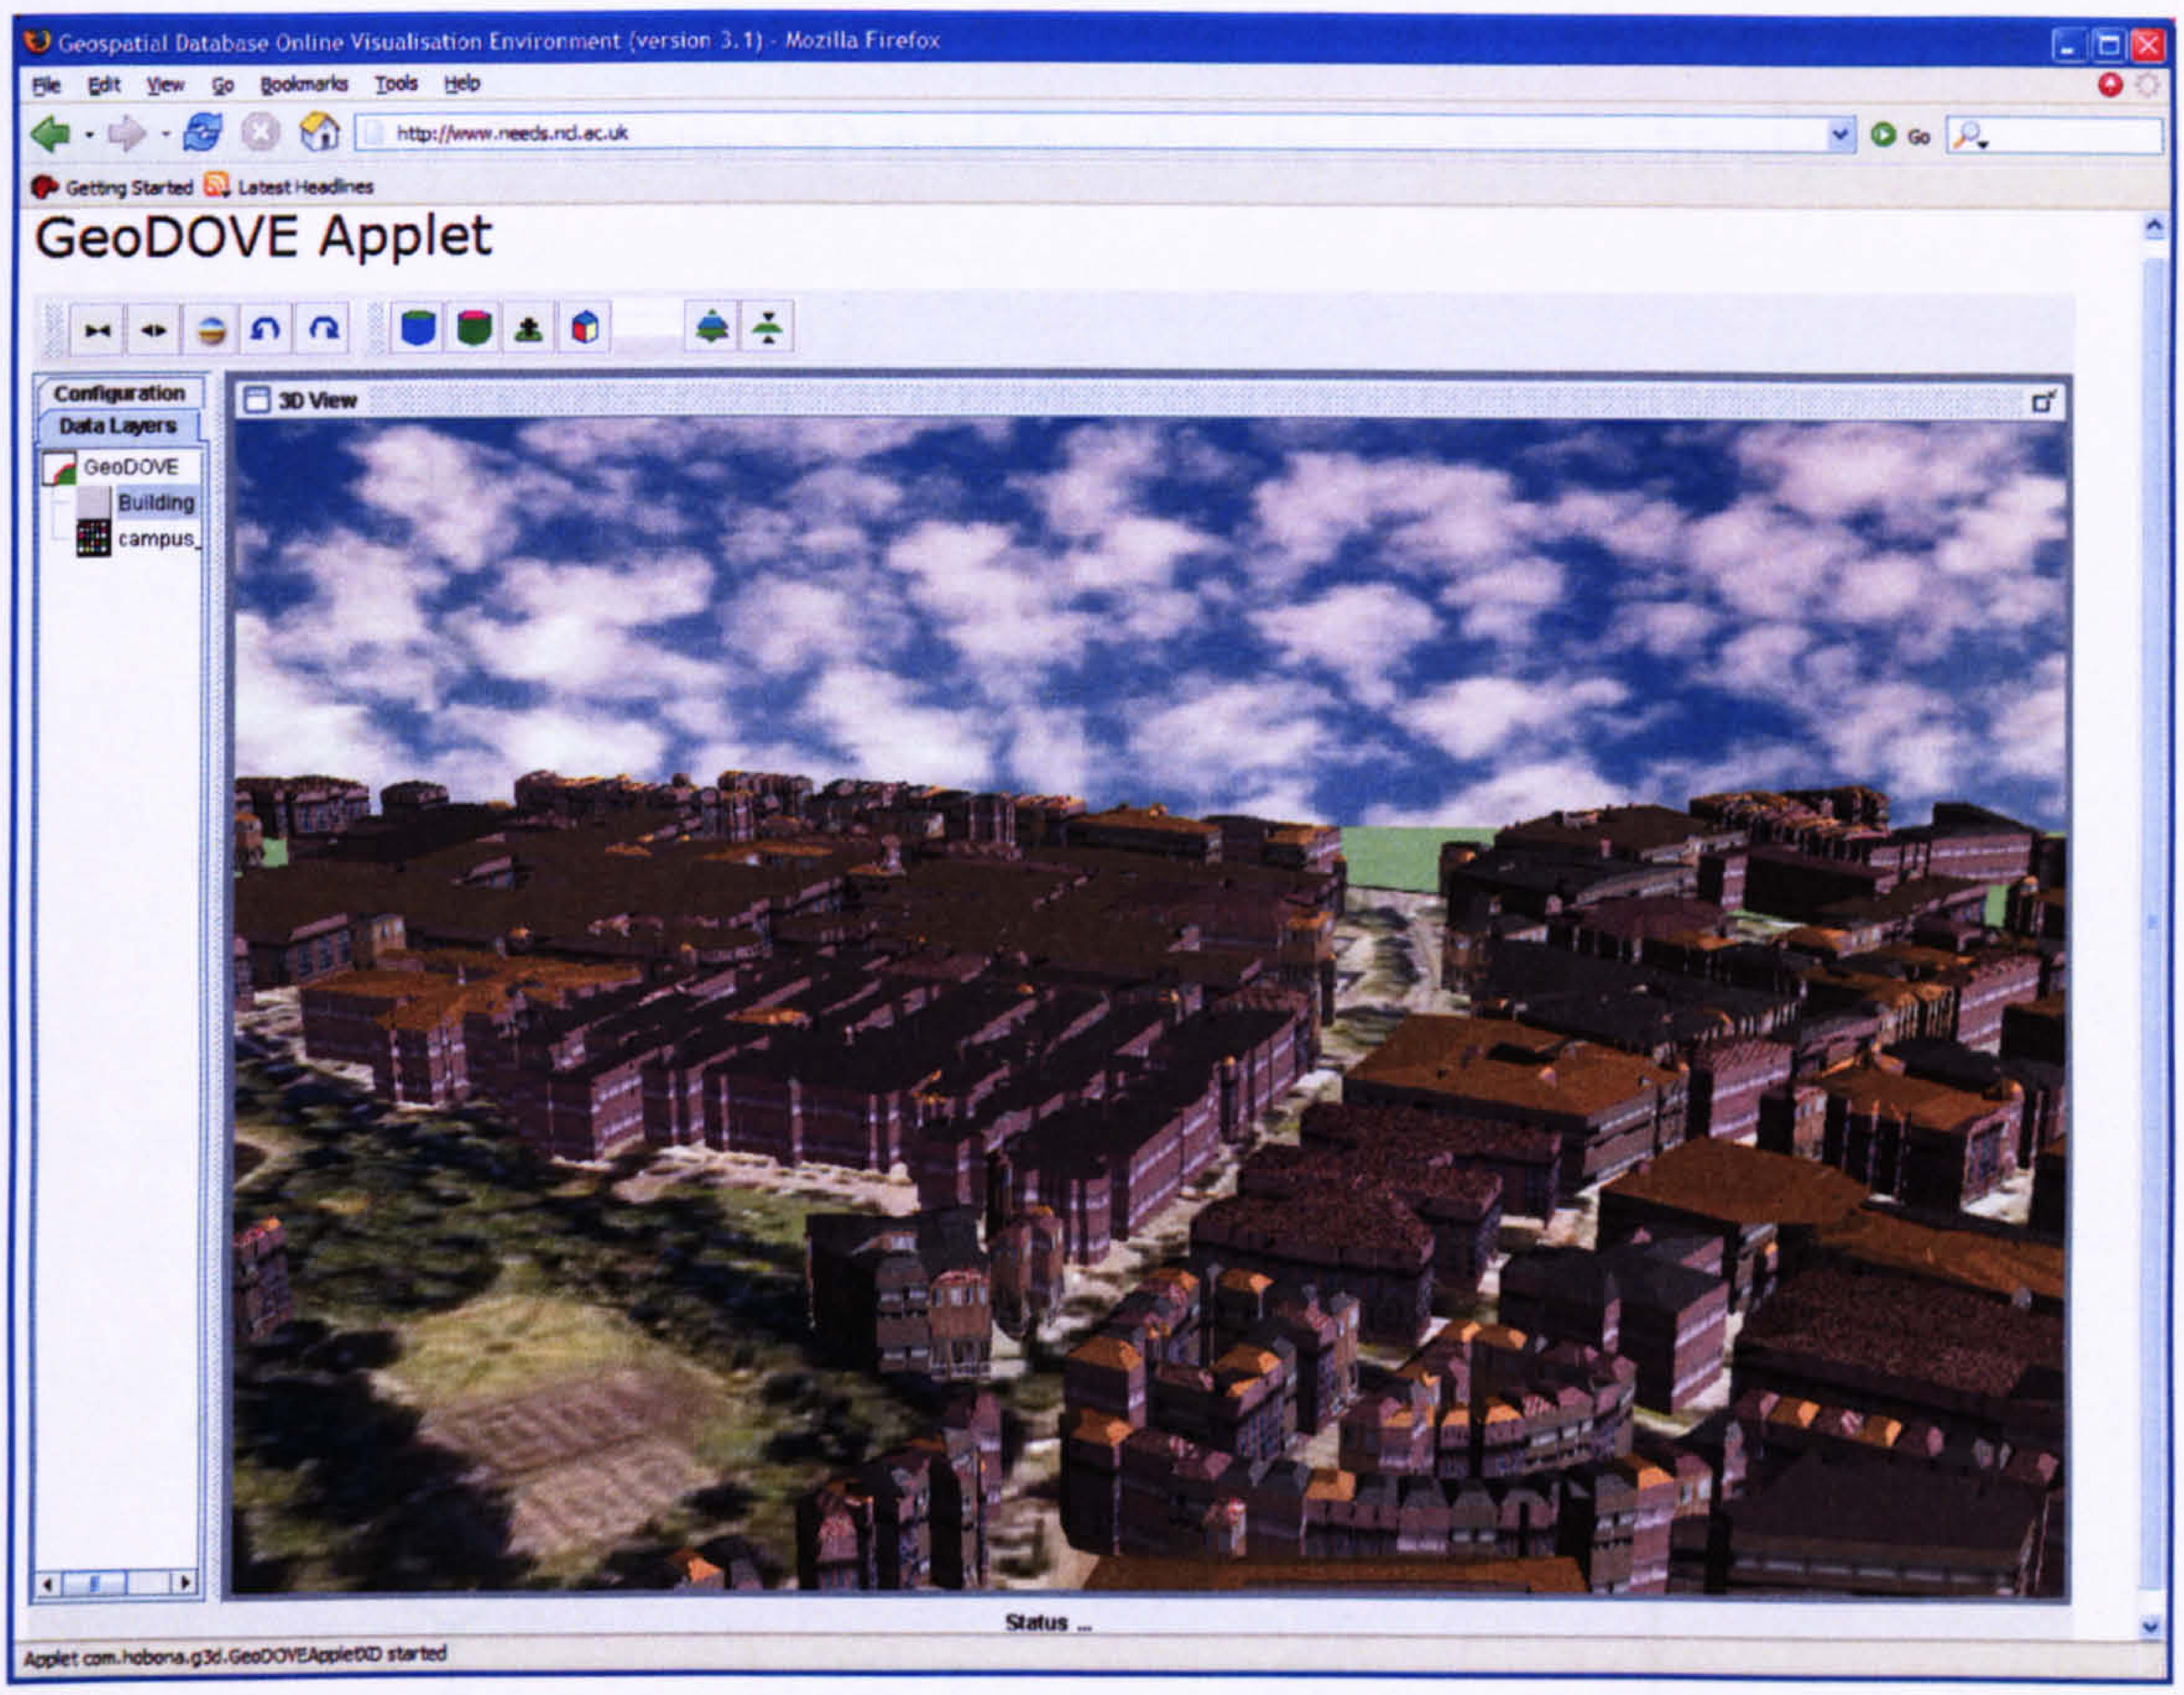

Figure 5.9 A screenshot of a geovisualisation from GeoDOVE. Ordnance Survey © Crown Copyright

5.3.2.1 Adding OGC Features to Java3D Scene-Graphs

Once a feature collection—an abstraction of the SFS `GeometryCollection` class— has been created by the data upload facility, the feature collection is then added to the instance of the `GeoPanel3D` class. This object handles all the visualisation processes including the 3D modelling of `Point`, `LineString` and `LinearRing` objects as per the procedure described in Section 3.4. The `GeoPanel3D` object retrieves metadata describing the types of SFS geometries held in the feature collection, the number of features and the spatial extent of the feature collection. If the feature collection holds `Polygon` objects, a loop is invoked that obtains the collection of `LinearRing` objects from each `Polygon` object. The collection of `LinearRing` objects is then used to create an instance of an `ExtrudedLinearRing` class, which then creates a 3D model of the

LinearRing from Java3D Shape3D objects. Figure 5.10 illustrates the post-upload workflow for creating 3D models within the GeoPanel3D object.

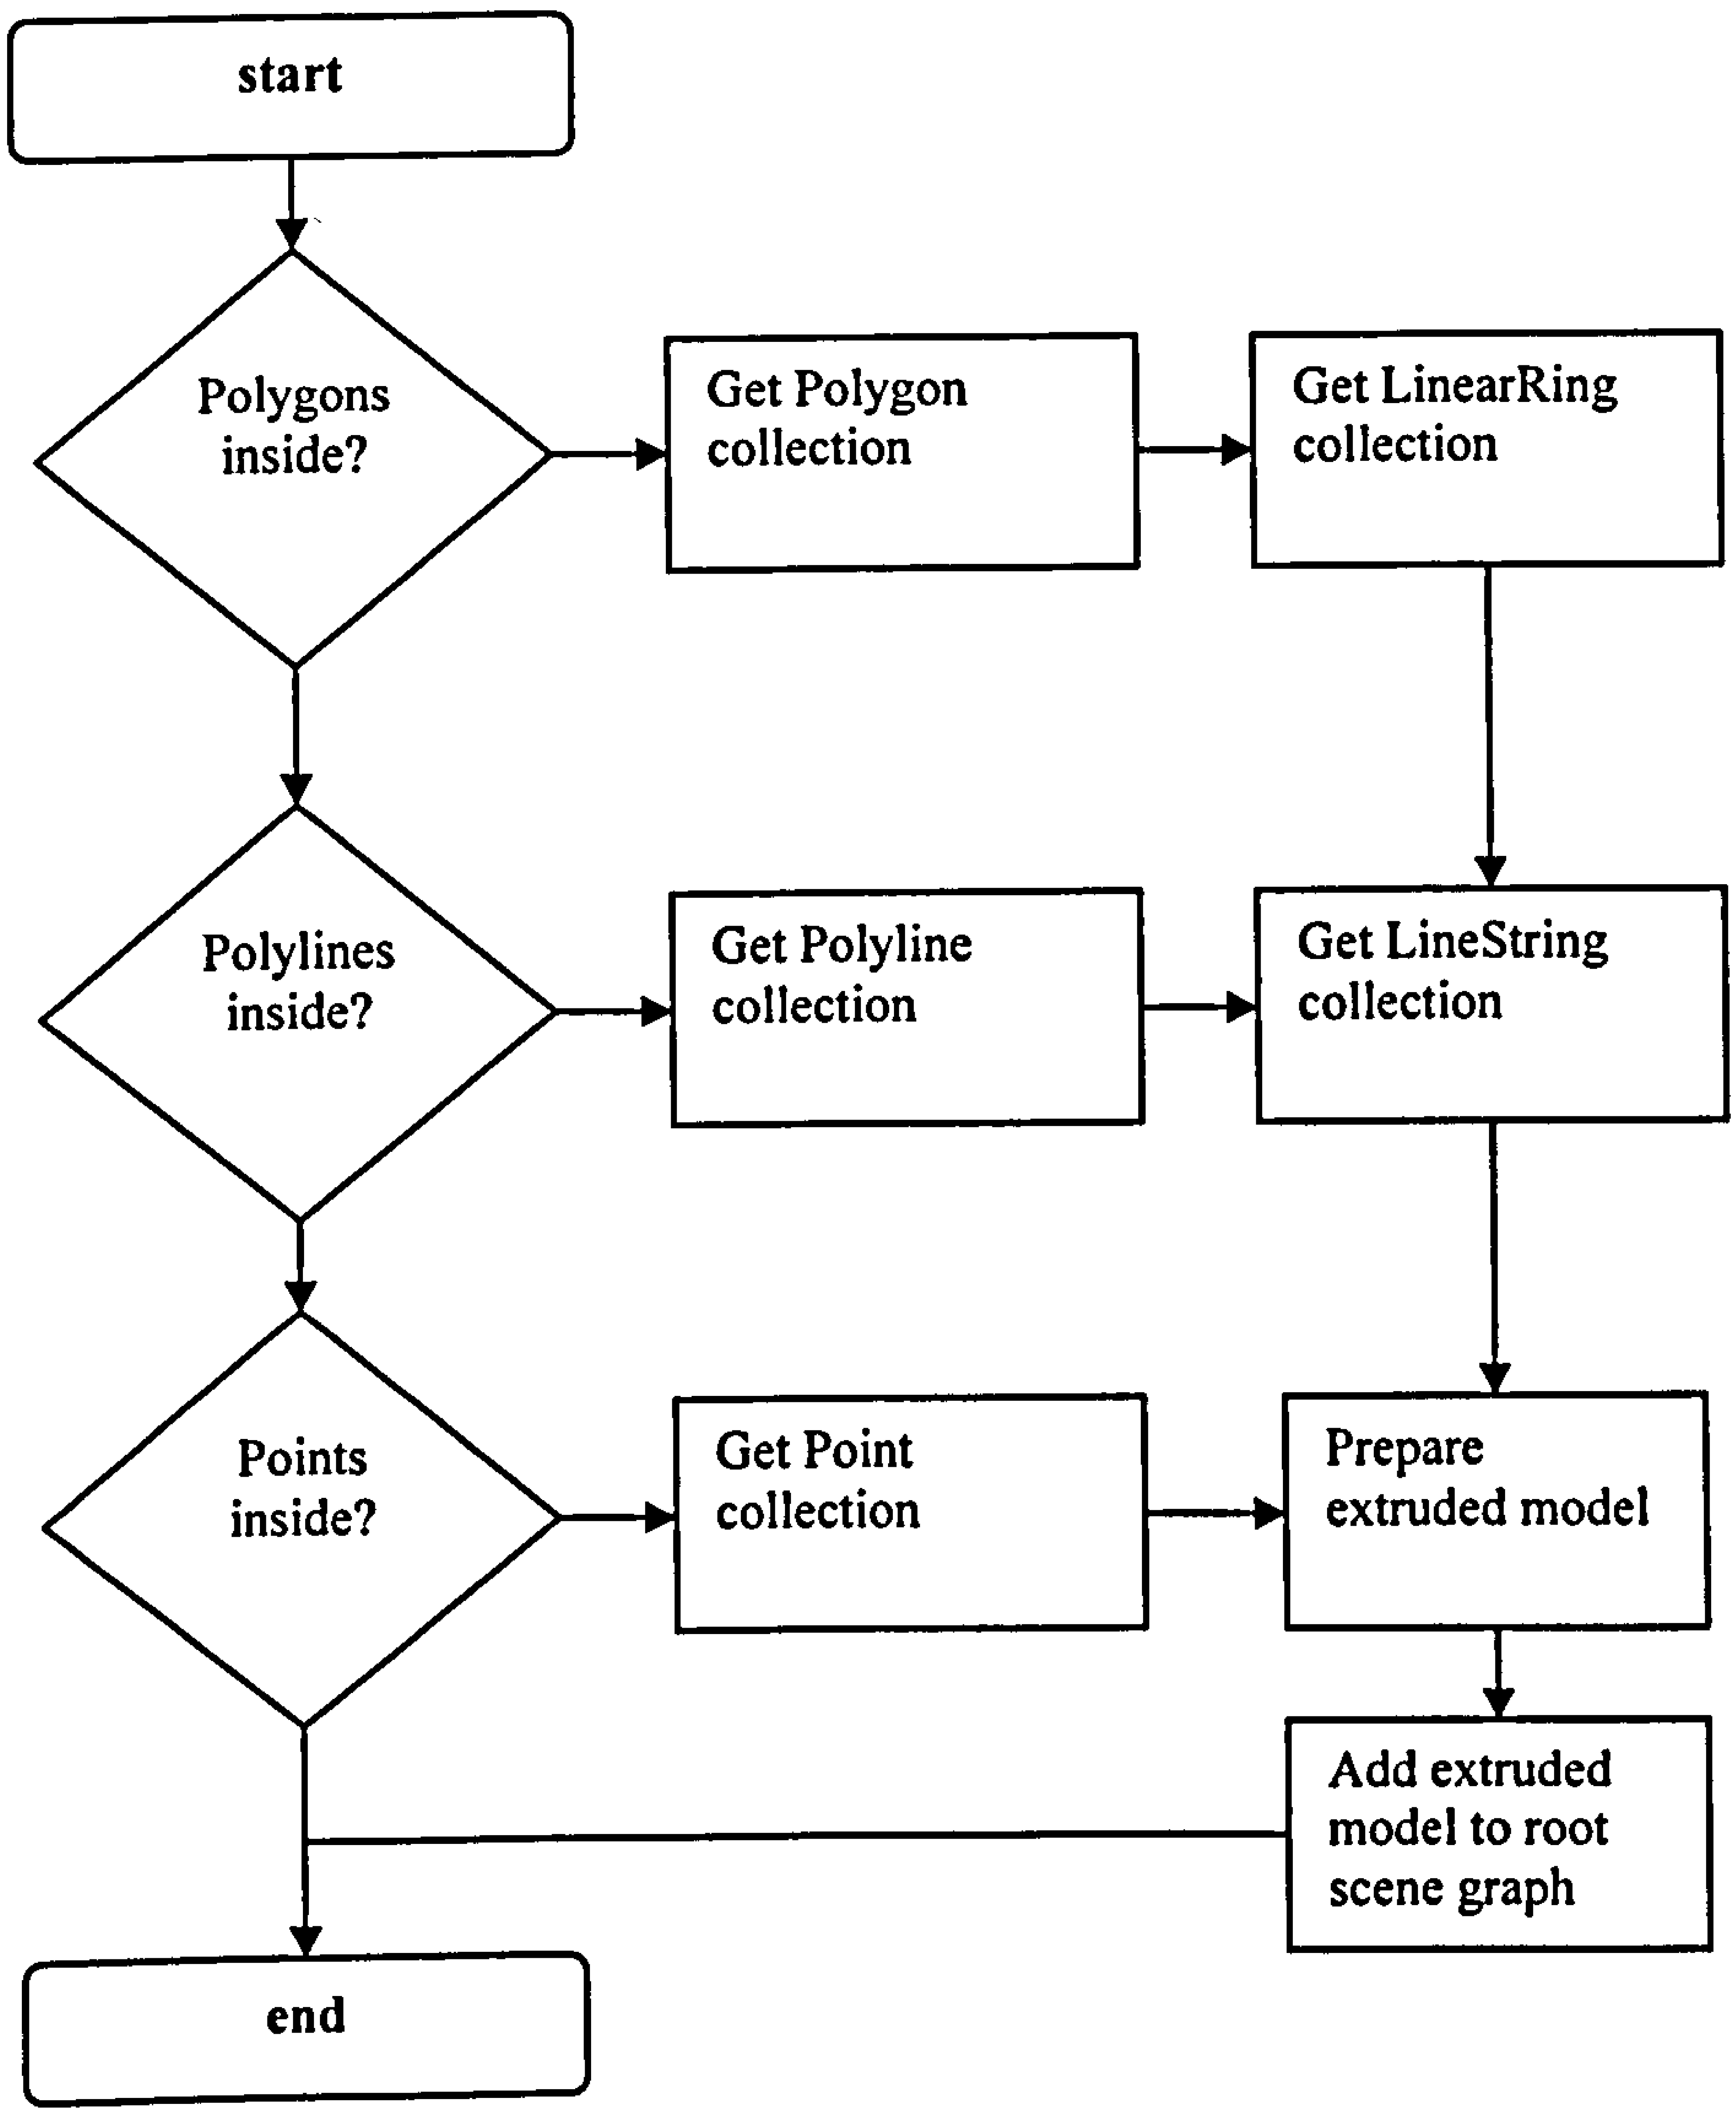

Figure 5.10 Flowchart of addition of SFS geometries into a 3D scene

If the GeometryCollection holds instances of the Polyline class (an abstraction of the SFS MultiLineString) then from each Polyline the collection of LineString objects is retrieved. An ExtrudedLineString is then created from the collection of LineString objects using the procedure described in Section 3.4. An obvious difference between an ExtrudedLineString and an ExtrudedLinearRing is that objects of the former class implement a roof structure whereas objects of the latter do not.

However, an additional difference is that, unlike the ExtrudedLineString, the back of the ExtrudedLinearRing is culled (not rendered) as it is never visible due to the roof overhead. Visualisation applications apply culling techniques to optimise the speed and efficiency of real-time 3D rendering. An alternative approach for abstracting LineString objects is through the creation of Java3D LineArray objects to represent features such as pipe-works or other tubular structures. By default, if the user chooses not to extrude a collection of LineStrings, then the geometries are modelled as tubular LineArray objects. Finally, the user is given the opportunity to specify the width of each LineArray.

As illustrated in Figure 5.10, Point objects are also retrieved from a GeometryCollection. If the user wishes to place a cylinder at the location of the point, then an ExtrudedLinearRing is created with vertices along the circumference of the point. Alternatively, if the user selects to create a spherical representation at the location of that point then a Java3D Sphere object is placed at the location. The applet also allows the user to specify the radius of the spherical or cylindrical objects. The Java3D Sphere and the Shape3D objects created from instances of the Point, Polyline and Polygon classes are then added to instances of the Java3D TransformGroup class. The TransformGroup class allows for the translation, rotation and scaling of coordinates, including the positioning of models in the scene. The TransformGroup object is then added to a Switch object, which offers a toggle to enable or disable rendering of child objects. Each Switch object is then added to a BranchGroup which represents a GeometryCollection. The BranchGroup class has the special property that it is the only subclass of the Group class that can be added to a compiled and live scene. Figure 5.11 illustrates the resultant scenegraph.

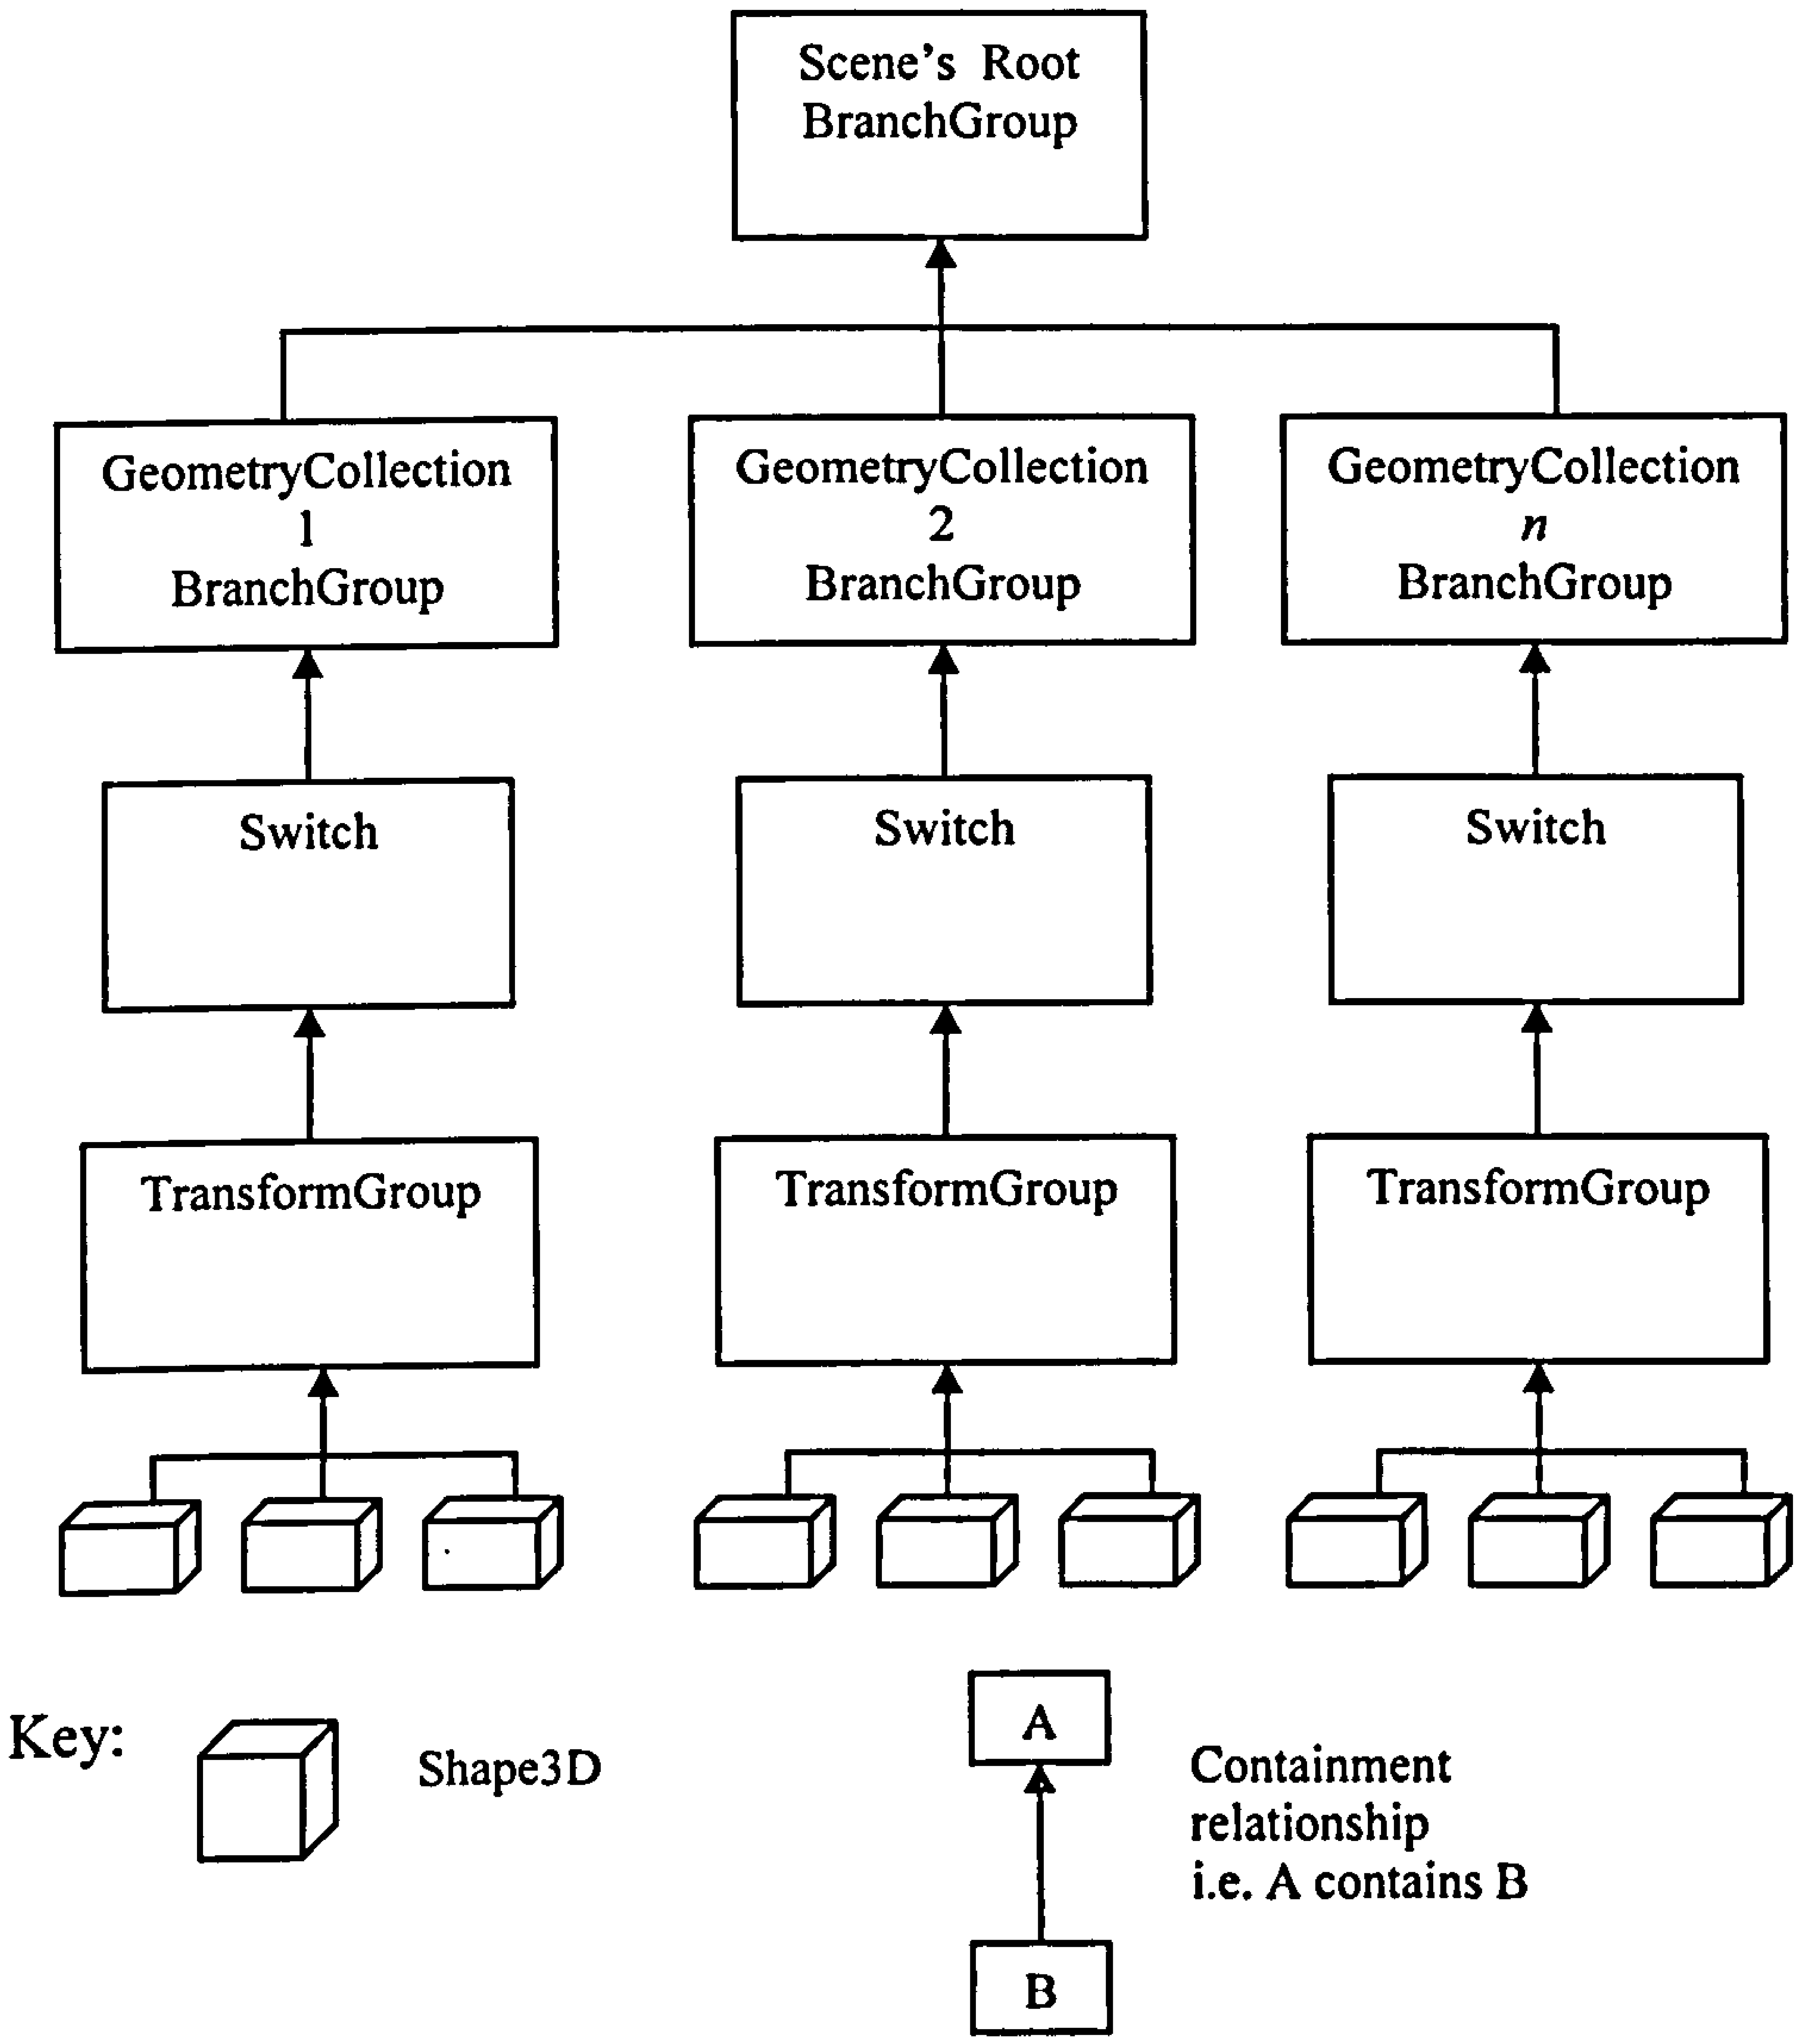

Figure 5.11 The GeoPanel3D scene-graph

So far we have discussed only how the SFS geometries are modelled, how they are retrieved and how they are added to a 3D scene (earlier in this section). An important design issue was the handling of colours and appearance of objects in the visualisation. "About four percent of the public (8.0% of men and 0.4% of women) are congenitally colour deficient"(Robinson et al., 1995: pp. 400). It was therefore necessary to include dynamic colouring and texturing in our design. Each Java3D `Shape3D` object is instantiated through the passing of two parameters, a `Geometry` and an `Appearance` object. Whereas the `Geometry` objects are created from SFS abstractions as discussed in Section 3.4, the appearance of each `Shape3D` is handled after the objects have first been uploaded. The appearance is defined through the specification of a `Material` or a `Texture` object. A `Material` object offers a homogeneous colour fill over a shape, whereas a `Texture` object allows an image or bitmap to be draped on geometry. Subclasses of the `Behavior` class were therefore implemented and set to drape an image on an extruded SFS geometry as a `Texture` object. In the following text, subclasses with a '`vior`' suffix are part of the Java3D API, those with a '`viour`' suffix have been implemented specifically for GeoDOVE.

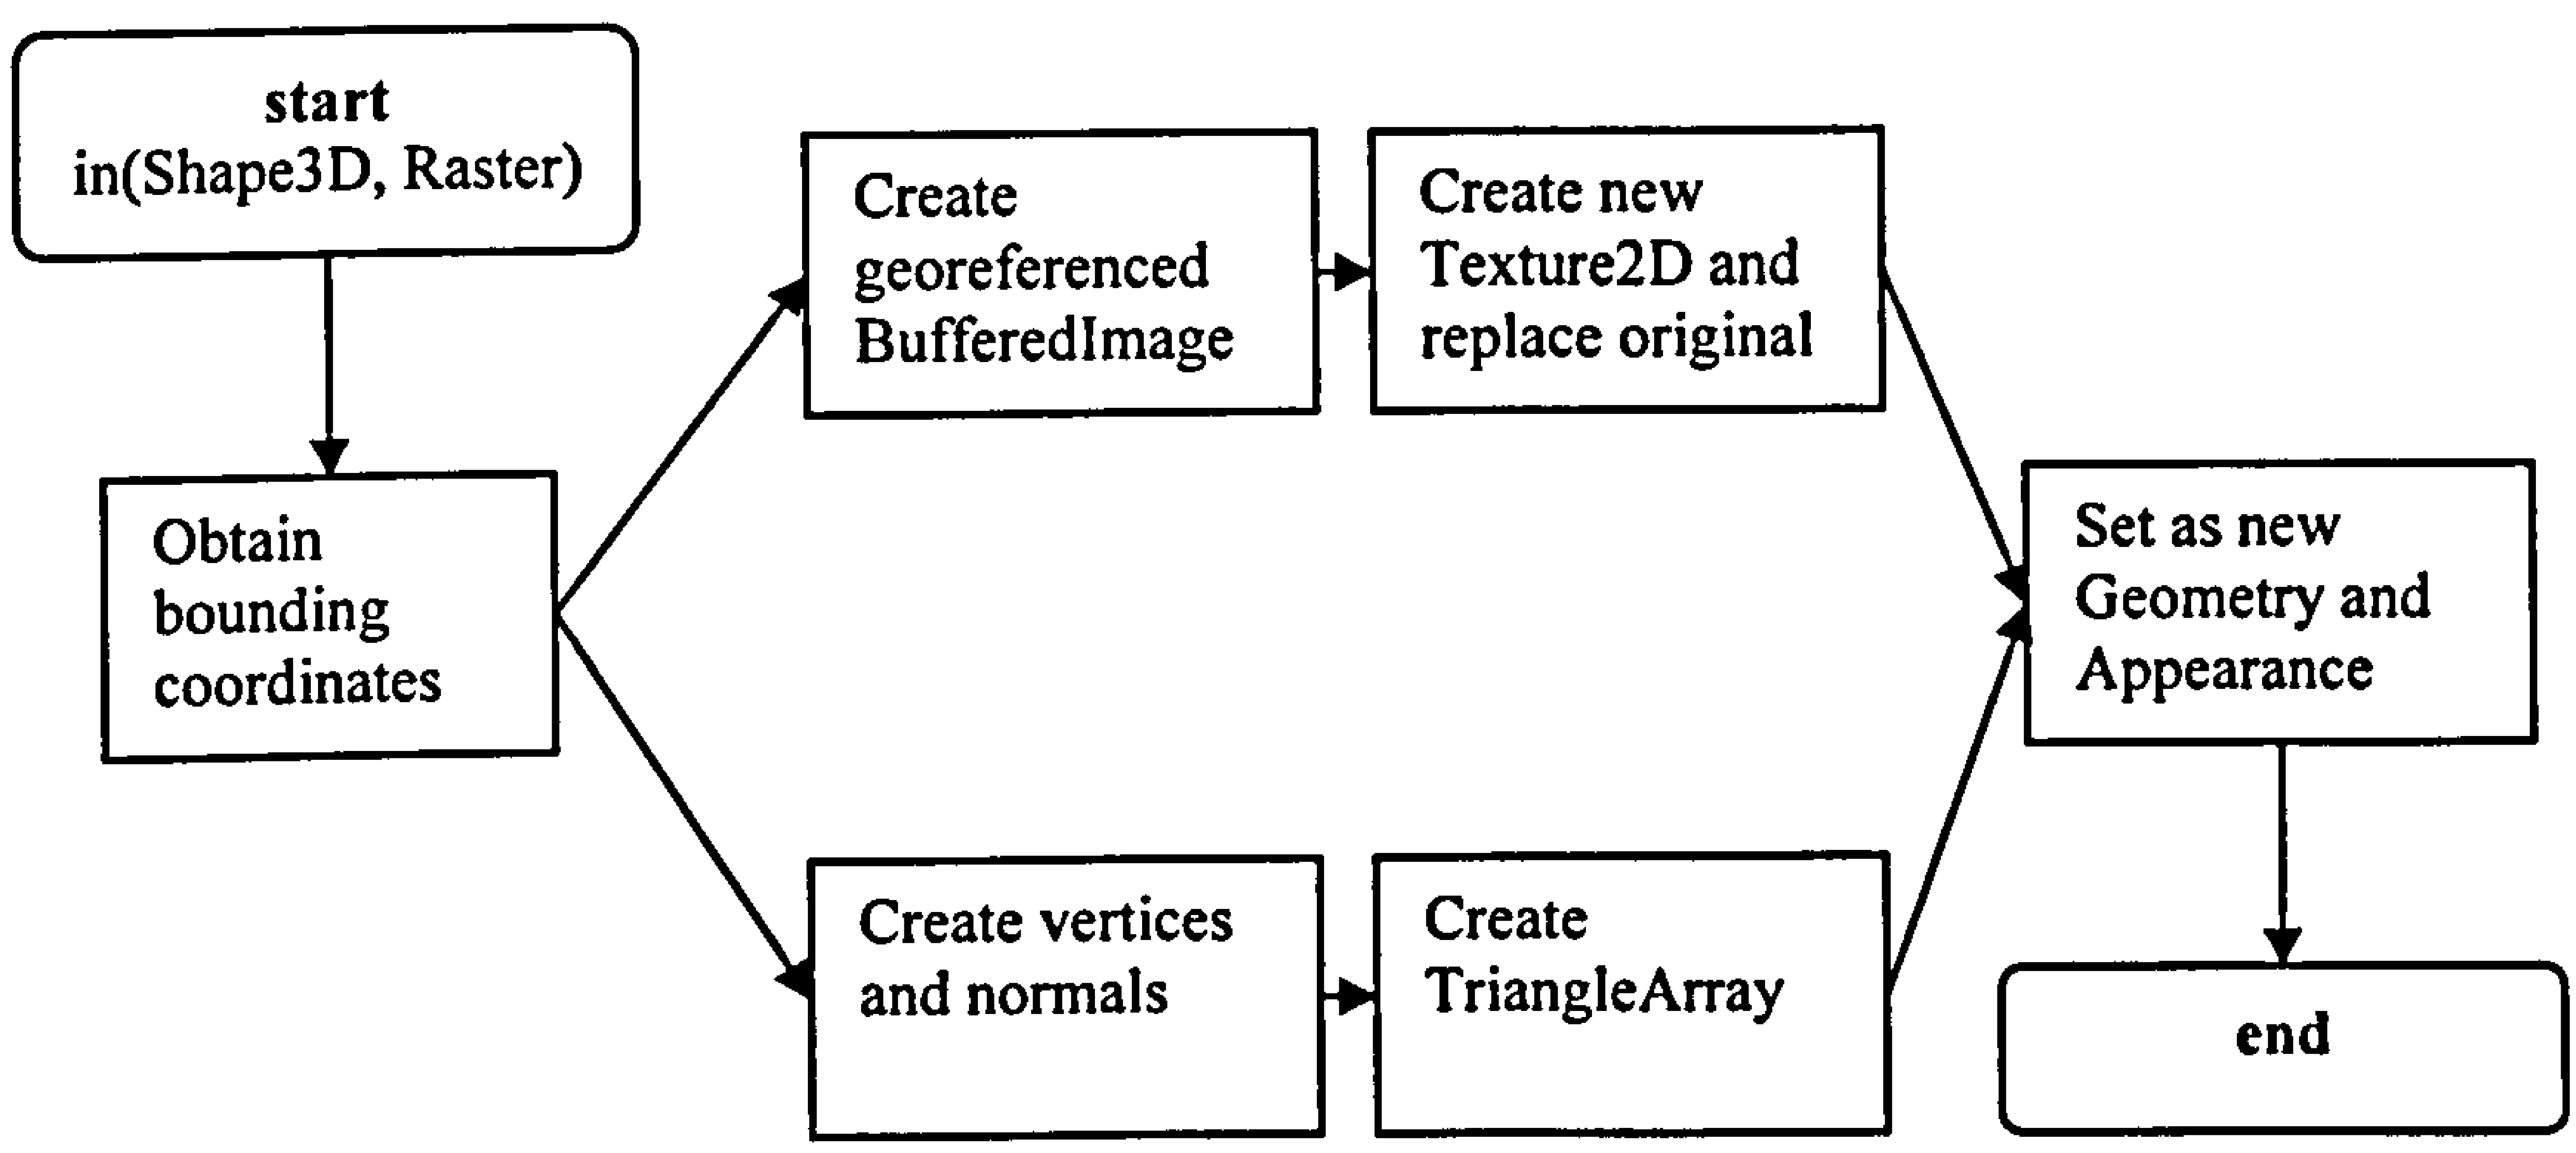

Figure 5.12 Flowchart of the DEM creation process

A separate behaviour subclass, called `DEMBehaviour`, was implemented for draping imagery over a DEM. The input parameters were the target `Shape3D` object and the input georeferenced raster dataset (an instance of the `GeoRaster` class). When the `DEMBehaviour` class is triggered, the process illustrated in Figure 5.12 takes place. First, the spatial extents and resolution of the input raster are retrieved. A regular grid of Java3D `Point3f` vertices covering the spatial extents and spaced according to the just-extracted resolution is created. Then the corresponding normals are created for each `Point3f` vertex constituting the DEM. An instance of the `TriangleArray` class is then created and set as the geometry of target `Shape3D` object. For the appearance of the target `Shape3D` object, the bounding coordinates of the input raster are used to geo-reference an instance of the Java image-handling class `BufferedImage`. Once the `BufferedImage` has been geo-referenced, a second input raster is then rendered on to the geo-referenced `BufferedImage`.

The second input raster may be a colour-coded version of the first raster dataset or an image that was draped on the target `Shape3D` object when the process was invoked.

5.3.2.2 Scene Initialisation

The scene is initialised by adding behaviours to enable user navigation within the 3D environment. Navigation is implemented by altering transformation properties of a Java3D `TransformGroup` object. Java3D offers a `KeyNavigatorBehavior` class for altering `TransformGroup` objects through a keyboard. It also offers `MouseZoom`, `MouseTranslate` and `MouseRotate` classes for altering `TransformGroups` through actions on a mouse device. During initialisation of a scene in GeoDOVE, the root `TransformGroup` of the viewing platform (the location of the viewing camera) is added to a `KeyNavigatorBehavior` object, which is then added to the root `BranchGroup`. This allows the user to navigate around a 3D scene through key presses. Rather than using the `MouseZoom`, `MouseTranslate` and `MouseRotate` classes to navigate using the mouse device, our study adopted the `OrbitBehavior` class. The `OrbitBehavior` class encompasses all 3D transformations (zooming, translation and rotation) in one class; in contrast to the aforementioned '`Mouse`'-prefixed classes which offer a single type of transformation per class.

In addition to the navigation behaviours, the `GeoPanel3D` class also prepares view attributes such as the background appearance and the initial position (and orientation) of the viewer. Java3D offers a `Background` class that allows the developer to place an image in the background of the visualisation. Unfortunately, the background image also does not move when the position of the viewing camera changes, which is not a realistic abstraction of spatial phenomena. An alternative is to create a spherical object, using the `Sphere` class, to contain the rest of the scene. This option allows all objects in the scene including the background to move relative to the viewing camera. The second option was implemented for the final version of the applet.

Figure 5.13 Orientation in 3D space illustrated with views of a teapot from above (a) the side (b) and the front (c)

Defining the initial position and orientation of the viewing camera is an important part of the pre-visualisation process. The scene is initiated with the initial position of the viewing camera located at the centre of the first uploaded dataset. In addition to the position, we consider the orientation of the viewing camera. In 3D geovisualisation, the viewing camera is allowed to rotate both horizontally and vertically, thereby altering its yaw and pitch respectively. Figure 5.13 illustrates orientation in 3D space through the alteration of yaw, pitch and roll. Similar to OpenGL and Direct3D, orientation in Java3D is defined through the specification of a quaternion – a four parameter set of numbers – that specifies an axis in 3D space and an angle of rotation about that axis. Java3D adopts a default coordinate system with the Z axis increasing towards the viewing camera. GeoDOVE therefore alters the orientation such that the z-axis (representing the height above ground level) points upwards when the scene is initialised.

5.3.3 Exploration Facility

Once the `GeometryCollection` has been added to the root `BranchGroup`, the user is then able to explore the 3D scene. Visual exploration capabilities implemented included a behaviour class, `LayerSeparator`, which separates the layers vertically into a layer stack. The behaviour may be triggered by a key press or through clicking a menu button. The menu button dispatches an artificial key press through the `dispatchEvent` method that the `GeoPanel3D` inherits from Java's GUI API. When triggered the `LayerSeparator` object traverses the scenegraph illustrated in Figure 5.11 and checks if each node is a `TransformGroup` representing a feature collection. If the behaviour class finds such a `TransformGroup`, the reference number (*i*) of the feature collection is obtained from the `TransformGroup`'s metadata and used to translate the `TransformGroup` by adding a vertical distance of *i* x *10*. This means that each of the feature collections is translated according to whether it was uploaded first, second, third and so on. The behaviour class also offers a reverse function that re-integrates the multiple layers together by subtracting *i* x *10* from the height value. Figure 5.14 shows roads (in green), buildings (in purple) and laser-scanned data layers separated into a layer stack, including a background image of the sky with some clouds.

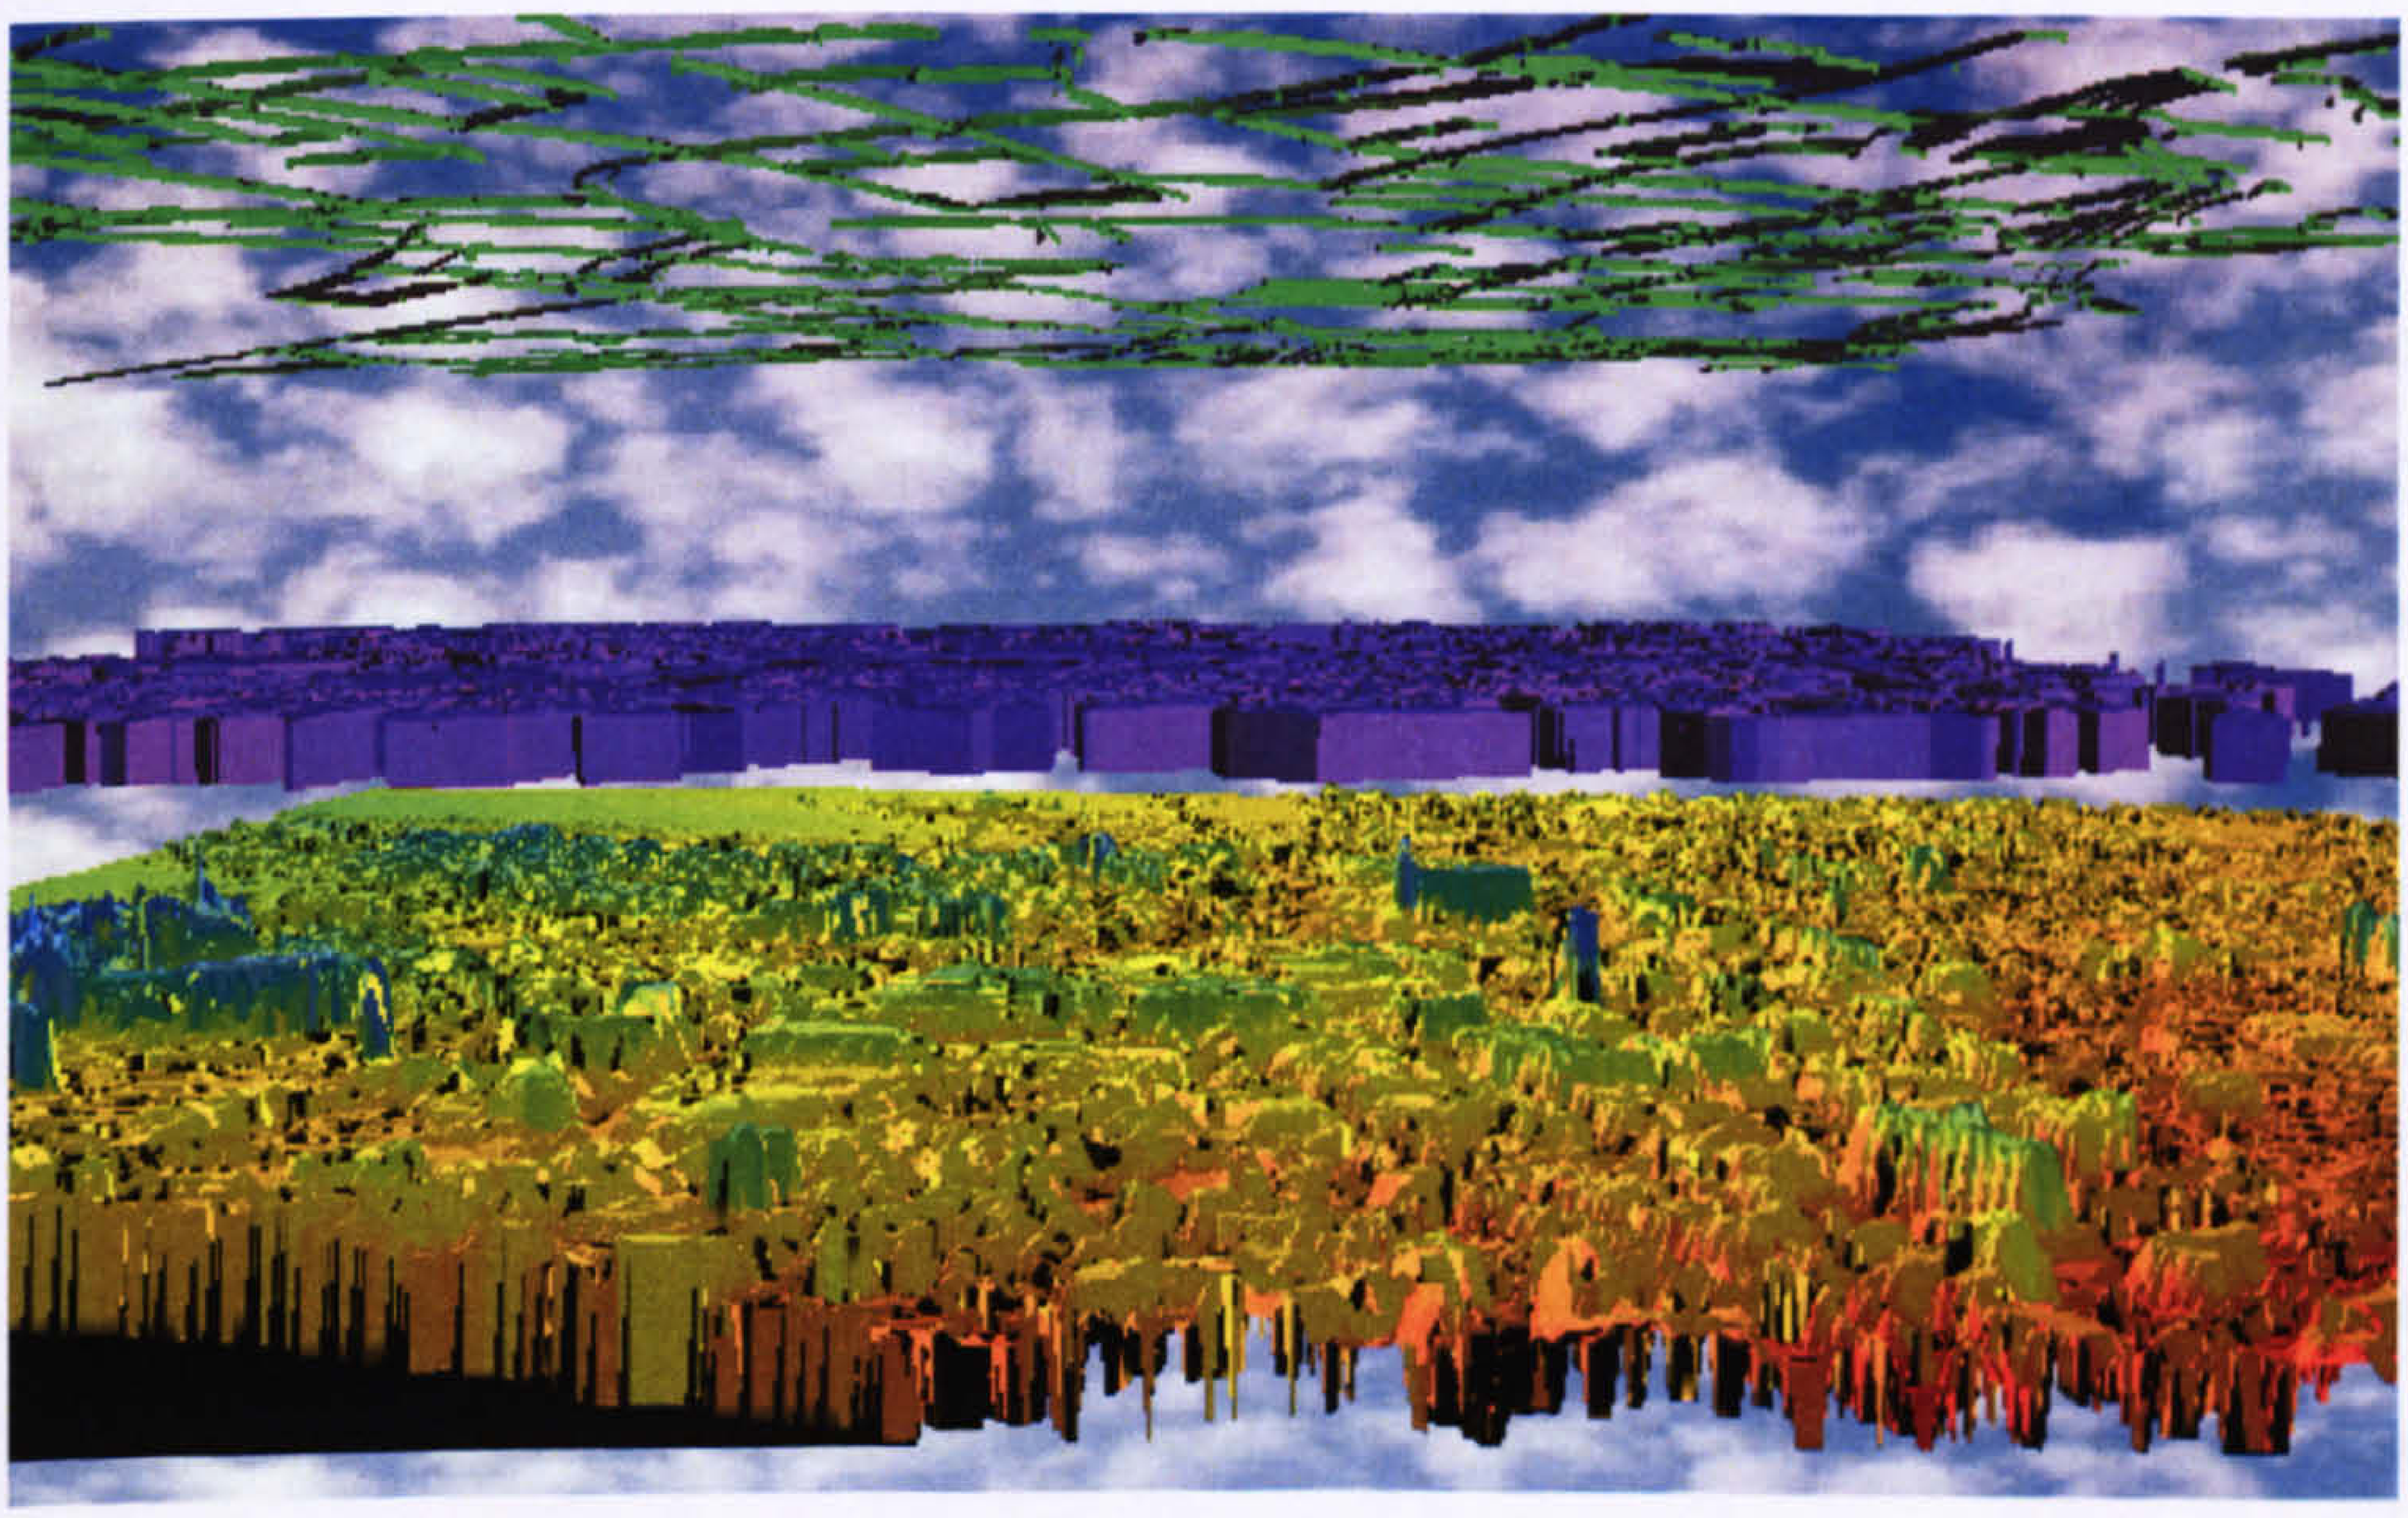

Figure 5.14 A layer stack created by the layer separation behaviour

Each feature's attributes can also be explored by highlighting the feature's record on an attribute table. A record can either be highlighted by selecting the associated geometric object in the 3D view, by selecting the associated record in the attributed table or by invoking an SQL query that selects the associated records. A geometric object representing a feature is selected using a `SelectionBehaviour` class, which extends the Java3D `PickMouseBehavior` class. The `SelectionBehaviour` is triggered when a user clicks on the `GeoPanel3D` panel. An instance of this class is added to the root `BranchGroup`. The `PickMouseBehavior` class and other subclasses of the Java3D `Behavior` class offer a mechanism for transmitting messages between the Java runtime and the lower-level OpenGL/Direct3D graphics engines. Different triggers (called `WakeUpOn` events) may be set to invoke an action to be performed by a `Behaviour` subclass.

Alternatively, the user may choose to select objects by clicking on a record in the attribute table. In this scenario, the `TableSelectionBehaviour` class handles the selection events. The attribute table passes the unique identifier of the selected record, which is then added to the list of selected identifiers in the `TableSelectionBehaviour` object. Unlike the `SelectionBehaviour` class, this class is triggered by elapsed time; specifically it is automatically triggered every half a second. It therefore, behaves like Java's built-in listener classes which continuously monitor events on GUI components. The `TableSelectionBehaviour` class also allows for the selection of objects through the execution of an SQL statement such as *SELECT * FROM MyTable WHERE UniqueID > 100* (that is, to select all features with an attribute called UniqueID that has a value greater than 100). The SQL query returns a list of records and their unique identifiers are passed to the `TableSelectionBehaviour` which then highlights the associated objects in the 3D scene.

Similar to the development of the STORM browser, the Behavior class played a key role in the implementation of the GeoDOVE applet. However, an additionally important component was the Java database connectivity API (JDBC). This section described, in detail, the processes through which the GeoDOVE applet retrieves geometry collections from a geospatial database and converts them into Java3D geometries. The next section describes how the STORM browser and the GeoDOVE applet are integrated into the NEEDS geoportal.

5.4 Developing the Geoportal

The previous two sections discussed the development of two visualisation applications. The first application is specifically for the visualisation of geographic search results. The second application is for the visualisation of multidimensional geographic data. With the visualisation software implemented, the next task was to integrate them into a single geoportal. For this to be achieved it was necessary to base the design of the geoportal on the OGC geoportal architecture presented in Figure 1.2 and the model suggested by Tait(2005) illustrated in Figure 1.3. The geoportal would have to support the Z39.50 protocol which is supported by over 400 spatial data clearinghouses distributed globally(FDGC, 2005). The architecture illustrated in Figure 5.15 was designed for the implementation of the NEEDS geoportal.

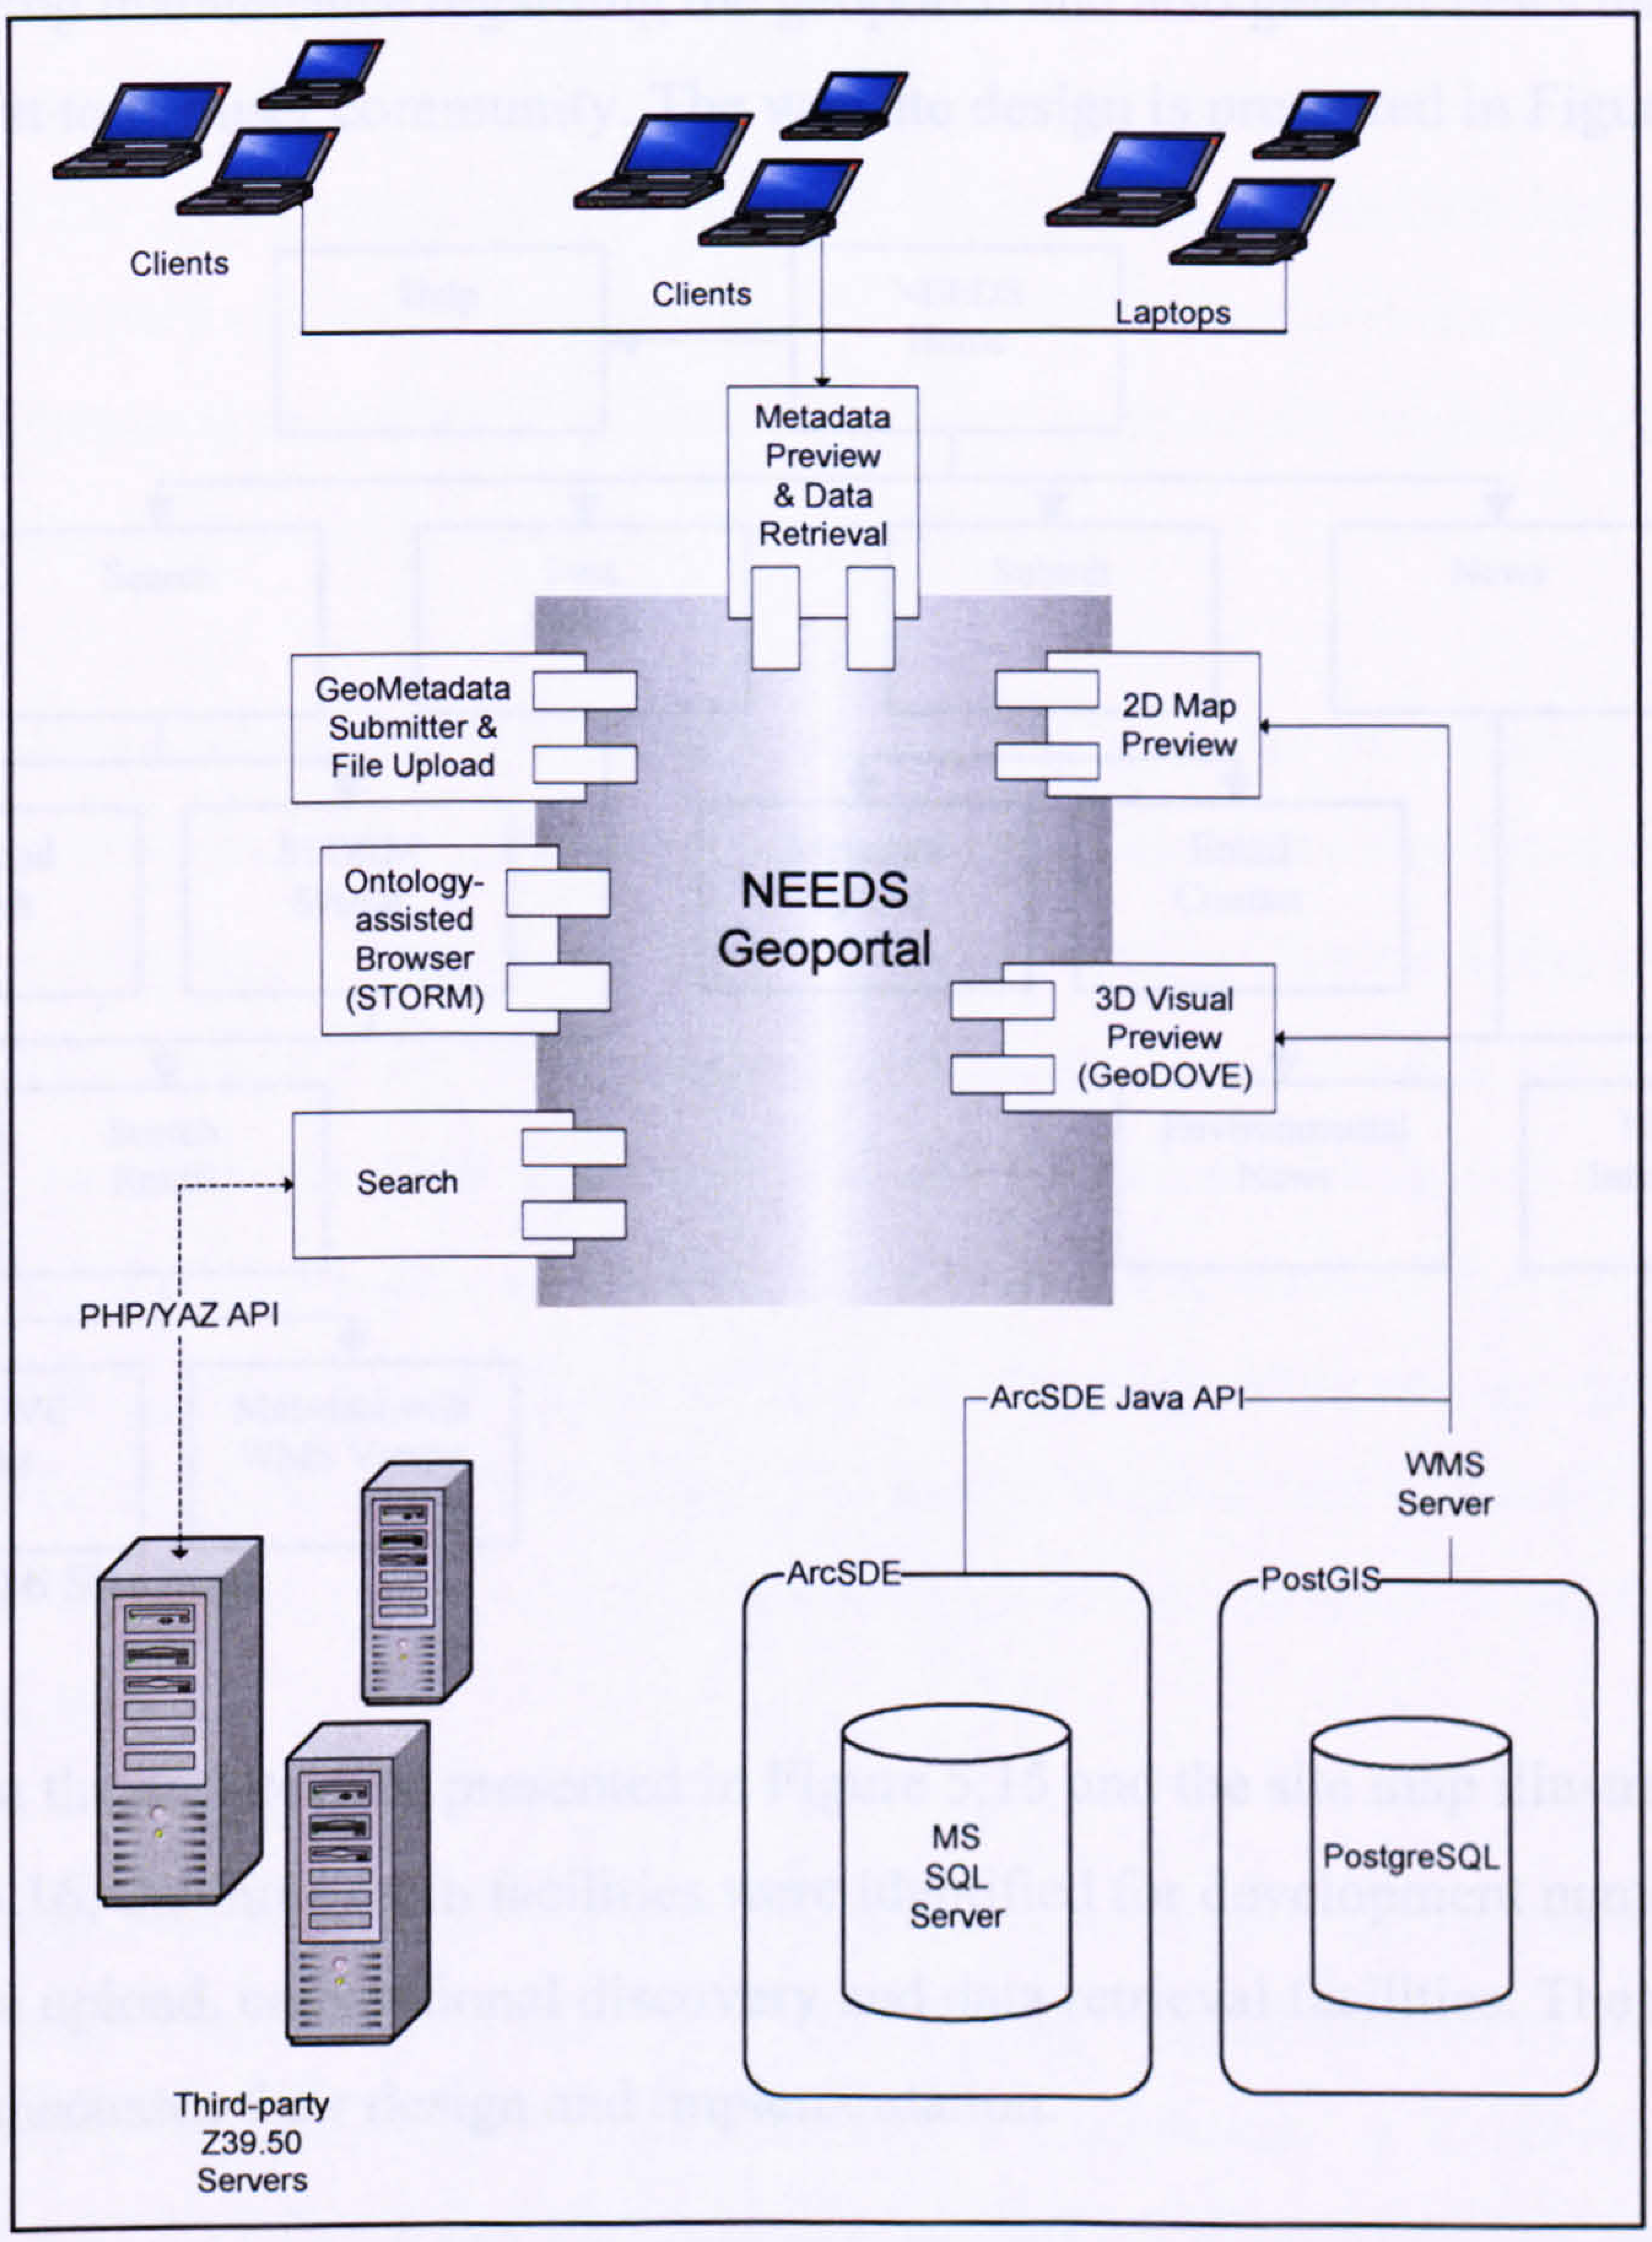

Figure 5.15 NEEDS architecture

The NEEDS geoportal is the platform from which the STORM browser and GeoDOVE are accessed. As these visualisation applications are primarily for enhancing geographic data discovery, they were made accessible from the geoportal's search pages and menus. Nielsen (2000) proposes that an intranet portal should always have three components: the directory, search and news sections. This thesis equates the directory to a catalogue as they both publish the existence and location of content. In addition to the search pages, the website structure was designed to include an upload section that included the metadata upload interfaces. The design of the NEEDS geoportal therefore included a news section for announcing maintenance regarding the geoportal and also general news that may be of interest to the user community. The website design is presented in Figure 5.16.

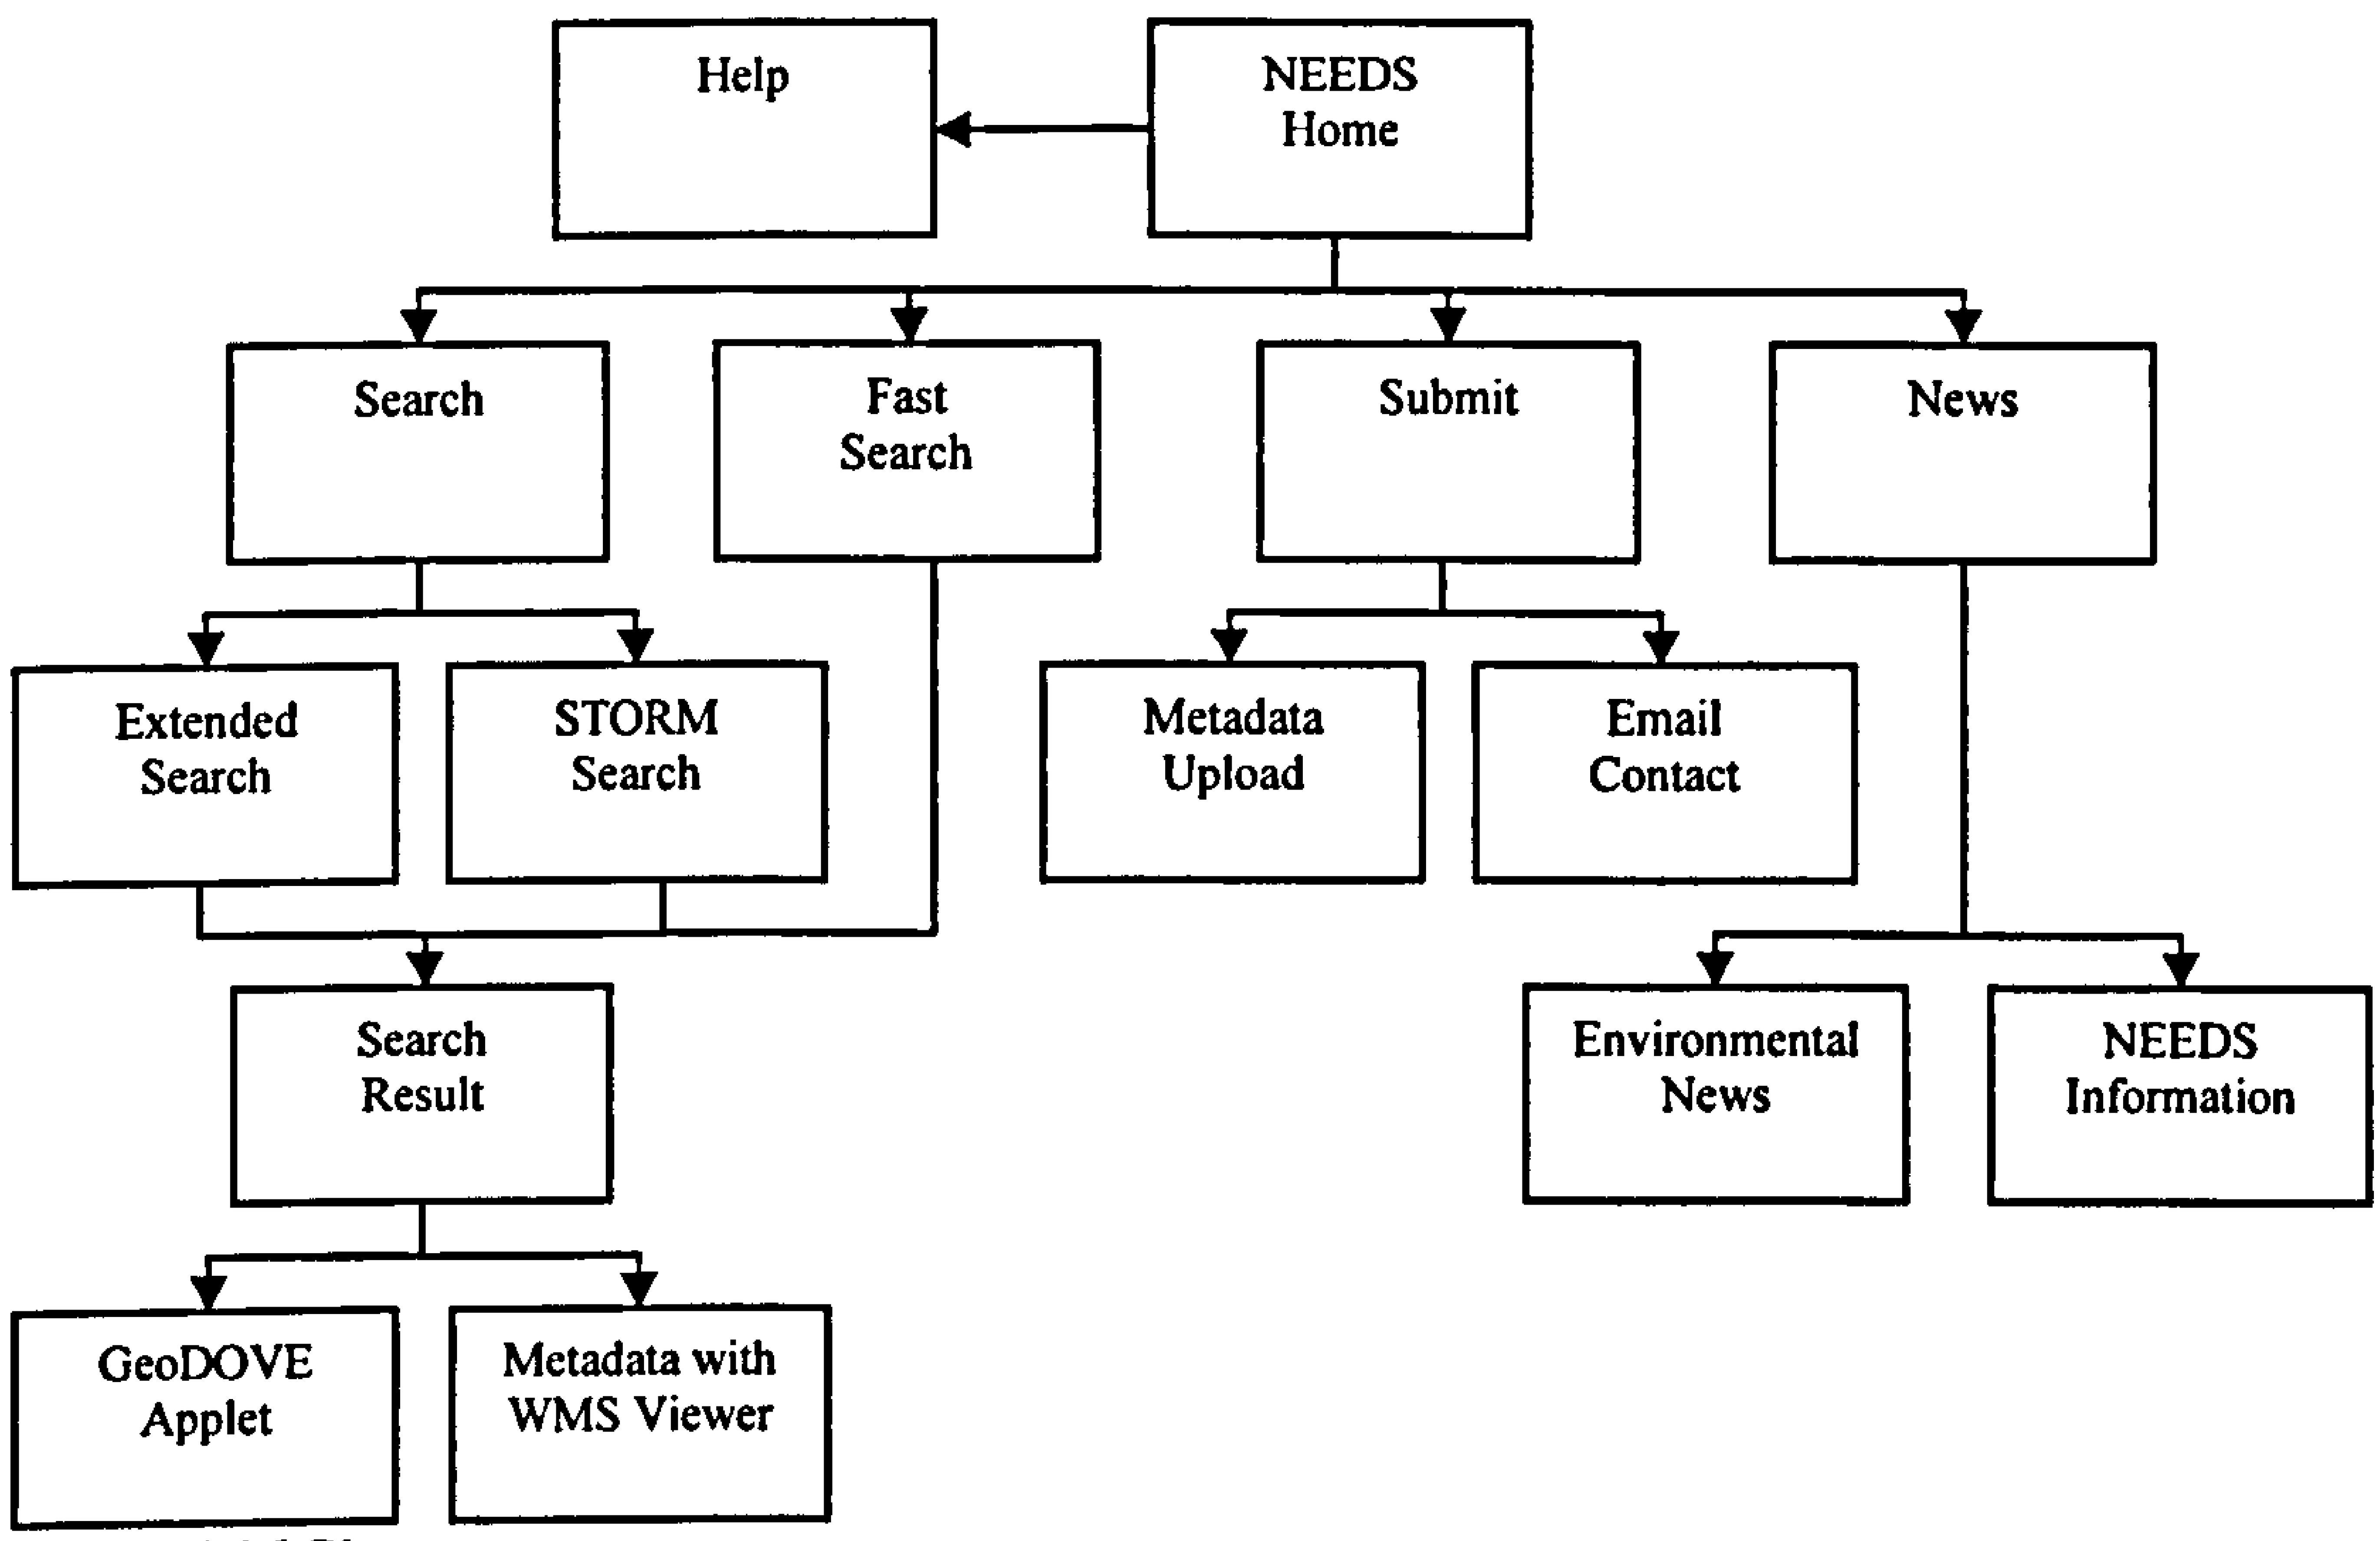

Figure 5.16 Site map

Based on the architecture presented in Figure 5.15 and the site map illustrated in Figure 5.16, the three main facilities were identified for development namely, metadata upload, conventional discovery and data retrieval facilities. The rest of this section discusses their design and implementation.

5.4.1 Metadata Upload

The metadata entries are entered through Java Server Pages (JSP). JSP offer dynamic content that depends on the parameters entered from the client-side. When a JSP page is visited, it is compiled into a Java Servlet which is a Java program that runs on a server and returns some output to a client. Servlets and JSP were selected for this project, over other server-side scripting platforms such as PHP and ASP.NET, because Java servlets have access to several API's such as the JDBC, Swing, Java3D and others. Furthermore, Java is free and some of its APIs, such as Java3D are now open source projects. The geoportal was designed to support two main scenarios for uploading metadata i) where the metadata references an external data source ii) where the metadata references an in-house data source. The first scenario is primarily meant for data vendors who wish to publish the existence of data, but without freely distributing the data itself. By referencing an external data source, the geoportal is able to publish thumbnails of data with access constraints controlled and enforced by the data vendor. The second scenario is primarily for data donors who do not have the resources to disseminate geographic data.

The metadata submission page implemented an HTML form for entering metadata fields. For the first aforementioned upload scenario, a textbox was created for entering a WMS `getMap` request URL. As was stated in Chapter 4, the `getMap` request is the interface for retrieving rendered vector or raster maps. For the second scenario, an additional field was made available for selecting a dataset to upload. The dataset had to be a compressed archive of the files being donated. For example, a ZIP file containing shapefiles (SHP) or GML files, alongside other non-geographic files such as PDFs or Microsoft Word files. The servlet unpacks the ZIP archive into a special folder on the server where all the files are made available for service-based dissemination.

In both scenarios all metadata entries are stored in a PostgreSQL relational table, which we have named 'geometadata'. The `getMap` request entered by the data donor, as per the first scenario, is stored in table geometadata as is. However, as the data donor in the second scenario does not have access to resources for disseminating data, a `getMap` request is automatically generated by the upload servlet. The generated `getMap` request contains parameters such as the bounding coordinates, a default height and width of 500 pixels and the URL of the geographic file uploaded. The geographic file is renamed to a 'safe' filename because certain characters are allowed for filenames but encoded on a URL, for example '(' is encoded as '%28' and ')' as '%2e'(W3C, 1999-2006). The `getMap` request in this case references a web map service built into NEEDS that allows dynamic references to data sources. The web map service is implemented in a class called `GeoRenderer`.

Once metadata has been uploaded, an administrator is given the opportunity to examine newly added records before they are published. This is achieved through the separation of the primary metadata storage database from the metadata publishing server. As stated, the primary storage is offered by a relational table called geometadata and the metadata server is implemented through an Isite installation. When an administrator selects a metadata record for publishing, the metadata is exported as an XML file into the Isite data folder. Once all new metadata records have been examined and exported to the publishing folder, the administrator is then able to invoke a re-indexing operation on Isite. All the administrative tasks, including XML export, starting, stopping and re-indexing the metadata server is controlled from the geoportal and hence the existence of the geometadata table and the metadata server is transparent to the user.

5.4.2 Conventional Discovery

The conventional discovery facility has two main tasks; to search and preview. We name this facility 'conventional' to distinguish it from the ontology-based approach adopted for the STORM browser. The search task involves the sending of a search request to a Z39.50 geographic metadata server and the extraction of metadata fields from the resulting XML into a formatted web page. The 2D preview task is only available for metadata records that include a `getMap` request on the 'online resource linkage' attribute of the ISO19115 or the CSDGM, therefore metadata records on the NEEDS server included a `getMap` URL. The objective of the preview task is therefore to extract the `getMap` request and automatically prepare a WMS client for previewing the `getMap` source.

Technology	Task
JSP	Presentation
Servlet	Request and response processing
PHP	Connection and retrieval of records from Z39.50 servers

Table 5.1 Summary of tasks performed by the three different technologies during search

During the search task a user enters search parameters on a JSP page. The search parameters are sent to a servlet, which then prepares and forwards a request to a PHP web application. The PHP web application adopts the Yaz toolkit for Z39.50 communication and hence retrieves the titles of metadata records returned by the Z39.50 servers. The servlet and JSP then convert the list of titles from XML to HTML for presentation to the user. When a user selects a particular record for preview, the request is sent to a second servlet which sends a request to another PHP web application. The web application then sends the complete title of the selected record to the Z39.50 server which returns a complete metadata record in XML or SGML, depending on the server's implementation. The XML document is forwarded to the second servlet which extracts the metadata fields from the XML document and stores them in an instance of the Java `Hashtable` class. A JSP page then retrieves the field-value pairs from the `Hashtable` and presents them to the user. A summary of the description of tasks performed by JSP, servlets and PHP is presented in Table 5.1.

The WMS client is built into the aforementioned JSP page, therefore the preview chronologically follows the search task. The main technologies adopted for the development of the preview facility include JavaScript, JSP and servlets. It should be noted that JavaScript programs run on a client web browser, whereas JSP and servlets run on a server. When the JSP page retrieves the `getMap` request from the aforementioned `Hashtable`, it separates the `getMap` request into the different URL parameters for example, the geographic layer, WMS server and bounding coordinates.

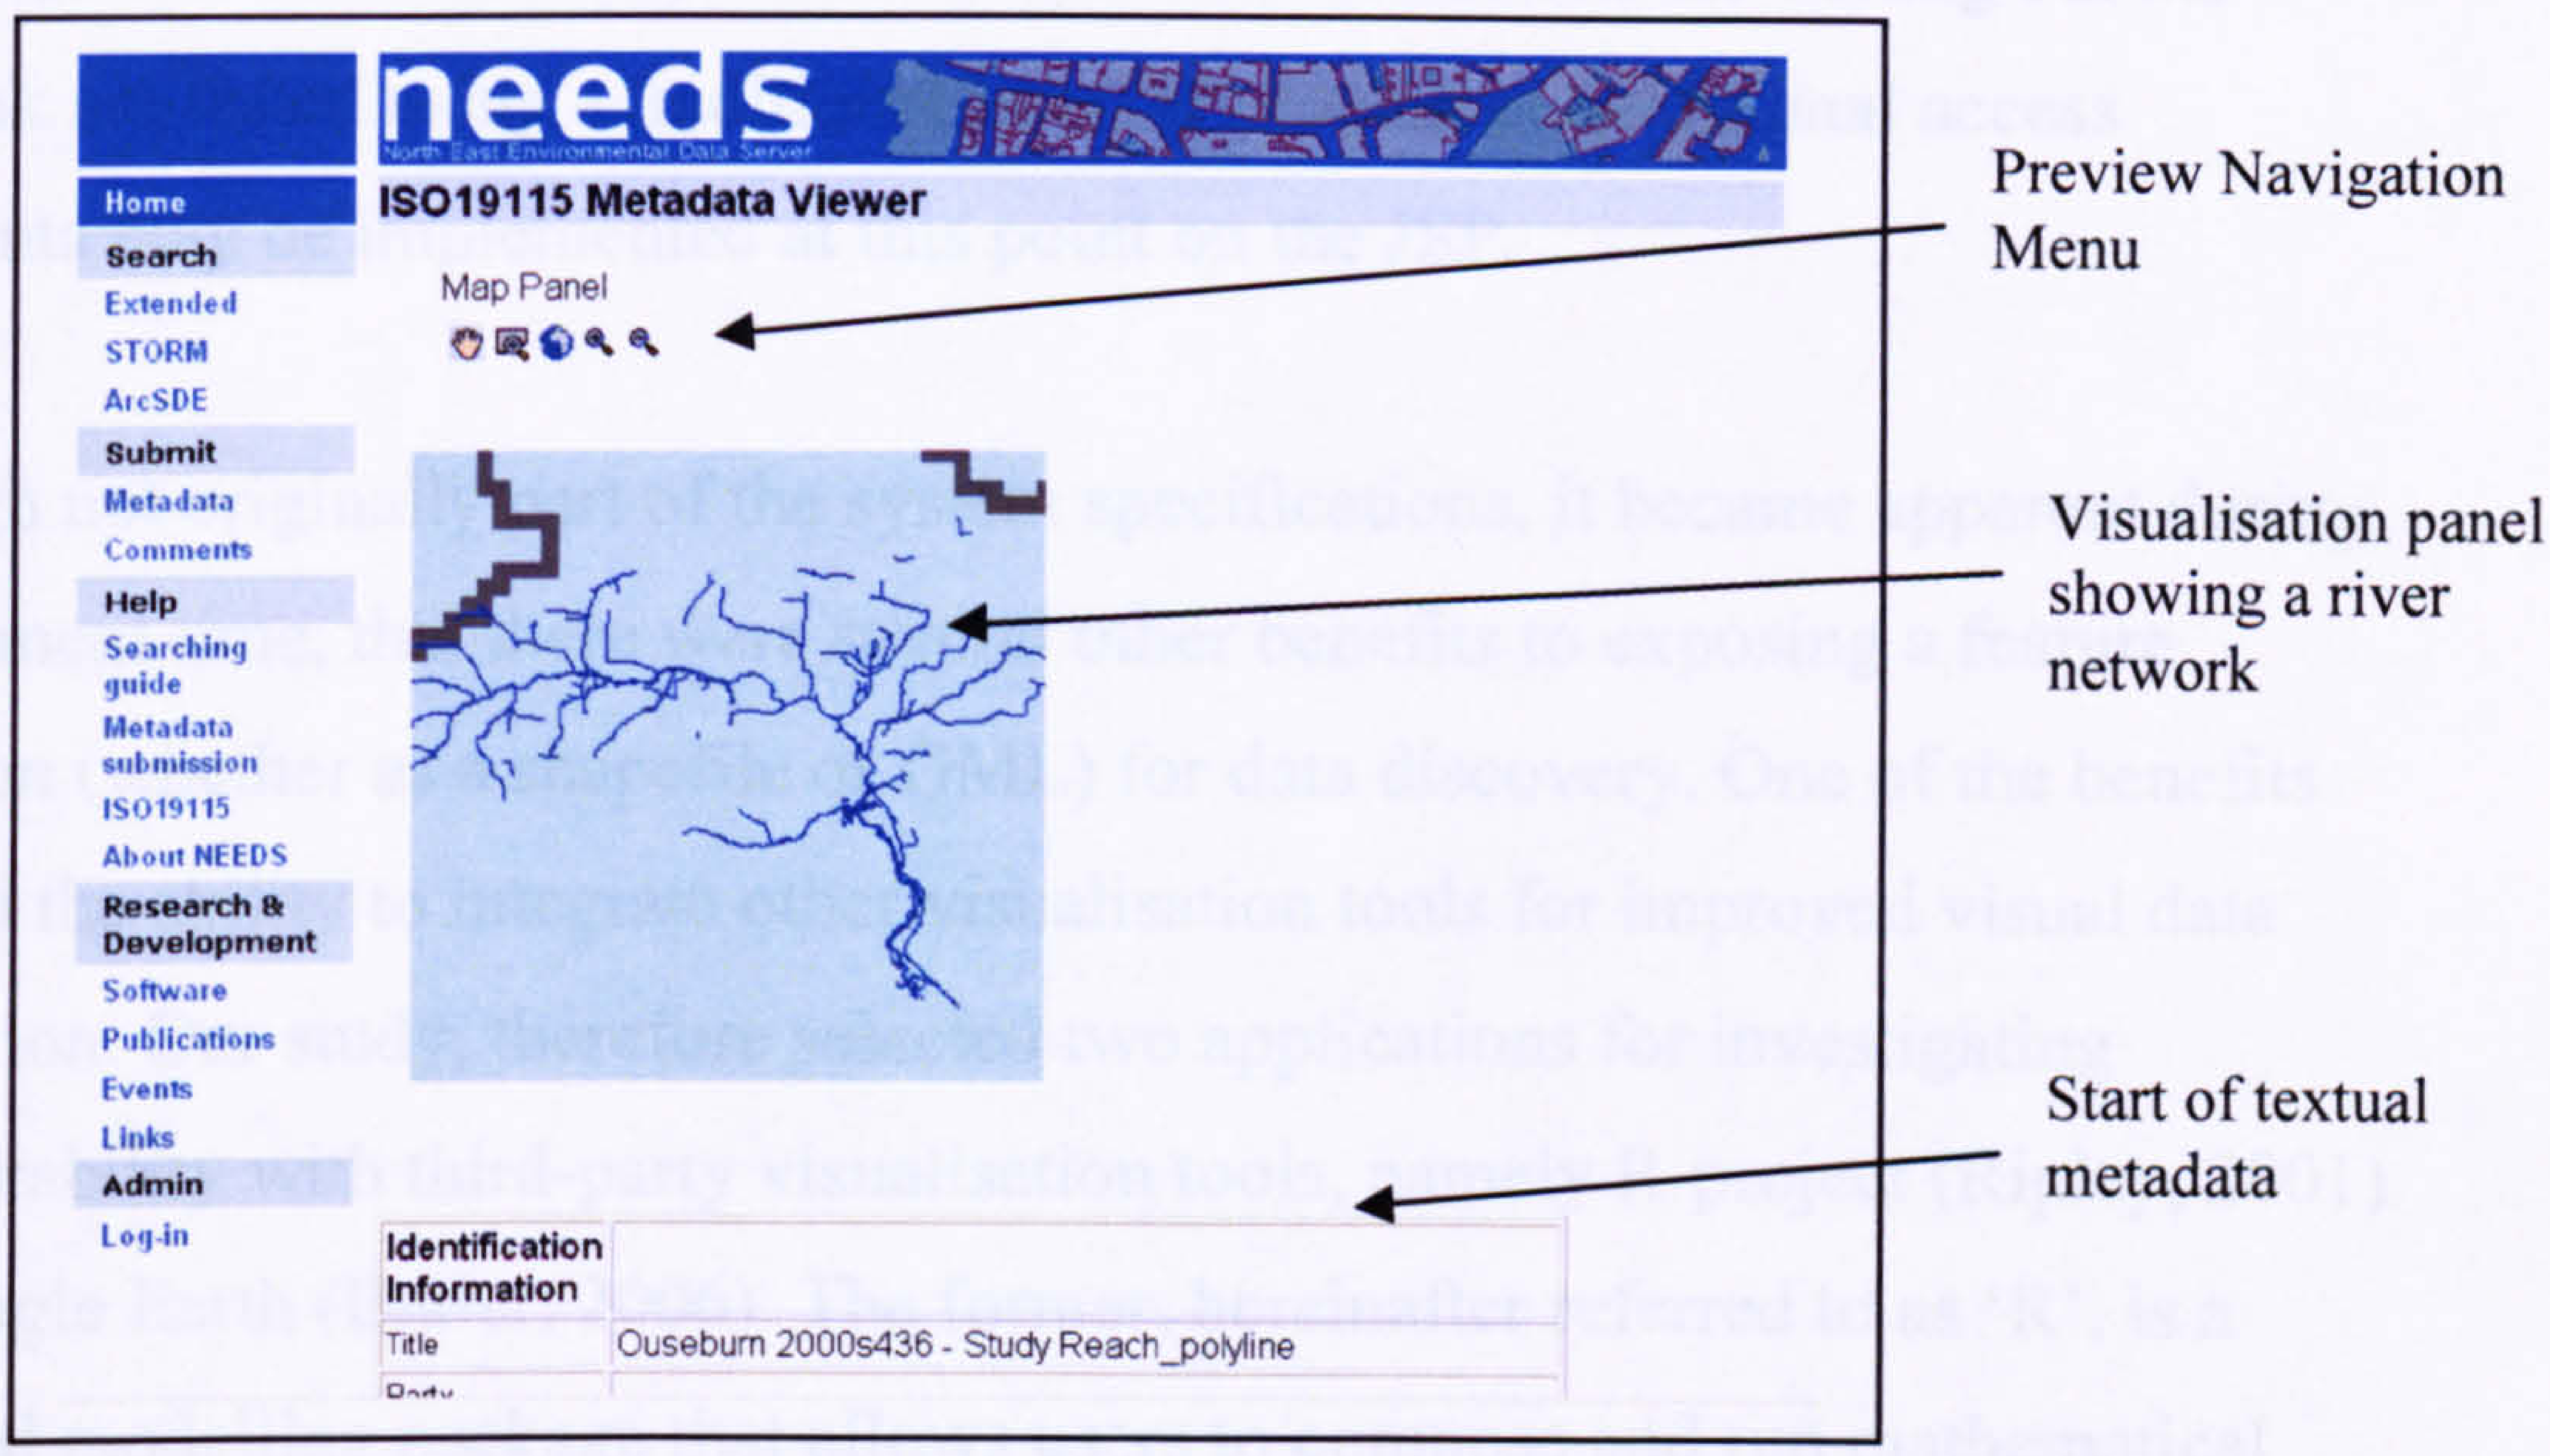

Figure 5.17 A screenshot of the NEEDS WMS client

The parameters are then sent to a third servlet which automatically prepares a JavaScript configuration file for the WMS client. Some of the parameters used for configuring the WMS client include the boundary coordinates, image format, image size and the names of layers to view. Other `getMap` parameters such as the style are ignored in the current implementation. The boundary coordinates are used to initialise the current viewable extents. Each click on the navigation menu alters the coordinates of the current viewable extent. A `paintComponent` function was implemented in JavaScript and is invoked every time the viewable extent is altered. The name of the method is meant to highlight its relation to the Java GUI components which implement a similarly named function. A screenshot of the WMS client is shown in Figure 5.17.

5.4.3 Data Retrieval

For metadata with accompanying datasets, the JSP script writes out a hyperlink that references the ZIP archive that was uploaded. If the metadata does not include an accompanying dataset, for example, if it references a third-party metadata server, then the hyperlink is not written out to the JSP page. The hyperlink is also not written out if the client's request has been sent from an unauthorised IP (internet protocol) address. The IP address is a unique sequence of numbers used to identify nodes (computers and other devices) on the Internet. As conditions for writing out the hyperlink are tested using a traditional Java 'if' statement, additional access constraints may be implemented at this point on the JSP.

Although not originally part of the system specifications, it became apparent during development time, that there were several other benefits to exposing a feature collection (whether as a shapefile or GML) for data discovery. One of the benefits included the ability to integrate other visualisation tools for improved visual data exploration. Our study, therefore selected two applications for investigating interoperability with third-party visualisation tools, namely R-project (Ripley, 2001) and Google Earth (Butler, 2006). The former, hereinafter referred to as 'R', is a statistical modelling package that allows users to compose and run mathematical functions from a GUI or the command line. One of the differences between the two systems is that the primary file format used by Google Earth, Keyhole Markup Language (KML), is a geospatial file format which borrows heavily from GML(Google, 2005, OGC, 2004). In contrast, R does not natively handle geospatial file formats and hence data held by NEEDS requires processing before use in R.

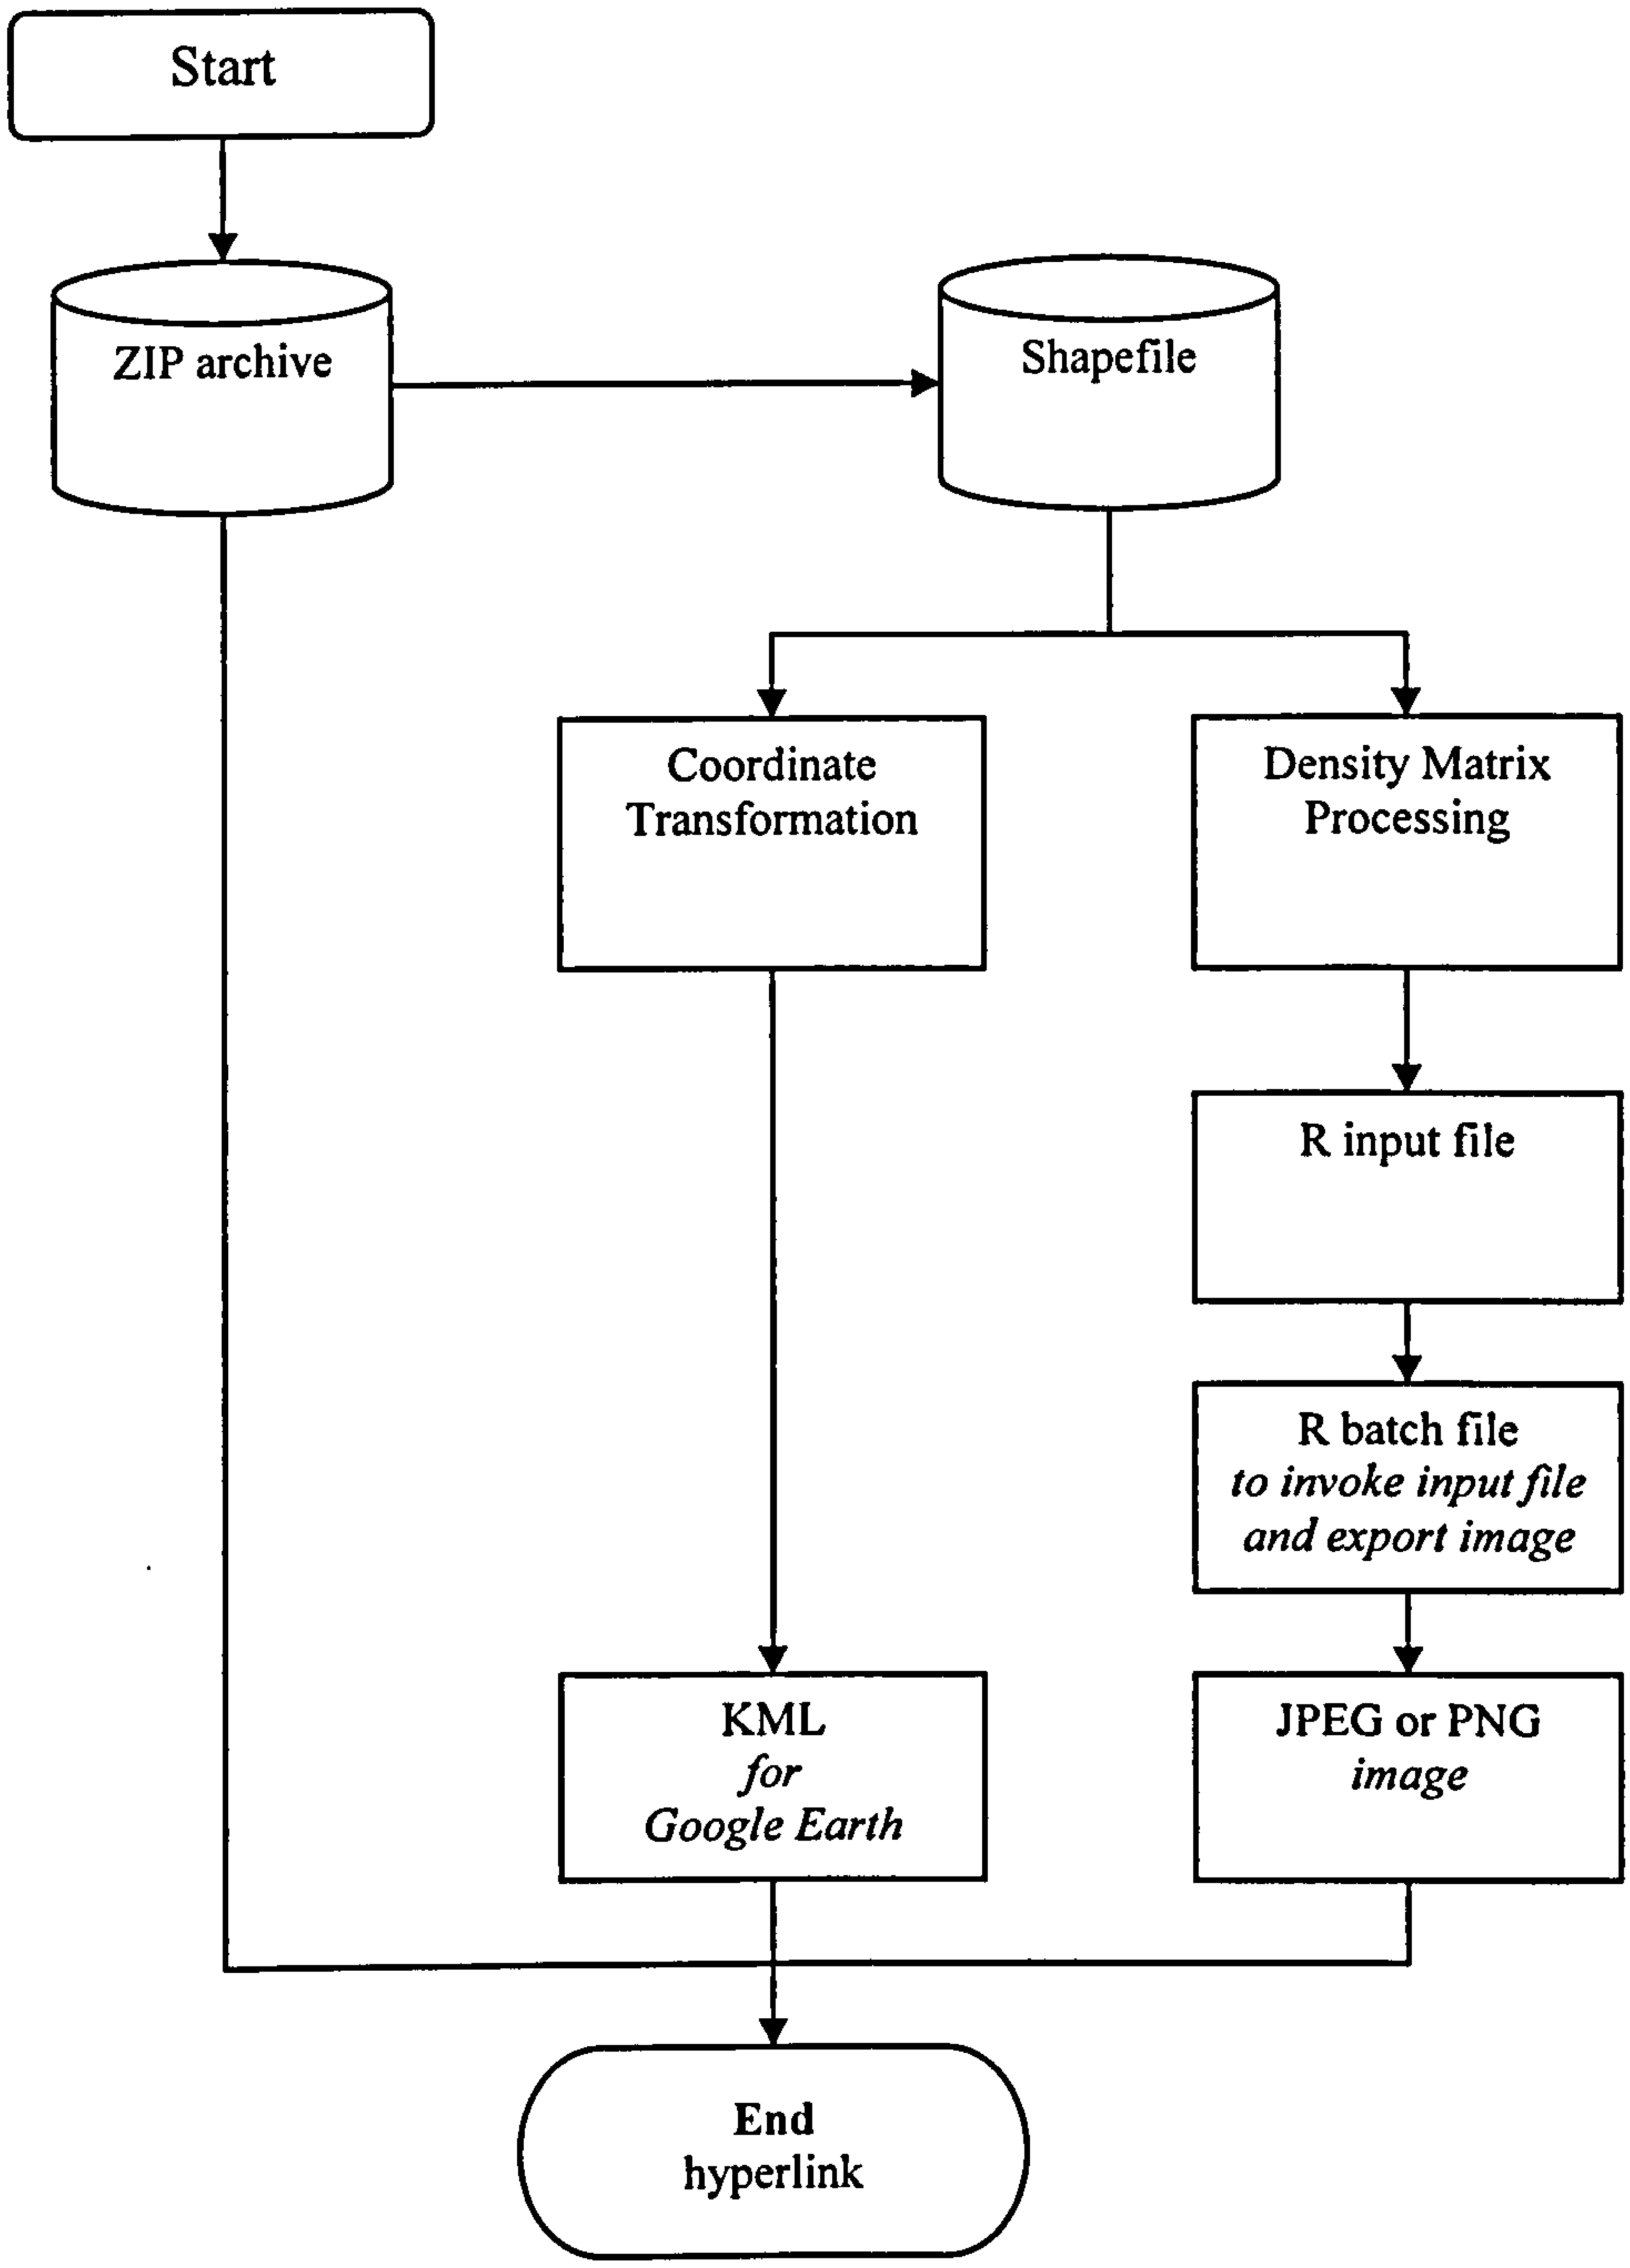

Figure 5.18 A flowchart of the data retrieval process

The Java platform offers a class called `Runtime` that allows an application to interface with the environment from which it runs. Commands that would normally be executed from the Windows DOS prompt or Unix shell prompt are passed to an instance of the Runtime class through the `exec` method. When the method is called, a process is created which the Java application then uses to determine if the method operation was successfully. To integrate R with NEEDS, we created a hyperlink to a JSP page that invokes a data processing servlet then writes an R input file. For illustrative purposes, we implemented a data processing function that determines the density of features in a ten by ten grid matrix. Using the `Runtime.exec` method, the JSP invokes an R batch script which references the R input file that has just been created. The R batch script creates a set of statistical graphs from the ten by ten matrix and writes the graphs to image files that are then referenced as HMTL IMG entries in the JSP output. This workflow is illustrated in Figure 5.18.

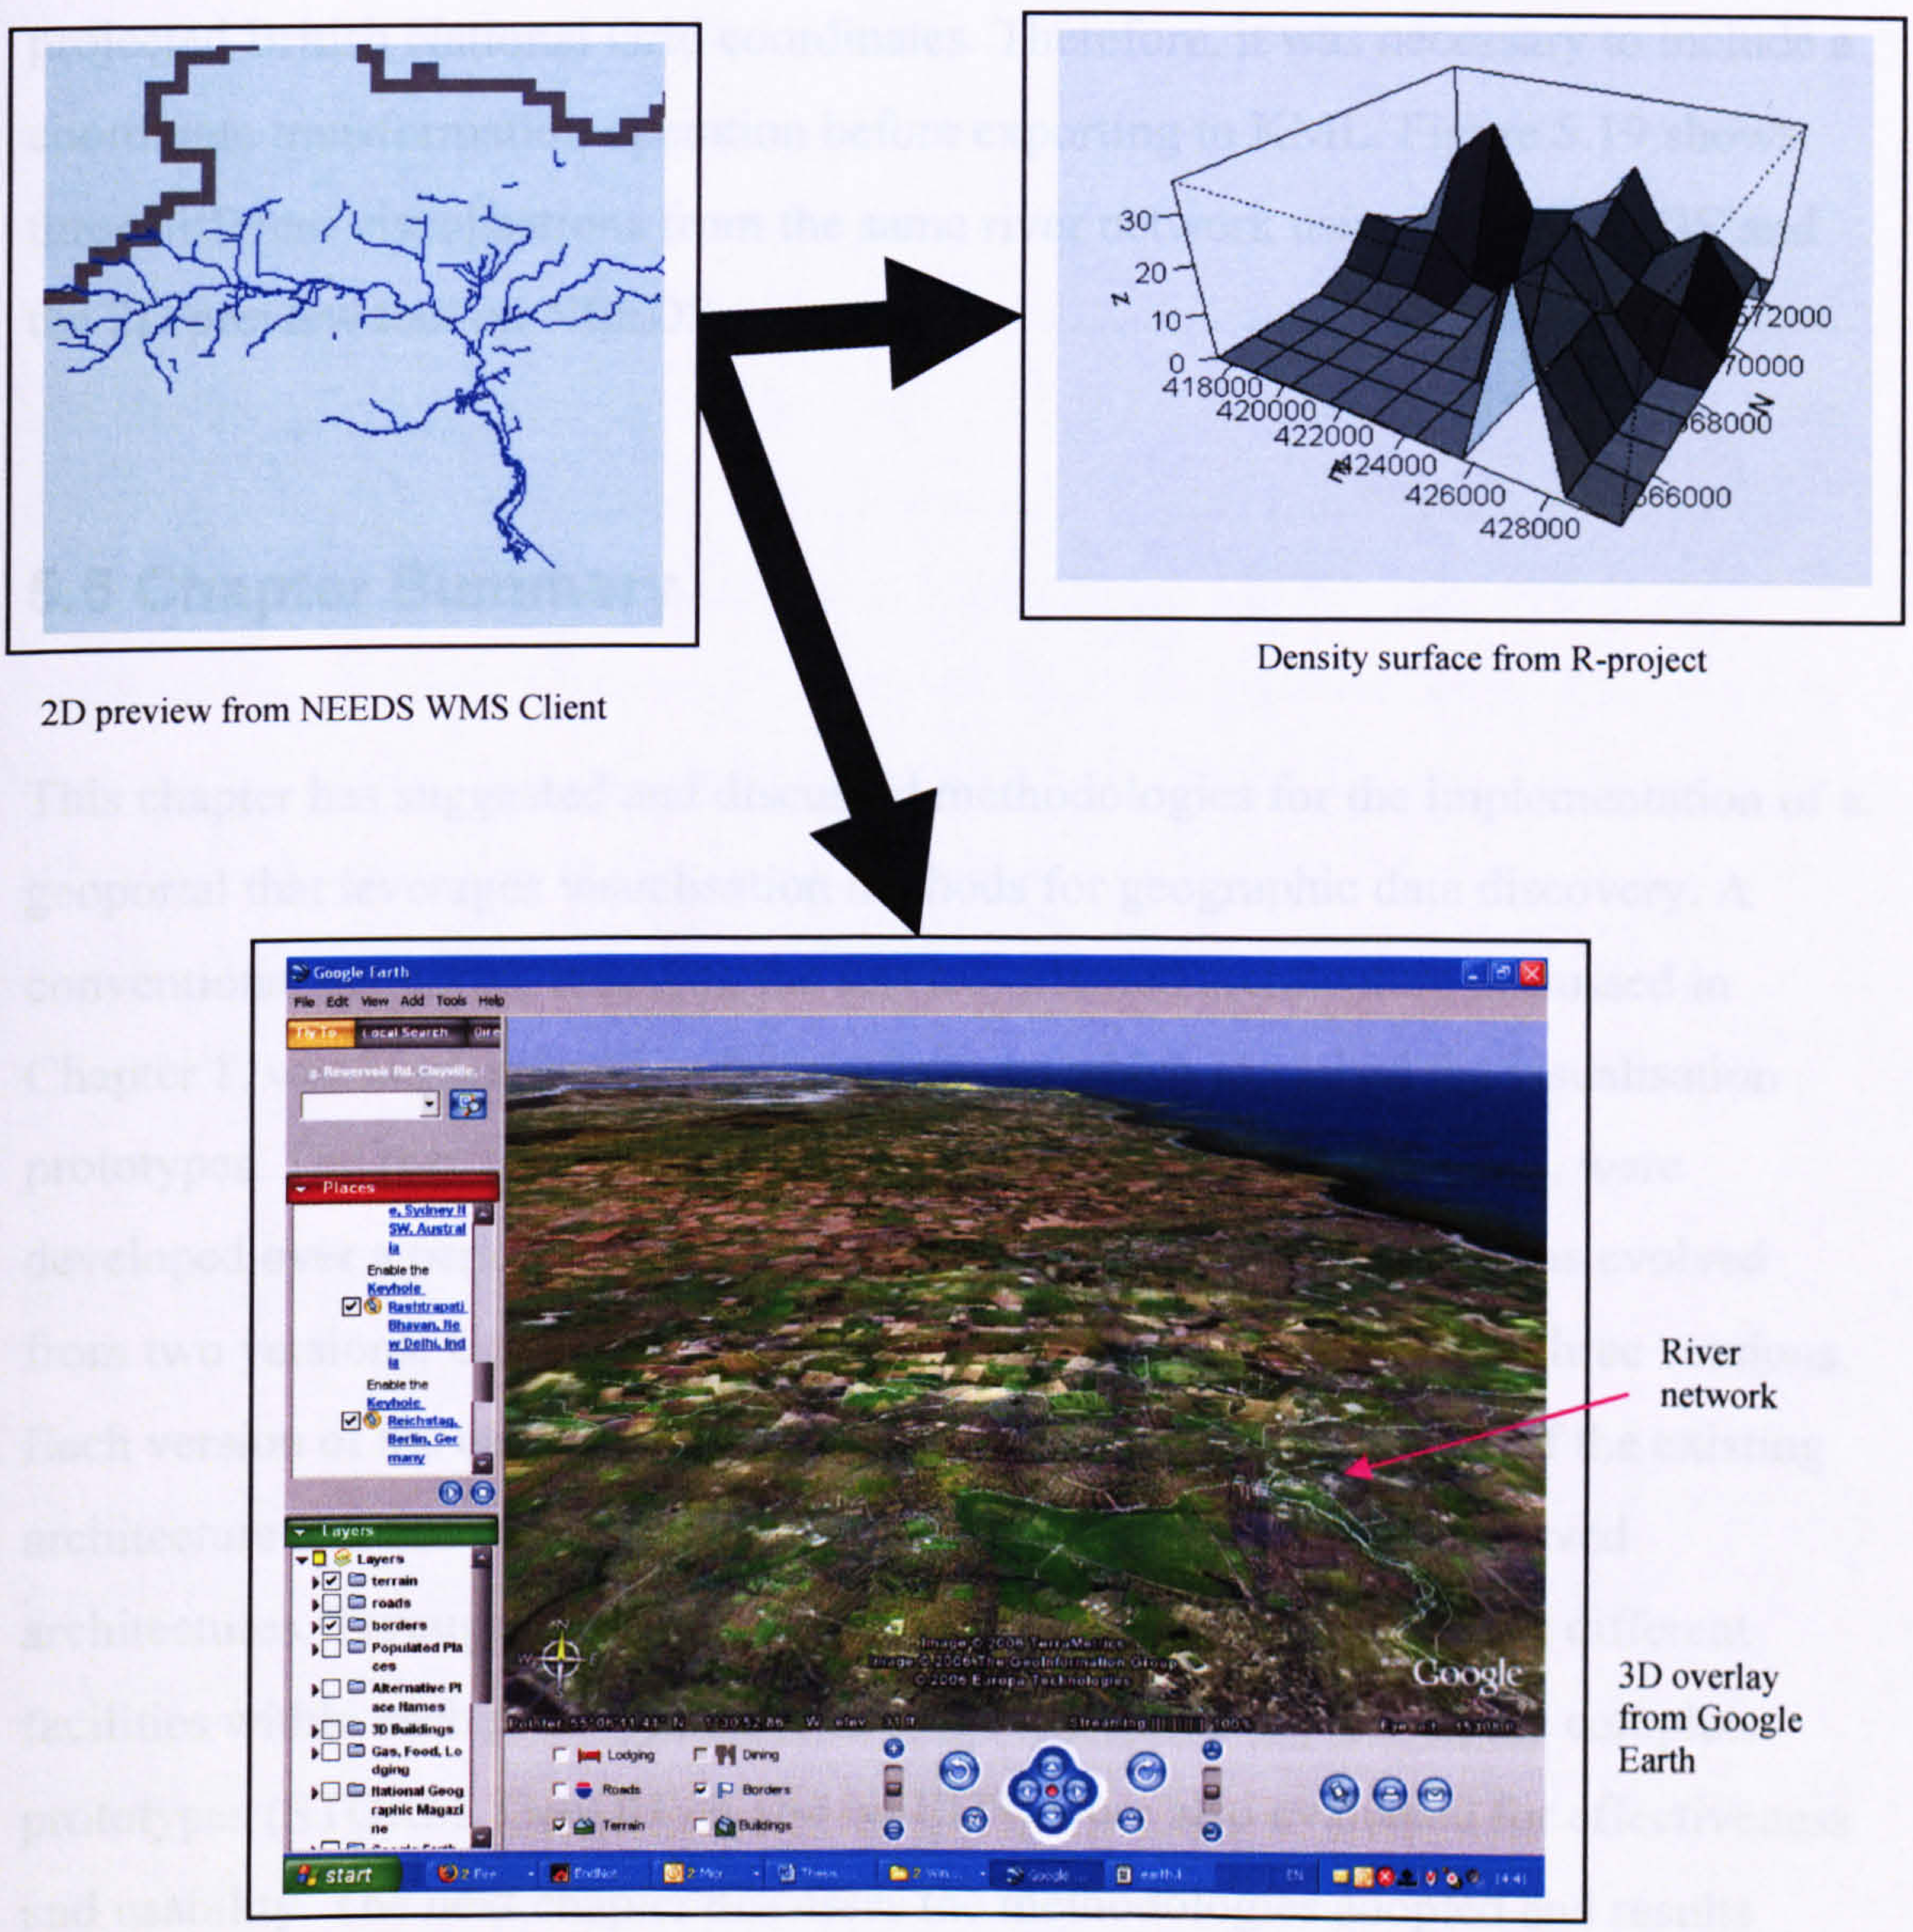

Figure 5.19 Screenshots of output from R-software and Google Earth interoperating with NEEDS

The integration of Google Earth (GE) with the rest of the NEEDS system follows a different data workflow, as illustrated in Figure 5.18. GE operates as a standalone application that is launched when a user clicks on a hyperlink to a KML file. If GE is already running, then clicking on a KML file triggers GE to upload the file into the current 3D scene. Therefore, to integrate GE with NEEDS, it was necessary to export a feature collection to a KML file. However, the main challenge with this process was that KML files only adopt the WGS84 coordinate system (Google, 2005: URL http://earth.google.com/kml/kml_tags.html). As the NEEDS system primarily holds data from the United Kingdom, most of the datasets held on NEEDS are based on projected British National Grid coordinates. Therefore, it was necessary to include a coordinate transformation operation before exporting to KML. Figure 5.19 shows three different visualisations from the same river network using R-project, GE and the 2D preview tool on NEEDS.

5.5 Chapter Summary

This chapter has suggested and discussed methodologies for the implementation of a geoportal that leverages visualisation methods for geographic data discovery. A conventional geoportal, based on the service-oriented architecture discussed in Chapter 1, was implemented as the platform on which to embed the visualisation prototypes. The two visualisation prototypes, STORM and GeoDOVE, were developed over a period of three years. During that period, STORM has evolved from two versions, GeoDOVE from three versions and NEEDS from three versions. Each version of the each prototype required a thorough examination of the existing architecture and this chapter has presented the current and more improved architectures. Throughout the development of the complete system the different facilities within each prototype were individually tested. However, the complete prototypes (STORM, GeoDOVE and NEEDS) were also evaluated for effectiveness and usability. The next chapter discusses the methodologies adopted and results observed during the evaluation studies.

Chapter 6 Evaluation

6.1 Introduction

One of the stated research questions is "what effect would the suggested approaches have on the performance of a geoportal?". In addressing this question, GeoDOVE and STORM are evaluated independently of the NEEDS geoportal. Thereafter, the geoportal (NEEDS) is evaluated with the visualisation applications integrated into it. "The determinants of system-acceptance are functionality (the degree to which the system provides functions which the users need to do their tasks) and usability (ease of use, a user-related and a task related concept)" (Benbunan-Fich, 2001: pp.152) . Therefore, the aim of the evaluation was to determine whether a geoportal that supports data discovery through multidimensional visualisation, as proposed in this thesis, could provide acceptable usability and performance. Related studies that have adopted both a user and system-oriented evaluations include for example, Lighthouse (Leuski and Allan, 2004), SPIRIT (Bucher et al., 2005) and the Integrated Thesaurus-Results Browser (ITRB) (Sutcliffe et al., 2000).

6.1.1 Empirical Testing in GIR

An overwhelming majority of IR evaluation studies focus on system-based empirical testing methods. Empirical testing involves the carrying out of a controlled experiment to gather evidence to support a hypothesis. Experiments in empirical evaluations therefore follow the same guidelines; first a hypothesis is presented then a set of controlled input variables that differ only in value are tested. The resultant observations are recorded and later analysed to determine any relations between changes in the values of input variables and the recorded observations. Within both IR and GIR, the two measures that are commonly evaluated are precision and recall; the latter being the proportion of retrieved relevant documents among all relevant documents in the document collection and the former being the proportion of relevant documents among retrieved documents (Kobayashi and Takeda, 2000). The following formulae are commonly used for calculating normalised precision P_i and recall U_i for a retrieved set i:

(a) $$P_i = \frac{|R_i|}{|R_i| + |N_i|}$$ (b) $$Ui = \frac{R_i}{W_i}$$

where:

- R_i is the number of relevant documents retrieved
- N_i is the number of non-relevant documents retrieved.
- W_i is the number of all documents within a corpus that are known to be relevant to a query

Although these have become standard measures for traditional IR evaluation, different studies have suggested that they offer an incomplete evaluation (Wang and Forgionne, 2006, Borlund, 2003b). The studies suggest that precision and recall present the following limitations: firstly, the measures are computed from topical relevance, making them static. This means that they assume the user will find a suitable document from the first query. Therefore, they do not recognise the interactive and iterative process of IR. Secondly, they do not include the user's view of relevance within the evaluation. This means that the evaluation is constrained by how well the system matches the query constraints to the document description (normally metadata). Lastly, they ignore the multidimensionality of relevance, and in particular geographic relevance which may be influenced not only by the topic of the document but by the scale and date of creation as well.

Evaluation of precision and recall in traditional IR has involved the use of very large document collections. Some of the earliest adopted document collections included the Text REtrieval Conference (TREC) and the Cross Language Evaluation Forum (CLEF) collections. Unfortunately these ignore the spatial (for example, scale and extent) and graphical (for example, user interface and graphic variables) characteristics of a GIR System. In response, a new initiative, called GeoCLEF, has been established to augment the multi-lingual retrieval framework CLEF with geographic considerations (Gey et al., 2005). Further justification for a geo-centric track within CLEF is offered by Martins et al. (2005), who emphasize that a GIR system integrates different components that interoperate but may require separate evaluations. They list the following specific considerations i) geographic ontology development, ii) geographical reference extraction from text, iii) assignment of geo-scopes(coverages), iv) ranking according to geographic relevance and iv) design and development of user interfaces for GIR. However, as GeoCLEF is still at an early stage of development, it will not be employed for this study. Instead, this study uses geographic metadata collections registered on the FGDC Clearinghouse registry (FDGC, 2005) as these are the primary catalogues within the geographic information industry. However, future studies should consider using GeoCLEF for similar evaluations.

6.1.2 Usability Testing

As mentioned earlier in this chapter, several studies have emphasized the importance of involving the user in any evaluation of an IR or GIR system. The three main approaches to generic usability testing are usability inspection, group walkthroughs and user testing (Brink et al., 2002). Usability inspections involve observations of specific aspects of a system by an expert. Although the expert can be the actual developer, it is highly recommended that the evaluator be someone else. Group walkthroughs involve a team of stakeholders in the system coming together to test and evaluate the system. The team is often made up of managers, clients, sales and support personnel and therefore excludes the developer or designer of the system. Both usability inspections and group walkthroughs are sometimes referred to as heuristic evaluations — a method for structuring the critique of a system according to a collection of recognized usability techniques. User testing involves observing users as they carry out prescribed tasks on the system. It is one of the most popular evaluation methods because of its low cost and involvement of members of the expected user community. Two main methods for collecting data during user testing are often applied namely, questionnaire filling and the 'think aloud' method.

Questionnaires allow the developer to assign fixed questions that refer directly to aspects of the system that are being investigated. The responses to the questions can be entered during or after the testing session. Although traditionally the responses were written on a sheet of paper, modern studies employ forms on websites that store the responses directly in an RDBMS. The ability to disseminate usability questionnaires through the web means it is possible to involve users from different countries in a single usability test. The questions posted allow for profiles of the user to be acquired by asking general questions about the user's past experience and education with respect to concepts being investigated. Dix et al., (2004) categorise the types of questions generally posted in a questionnaire into the following five classes:

- General: These are questions that obtain a profile of the user such as past experience, education and country of residence (if international users are involved).
- Open-ended: These ask the user their own opinions without restricting the possible responses. Popular open-ended questions include for example "what improvements to this system do you recommend?"
- Scalar: These ask the user to specify the degree to which a statement relating to the system is correct. The degree is usually specified on a 7-step scale of integers. Even numbered scales, for example 4-steps, are usually avoided as they do not cater for undecided judgements (i.e. in the middle of the range). Users tend to find it difficult to judge when responding to scales greater than 7-steps.
- Multi-choice: These questions offer the user a set of explicit and independent responses.
- Ranked: These are primarily comparative questions that ask the user list a set of items according to user preference.

Another popular method for collecting data for usability testing is the 'Think Aloud protocol' approach, which is also known as Protocol Analysis (Owen et al., 2006, Lewis, 1982, Nielsen, 1994). This method requires users to verbalise their thoughts as they interact with the system. The verbalisations, sometimes with screenshots of the application, are often recorded on audio or video storage media. By recording the verbalisations assumptions, inferences and problems encountered during user-system interaction can be identified during later analysis. This provides a rich set of data for the developer to analyse. Hence, it is often considered that from five to ten users are sufficient for an effective protocol analysis (Nielsen, 2000, Brink et al., 2002). Although highly valuable in collecting data relating to user experience with an application, a think-aloud evaluation can be expensive to run due to the need for audio-visual recording equipment. The recording equipment can also be time-consuming to set up. Further, Nielsen (2000) warns that the video camera can be intimidating to users. However, other recording devices such as dictaphones may be less intimidating.

Each of the recorded sessions is transcribed into text. The quality of the transcription much depends on the clarity of the recording. From the transcriptions, irrelevant verbalisations are omitted and relevant ones are segmented (categorised) according to the objectives of the evaluation exercise (Blok, 2005, van Someren et al., 1994). For example, general comments about weekend sports or the morning's news are seldom important for a usability study and hence may be omitted. However, comments about the appearance of buttons and the navigation of the application under examination may be segmented into 'layout' and 'interaction' categories respectively. From those comments that are regarded as relevant, their frequency is often a useful indicator of the importance of the subject of the comments (Benbunan-Fich, 2001). Additional data is obtained from the time taken to accomplish each task when using the system. This data may be used to identify lengthy procedures that may require revision as "people have trouble with long procedures unless there are memory aids provided on paper or by the system" (Brink et al., 2002: pp. 112).

This section has explored the different evaluation methodologies employed in IR, GIR and general website evaluations. Whereas traditional IR employed empirical testing, there is a growing recognition that user-system interaction is a significant factor in system acceptance. The section has also presented arguments from related GIR studies that emphasize the need for geo-centric methods of evaluating GIR systems as opposed to traditional IR techniques. The rest of this chapter therefore discusses the evaluation of STORM, GeoDOVE and NEEDS.

6.2 Evaluating STORM

The aim of the evaluation of STORM was to determine the retrieval performance of the STORM browser, in comparison to a traditional geoportal (i.e. GIgateway). Users were asked to carry out a set of specific queries on each tool. For each query, they were to record the number of datasets retrieved and the number they found relevant to the query. The following subsections discuss this experiment further.

6.2.1 Experimental Setup

It is trivial that when comparing two systems A and B, if A returns more relevant documents and fewer non-relevant documents than B, then A is more effective than B. In traditional IR evaluation, the number of relevant documents for particular queries is known prior to evaluation. The formula for computing recall requires that the number of all datasets in the corpus known to be relevant be determined. However, due to limitations of the Z39.50 protocol, adopted for our study, this could not be determined. This means the study could only determine the precision of each search. "However, in many interactive settings, users require only a few relevant documents and do not care about high recall to evaluate highly interactive information access systems" (Hearst, 1999: pp. 262). Similarly, Leuski and Allan(2004) base their evaluation of the Lighthouse system on precision without recall.

Keyword	Spatial footprint	Collection
Borehole	north=54.955 south=53.024 west=-3.244 east=-0.87	BGS
Census	north=56.624 south=54.67 west=-3.178 east=-1.51	Central Government (IGGI)
Seismic	north=55.592 south=52.628 west=-3.947 east=-0.895	BGS
Macrofossils microfossils	north=55.329 south=52 .65 west=-3.508 east=0.597	BGS

Table 6.1 Some of the queries used in the evaluation

The United Kingdom's national spatial data clearinghouse, GIgateway, was selected as a benchmark for this study because *i*) it is the national spatial data clearinghouse; *ii*) it adopts the widely-used YAZ Z39.50 client (West and Schofield, 2004); *iii*) it harvests from clearinghouse nodes that implement the widely-used CNIDR zserver

(FDGC, 2005). Both GIgateway and the STORM browser were connected to various document collections, including the main GIgateway catalogue, British Atmospheric Data Centre (BADC) catalogue, British Geological Survey (BGS) catalogue, Central Government metadata archive (IGGI) and others. It was envisaged that the use of these archives would add to the realism and applicability of this study, as they are used within the geo-scientific industry.

A group of ten users with highly varying computing and geospatial backgrounds agreed to take part in the study. The handout given to the users is presented in Appendix C. The users had the following profiles:

- User A holds a first degree in Computing Science and is studying for a higher degree in Geomatics. He is less than 25 years of age.
- User B holds a basic degree in Actuarial Science. She is less than 25 years of age.
- User C holds a first degree in Surveying. She is aged between 25 and 30 years of age.
- User D holds a first degree in Civil Engineering and is studying for a higher degree in Geomatics. He is over 30 years of age.
- User E holds a first degree in Agriculture and is studying for a higher degree in Geomatics. She is over 30 years of age.
- User F holds a first degree in Surveying. She is aged between 25 and 30 years of age.
- User G holds a first degree in Archaeology and a higher degree in History. She is less than 25 years of age.
- User H holds a first degree in Geographic Information Systems. He is aged between 25 and 30 years of age.
- User I holds a first degree in Geographic Information Science. He is less than 25 years of age.
- User J holds a first degree in Geography. He is studying for a higher degree in Geomatics. He is less than 25 years of age.

Each user was tasked with carrying out a set of queries— which were consistent across all users. For each query, the users independently noted down the numbers of relevant datasets in the set of top fifty ranked datasets. Only the top fifty ranked datasets were considered because any more datasets significantly reduce the clarity of the visualisation. Considering only the top few datasets is also consistent with studies by Leuski and Allen (2004) and Bucher et al (2005). It was expected that similar queries carried out on both STORM and GIgateway would return varying sets of datasets because the STORM browser expands a user's query with additional search terms. Further, variation of search results was expected between users because of varying expertise. Table 6.1 shows some examples of the queries that were used during the study. As illustrated, these are conventional queries that can be entered into most geoportals. In all queries, the temporal ranges were from 1995 to 2005.

6.2.2 Results

Table 6.2 shows observations from two of the users. The difference in precision (i.e. the percentage of returned relevant datasets to all returned datasets) is presented in the *Diff* column. Positive values indicate more relevant datasets from STORM and negative values indicate more relevant datasets from GIgateway. One of the ten users completed only five of the seven controlled queries hence there were a total of 68 queries completed by all ten users. It was observed that 29 of all 68 queries exhibited higher precision by STORM; only 5 of all 68 queries resulted in higher precision by GIgateway; the remaining queries achieved equal precision measures. A summary of all observations for all ten users is presented in Table 6.3. It should be noted that the integrity of the results is upheld by the following points:

- Eight of the ten users show an improvement in favour of STORM
- The same metadata collections are searched by both systems
- The same queries, in number and form, are entered into both systems
- The users have varying degrees of both geospatial and computer-based skills
- The same users take part in the questionnaire-based evaluation

(i) (ii)

Query	STORM	Gig	Diff
1	100	100	0
2	60	4	56
3	20	20	0
4	100	100	0
5	54	44	10
6	54	34	20
7	48	8	40
Average			18.000

Query	STORM	Gig	Diff
1	100	100	0
2	0	0	0
3	22	20	2
4	100	100	0
5	46	40	6
6	0	0	0
7	34	4	30
Average			5.429

Table 6.2 Sample results from two of the users showing the difference in precision

The results of the comparison between STORM and GIgateway are presented in Table 6.3. Positive values for the differences in precision represent an improvement and negative values suggest a reduction in precision when using STORM. The results were assessed using a t-test to statistically compare the means of values of precision. We calculated 18 degrees of freedom from our two samples of 10 observations each. The results show a calculated t-value of 2.34, with a probability of error of 0.03. Looking up the t-value in a t-table, we observe that our calculated t-value is higher than the tabulated t-value for 18 degrees of freedom, with a critical value of 0.02. This suggests that there is a 98% probability that the difference between the means is not through chance and thus the t-test is statistically valid. Variances of 104.67 and 75.15 were calculated for STORM and GIgateway respectively, resulting in a variance ratio of 1.34. Conducting an F-test using the variance ratio, we observe that the difference in variances is not significant. This observation suggests that there was consistency in how users identified relevant datasets from a returned collection on both STORM and GIgateway.

User	STORM	GIgateway	Difference
A	37.71428571	30	7.714285714
B	42.57142857	38.2857143	4.285714286
C	37.42857143	37.7142857	-0.285714286
D	40.57142857	40.8571429	-0.285714286
E	62.28571429	44.2857143	18
F	43.14285714	37.7142857	5.428571429
G	63.42857143	51.1428571	12.28571429
H	45.42857143	38.8571429	6.571428571
I	54.4	47.6	6.8
J	59.42857143	20.5714286	38.85714286
		Average	**9.94**
		Standard Deviation	**11.52**

Table 6.3 Differences in the percentage of relevant datasets to all returned datasets (i.e. precision)

The average difference in precision was approximately 9.9%, with a standard deviation of 11.52, in support of our hypothesis. The only observation that is more or less than twice the standard deviation is that of User J (who is studying for a higher degree in Geomatics). Considering results for User J as an outlier, we calculate an average precision of 6.72%, which still supports our hypothesis. For similar queries, the majority of users achieved higher precision values on STORM than on GIgateway. The negative precision values recorded (for User C and D) were both -0.3%; this was significantly lower than the positive precision values which ranged from 4.3% to 38.9%. It was also observed that queries with domain-specific search terms such as 'macrofossils' or 'seismic' achieved higher improvements in precision than the more general search terms such as 'census' or 'borehole'. This suggests that use of a more domain-specific ontology could result in higher improvements in precision.

6.2.3 Reflective Summary

This section has presented an evaluation of the STORM browser based on the IR performance metric of precision. The results showed an improvement in precision of 9.9% in support of STORM over a traditional GIR system. Removing the outlier (User J), the improvement in precision reduces to 6.7%. Although this value also supports our hypothesis, the decline in precision illustrates the sensitivity of the results to the response of a single user to the STORM system. However, the variances for both STORM and GIgateway do not differ significantly as suggested by the variance ratio of 1.39. This value suggests that the users identified relevant datasets on GIgateway more consistently than on STORM. We attribute this observation to possible confusion in some visualisations presented on STORM.

6.3 Evaluating GeoDOVE

The previous section was concerned with the STORM browser, and addressed issues of system and user-based performance. This section will discuss the evaluation of the GeoDOVE applet. GeoDOVE is an applet embedded within the NEEDS geoportal; therefore, a user evaluation of the applet without the geoportal would not be appropriate. Consequently, this section centres on system-based performance of the applet as it is intended for use from within a hosting geoportal. The tasks in this experiment were to retrieve a dataset and render it in a 3D. These operations were timed to determine the performance of the approach for direct database connectivity proposed in Chapter 3.

This is in contrast to the STORM browser which uses Java webstart to run as a standalone application, though it is also launched from the web. It should be noted that we are not evaluating the visualisation metaphor offered by GeoDOVE as this has already been implemented in previous related studies. Instead we are evaluating the approach in which the geovisualisation are created, that is through direct connectivity to an RDBMS serving OGC simple feature objects. A usability evaluation of GeoDOVE is not offered at this stage, as the applet is to be embedded in a geoportal. Instead, the applet's usability is assessed as part of the NEEDS evaluation.

6.3.1 Experimental Setup

To determine whether the architecture presented in this study offers acceptable system-performance, the evaluation was based on the suggestion that a ten second limit is required for users to keep their attention on a task (Nielsen, 2000). Therefore, the evaluation aimed to show that GeoDOVE performs tasks within this ten-second limit. It was important to test the application's ability to upload and model vector and raster datasets. Further, the size of a geographic dataset can have a significant bearing on the performance of an application; the study therefore involved the upload and visualisation of geographic datasets of varying sizes and extents. An additional performance indicator for 3D applications is the frame rate; therefore our study also monitored the frame rate as the user navigated through the visualisation.

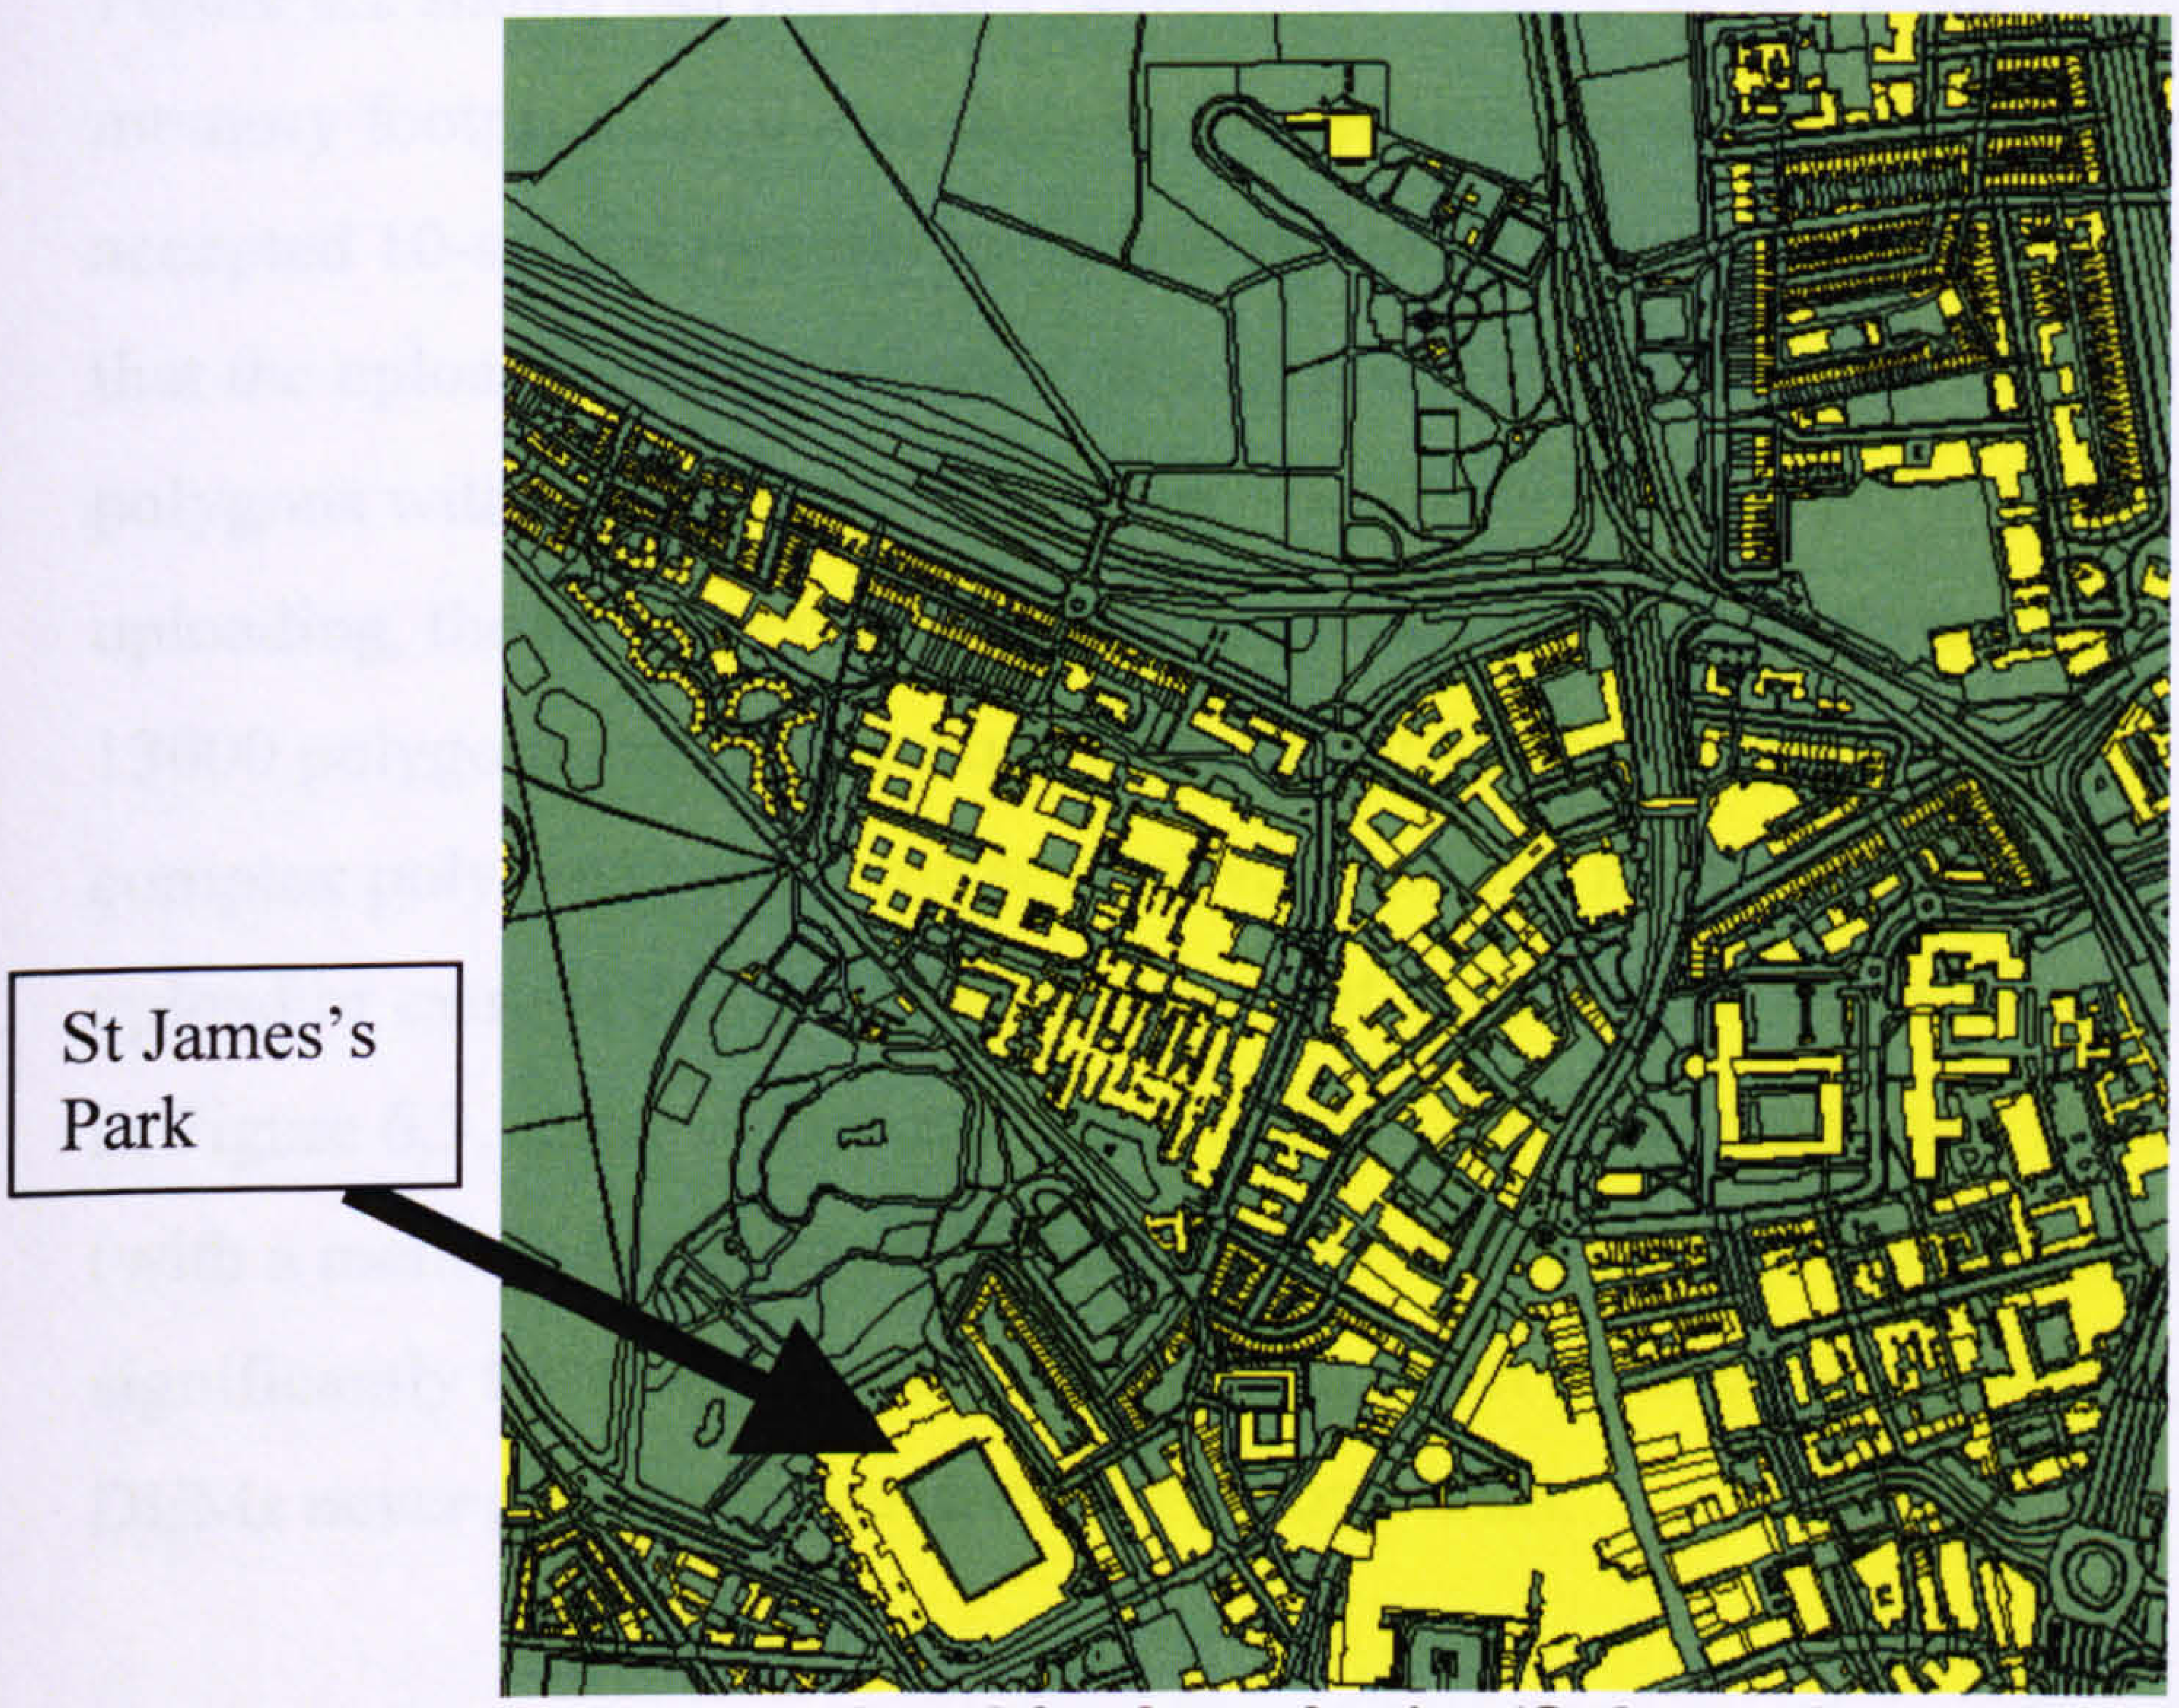

Figure 6.1 The study area selected for the evaluation (Ordnance Survey © Crown Copyright)

The experimental setup included Windows PCs with 3.0 GHz Intel processors. Each PC hosted a Java 1.5 virtual machine from which the applet was run. The upload task was timed from the moment a request for data was sent to the server. The visualisation task was timed from the moment the user instructed the system to create a 3D model of the vector or raster data. The vector data was modelled as extruded polygons and the raster data as DEMs, using the methodologies described in Chapter 3. The results of the upload and visualisation of vector data are presented in Figure 6.2. The vector datasets had differing file sizes on the database server, ranging from 0.4 megabytes for 1000 polygons to 15.2 megabytes for 14000 polygons. The vector datasets had coincidental centres to improve objectivity, thus every dataset contained all polygons that were contained in smaller datasets. Due to its high density of buildings, shown in Figure 6.1, the city of Newcastle upon Tyne was selected as the study area. The results for the upload and visualisation of raster data are presented in Figure 6.3. All the raster datasets covered the same geographic extent; however, they had varying pixel dimensions.

6.3.2 Results

Figure 6.2 shows that for vector datasets containing up to 12500 polygons (with a memory footprint of 10 megabytes), the application extruded polygons within the accepted 10-second (10 000 milliseconds) limit. It can be observed from the figure that the uploading of each vector dataset is slightly faster than the extrusion of polygons within the dataset, particularly as the number of polygons increases: for uploading, the limit specified is only exceeded when the dataset contains more than 13000 polygons (with a memory footprint of 11 megabytes). Unsurprisingly, complex polygons containing several vertices and multiple parts took longer to upload or extrude due to higher numbers of vertices and inner polygons. As indicated in Figure 6.3, raster upload exceeds our limit for files larger than 2100 x 2100 pixels (with a memory footprint of 15 megabytes); whereas creation of DEMs was significantly faster and more constant than raster upload. Surprisingly, creation of DEMs never appears to reach the 10-second limit.

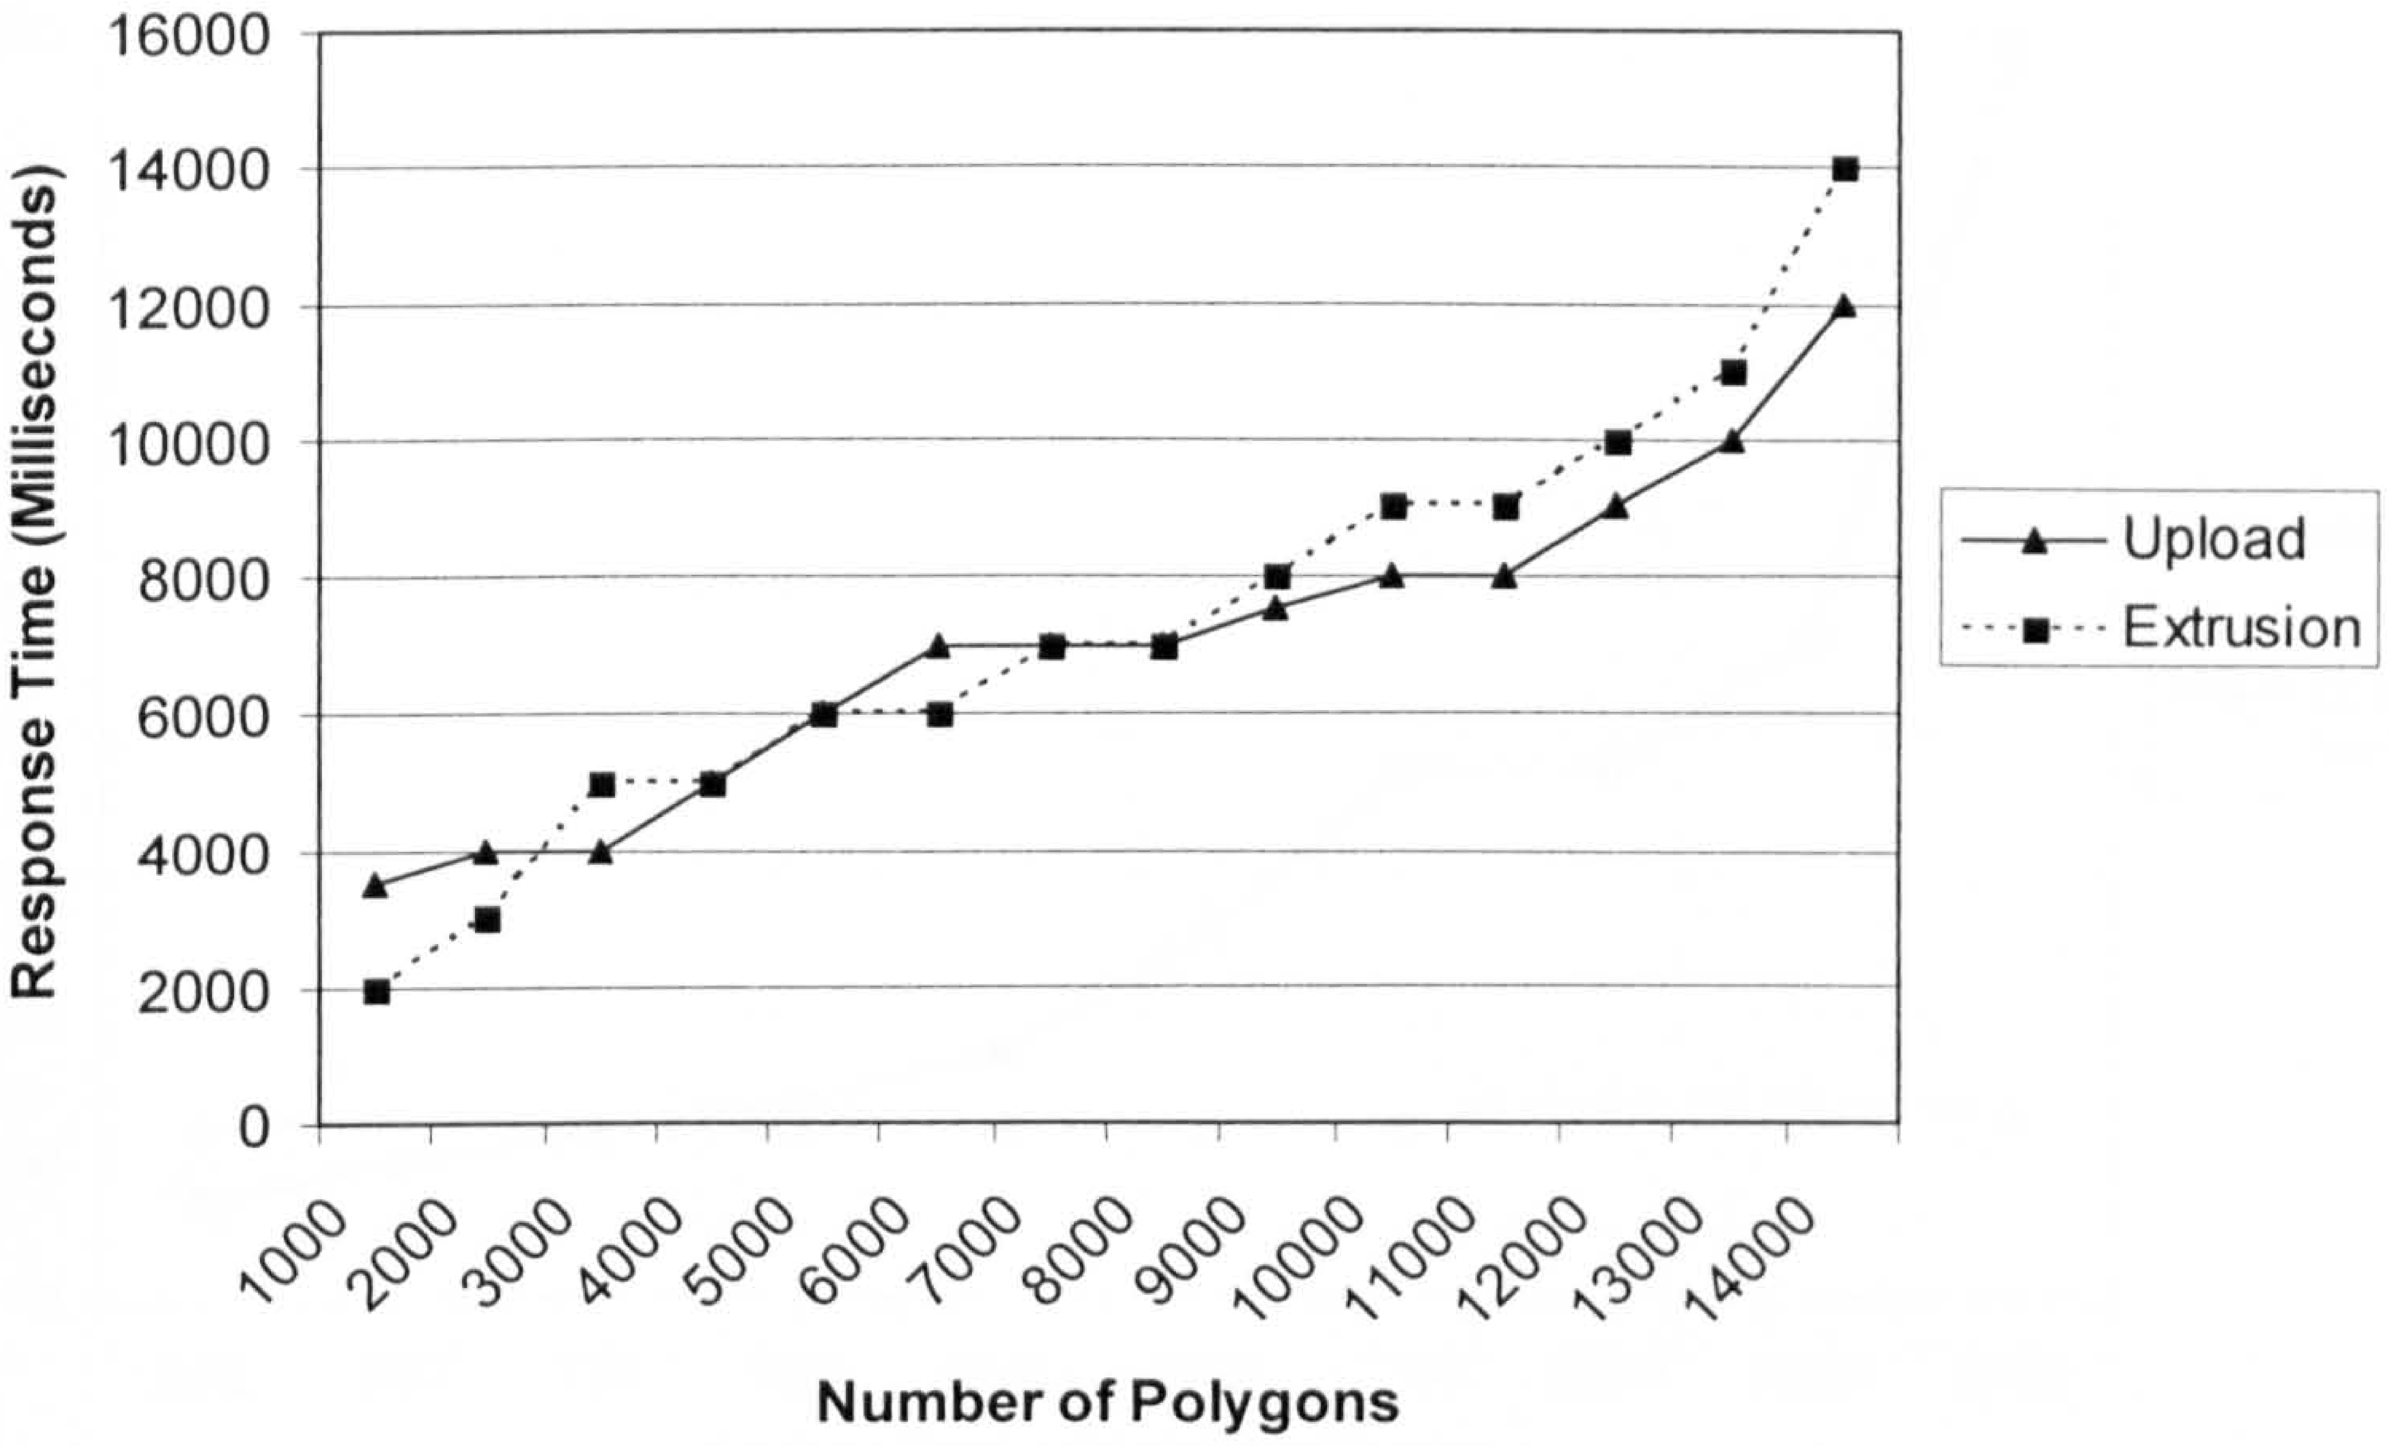

Figure 6.2 Response times for the upload and extrusion of vector layers

During the rendering of vector datasets, a frame rate ranging from 19 to 34 frames per second (FPS) was observed. The lower frame rate was recorded only when all extruded polygons were visible within the same field of view. The higher frame rate was observed when 'flying through' at low altitude as there were fewer extruded polygons within the field of view. "For smooth animation, a rate of 20 frames per second or more is desirable"(Selman, 2002: pp.104). As navigation through the visualisation requires smooth animation of the position of the viewer, the observed range in frame rate is acceptable. A higher frame rate was achieved by reducing the size of the applet to approximately 500 x 500 pixels.

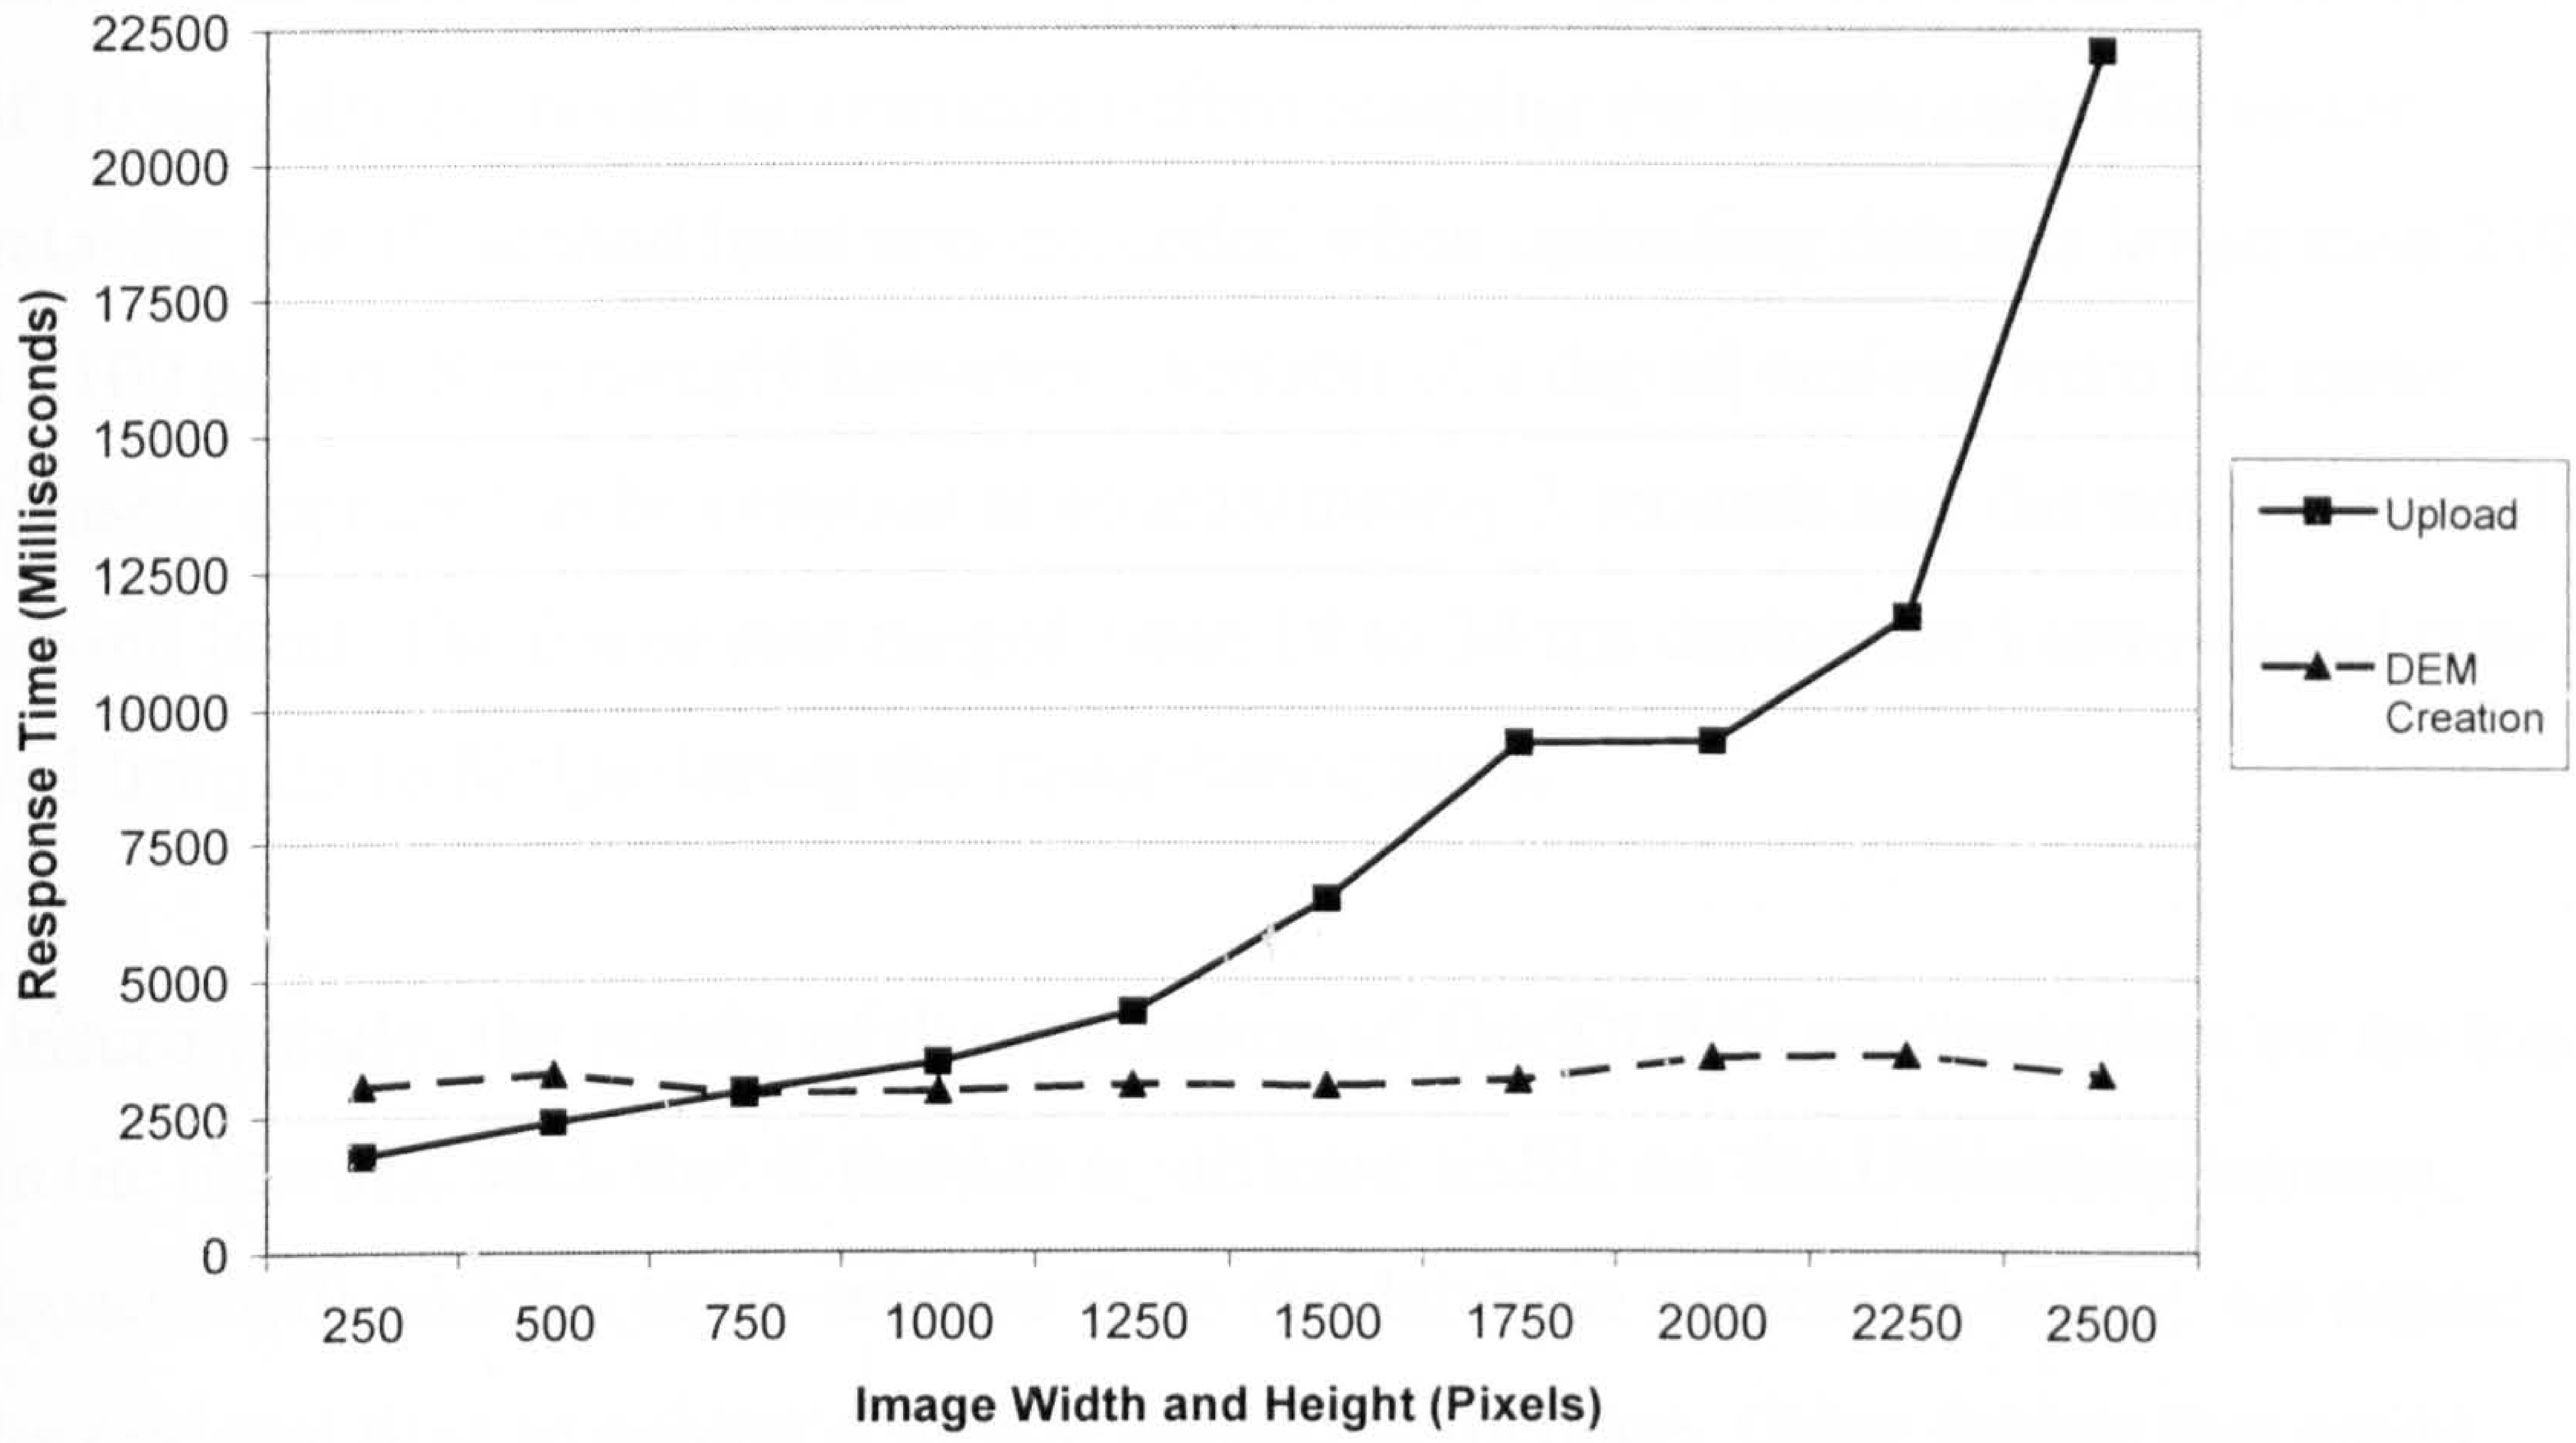

Figure 6.3 Response time for the upload and creation of DEMs from raster layers

The rendering of raster datasets on rectangular planes (flat polygons) produced the highest frame rate at 61 FPS. The rectangular planes were created prior to DEM creation to allow the user to select different visual properties before DEM creation. Unsurprisingly, draping raster datasets on DEMs reduced the frame rate to approximately 23 FPS. This decrease in frame rate is due to the higher number of triangles inherent in a DEM. As before, the frame rate increases when 'flying through' the DEM, depending on the number of triangles rendered in the field of view.

6.3.3 Reflective Summary

This subsection has presented a performance evaluation of GeoDOVE. As GeoDOVE runs embedded within a geoportal webpage, a usability study will be carried out as part of the evaluation of NEEDS. The performance evaluation involved the determination of response time upon loading geographic datasets of increasing size and geographic extent. The frame rate of the applet was also monitored to determine the 'smoothness' of animating the movement of the viewer. It was observed that up to 13000 vector polygons (with a memory footprint of 11 megabytes) could be uploaded before the response time exceeded a 10-second benchmark adopted for the study. Up to 12500 polygons (with a memory footprint of 10 megabytes) could be extruded before reaching the benchmark. For raster datasets, the 10-second limit was exceeded when uploading datasets larger than 2100 x 2100 pixels. Surprisingly however, creation of a digital surface from the raster datasets appeared to be constant at approximately 3 seconds and did not reach the 10 second limit. The frame rate ranged from 19 to 34 fps during the vector-based tests and from 23 to 61 fps during the raster-based tests.

Unsurprisingly, the results of the evaluation of GeoDOVE are dependent on the load on the network, such that if there is significant traffic on the University intranet, datasets will take longer to retrieve from the database engine. However, we expect the retrieval time to exhibit a similar pattern of increase. Other factors that could influence the results include the processes run by other programs on the testing machines. We have attempted to address both of these concerns by performing the tests after working hours when there is the least amount of network traffic on the intranet and when there are fewer processes running on the database server. However, these are noted as possible factors on the repeatability of the results.

6.4 Evaluation of NEEDS

This chapter has, so far, only discussed the evaluation of the individual visualisation components (STORM and GeoDOVE). However, as these components were intended to be integrated into a geoportal, it was necessary to evaluate the complete geoportal (NEEDS) as well. Consistent with studies by Benbunan-Fich (2001) and Blok (2005), the Think-Aloud approach was adopted for evaluating the usability of the NEEDS geoportal. The rest of this section describes the evaluation experiments and presents the observed results.

6.4.1 Experimental Setup

The evaluation sessions were conducted at the University of Newcastle upon Tyne. Each session included a notebook PC running an Intel 1.3 GHz processor. The notebook was connected to an audio/visual projector and the output from the projector was recorded using a video camera. In contrast to the experimental set up adopted by Blok (2005), this study did not include a video recording of the physical behaviour of the user because users can be intimidated by a video recorder, as suggested by Nielsen (2000). Instead the video recorder was intended to synchronize the verbalisations to events of the interaction between the user and the system. Figure 6.4 illustrates the experimental setup. Although audio/visual capturing software exists that can capture the output from the screen and optionally record the surrounding sounds, it was decided not to use the software to avoid overloading the PC with several concurrent processes. Use of a video camera therefore ensured that the PC's resources were devoted to running only the software (geoportal) under examination, in addition to the operating system and other 'house keeping' applications.

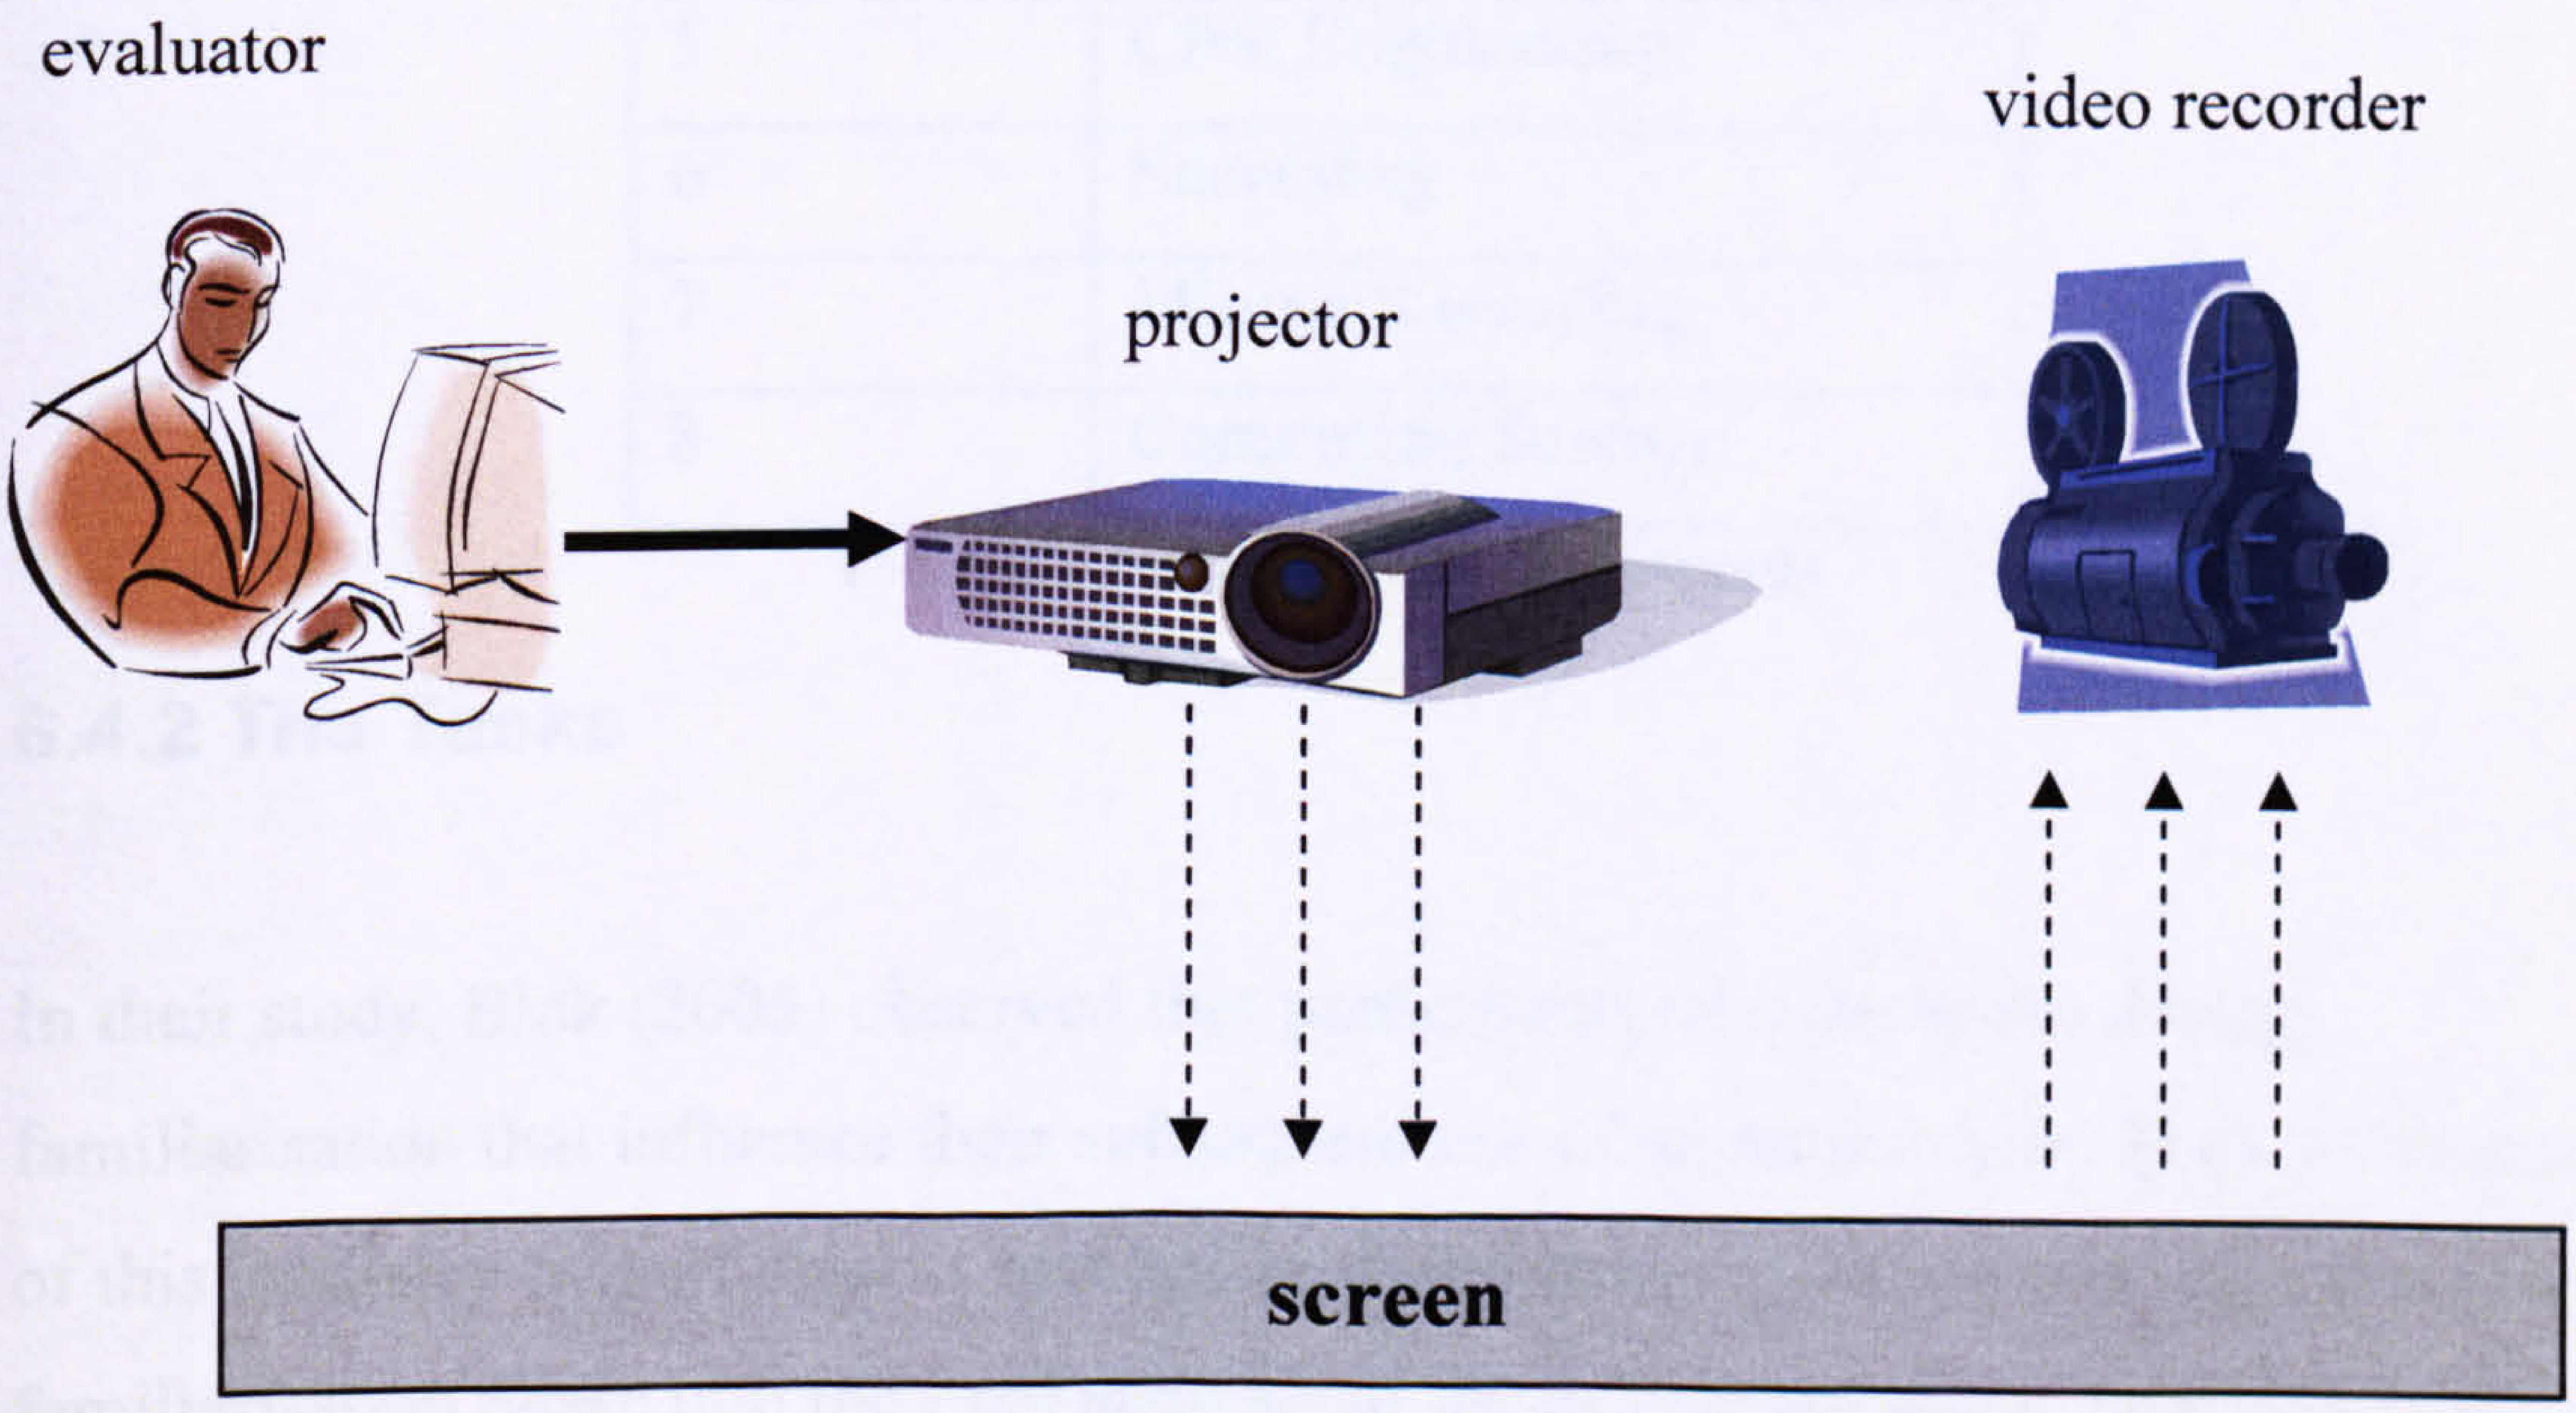

Figure 6.4 An illustration of the experimental setup of the think-aloud sessions

Participants were selected from the intended user community; this included students on geospatial or related courses, as listed on Table 6.4. This meant that participants were likely to be familiar with the concept of spatial containment, which the geoportal uses as a search constraint. A classic study by Jakob Nielsen (1994) recommends using five participants for a think-aloud evaluation with any additional users depending on i) the skills and experience of the experimenter ii) the number of iterations planned for the system, and iii) the financial impact of the use of the system. Consistent with recommendations in usability literature (Brink et al., 2002, Nielsen, 1994), a group of eight users was enlisted for our think-aloud evaluation. Four of the enlisted users had taken part in the STORM usability study and were thus familiar with the STORM browser. However, none of the participants had had prior exposure to GeoDOVE.

Participant	Discipline
1	Transport Engineering
2	Geography
3	Geographic Information Science
4	Mathematics and Computing Science
5	Civil Engineering
6	Surveying
7	Mining Surveying
8	Computing Science

Table 6.4 Participants' backgrounds

6.4.2 The Tasks

In their study, Blok (2005) observed that participants take decisions during familiarisation that influence their subsequent use of an application. Their realisation of this tendency by participants resulted in Blok having to incorporate data from the familiarisation phase into the final analysis of the evaluation study. Learning from their experience, the evaluation of NEEDS comprised of a single recorded session per user, which included both a familiarisation and a main element (comprising of tasks). To support the user in their familiarisation of the system, the documentation accompanying the evaluation exercise included explanations of the navigation controls. Integrating familiarisation with the actual experiment is supported in a study by Owen et al (2006). Their study aimed to examine how software designers employ available documentation for designing systems from existing components. They observed that training, prior to running a think-aloud evaluation, could have had an effect on the final outcome due to the related knowledge transfer from the experimenter to the user and hence they carried out their experiments without providing prior training. They contend that the integrity of their experiments was upheld by the lengthy think-aloud sessions they undertook.

Four tasks were prescribed for our experiment. The first two tasks required use of STORM and the last two tasks required use of GeoDOVE. However, all the tasks required starting from the NEEDS home page, navigating to the assigned tool (GeoDOVE or STORM) and interacting with that tool to identify datasets relevant to the scenarios presented in Appendix A. This process workflow ensured that each of the visualisation tools was examined with respect to its incorporation on the geoportal, in contrast to acting as a standalone application. In order to improve the objectivity and repeatability of the experiments, only datasets applicable to the four tasks were made accessible through the geoportal. Further, the users were prescribed specific queries to search on the geoportal. This ensured that the evaluation remained focussed on usability and not the retrieval performance (which is addressed in earlier sections of this chapter).

For the interaction with STORM, users were asked to verbalise their comprehension of the STORM visualisations. The researcher's comprehension of the visualisations presented during task 1 and 2 was used as a benchmark of the minimum information participants were expected to discover. Each task was considered achieved if a user correctly identified the most relevant datasets for each of the three axes. Additional reasonable interpretations of the visualisations from each user were added to the benchmark interpretation, for example identification of clusters. Therefore, the benchmark served as a guide to 'what the researcher expected the participants to discover'. With regard to GeoDOVE, participants were asked to identify datasets that were relevant to specified scenarios. For each of the two tasks involving GeoDOVE, a participant had to 1) examine the metadata of a dataset from the NEEDS interface 2) then navigate to the GeoDOVE page for that dataset 3) add a DEM (for Task 3) and a 3D model (for Task 4) then decide if that particular dataset was relevant to the specified scenario. Together the two tasks (3 and 4) required the user to examine approximately 10 datasets using GeoDOVE.

6.4.3 Results

In this subsection, a summary of results from the usability exercise is presented. As aforementioned, four tasks were prescribed for each participant. During execution of each task, the verbalisations and machine-user interactions were recorded using a video recorder. After each session the verbalisations were then analysed to identify usage with respect to i) navigation from the geoportal to the applets, ii) layout of the geoportal or the applets, iii) search capabilities and iv) interpretation of visualisations. Consistent with related think aloud studies by Blok (2005) and Fabrikant (2000), the time to completion of tasks was also recorded. Recording time to completion of tasks provided an indication of the usability of the system, allowing for lengthy procedures to be identified.

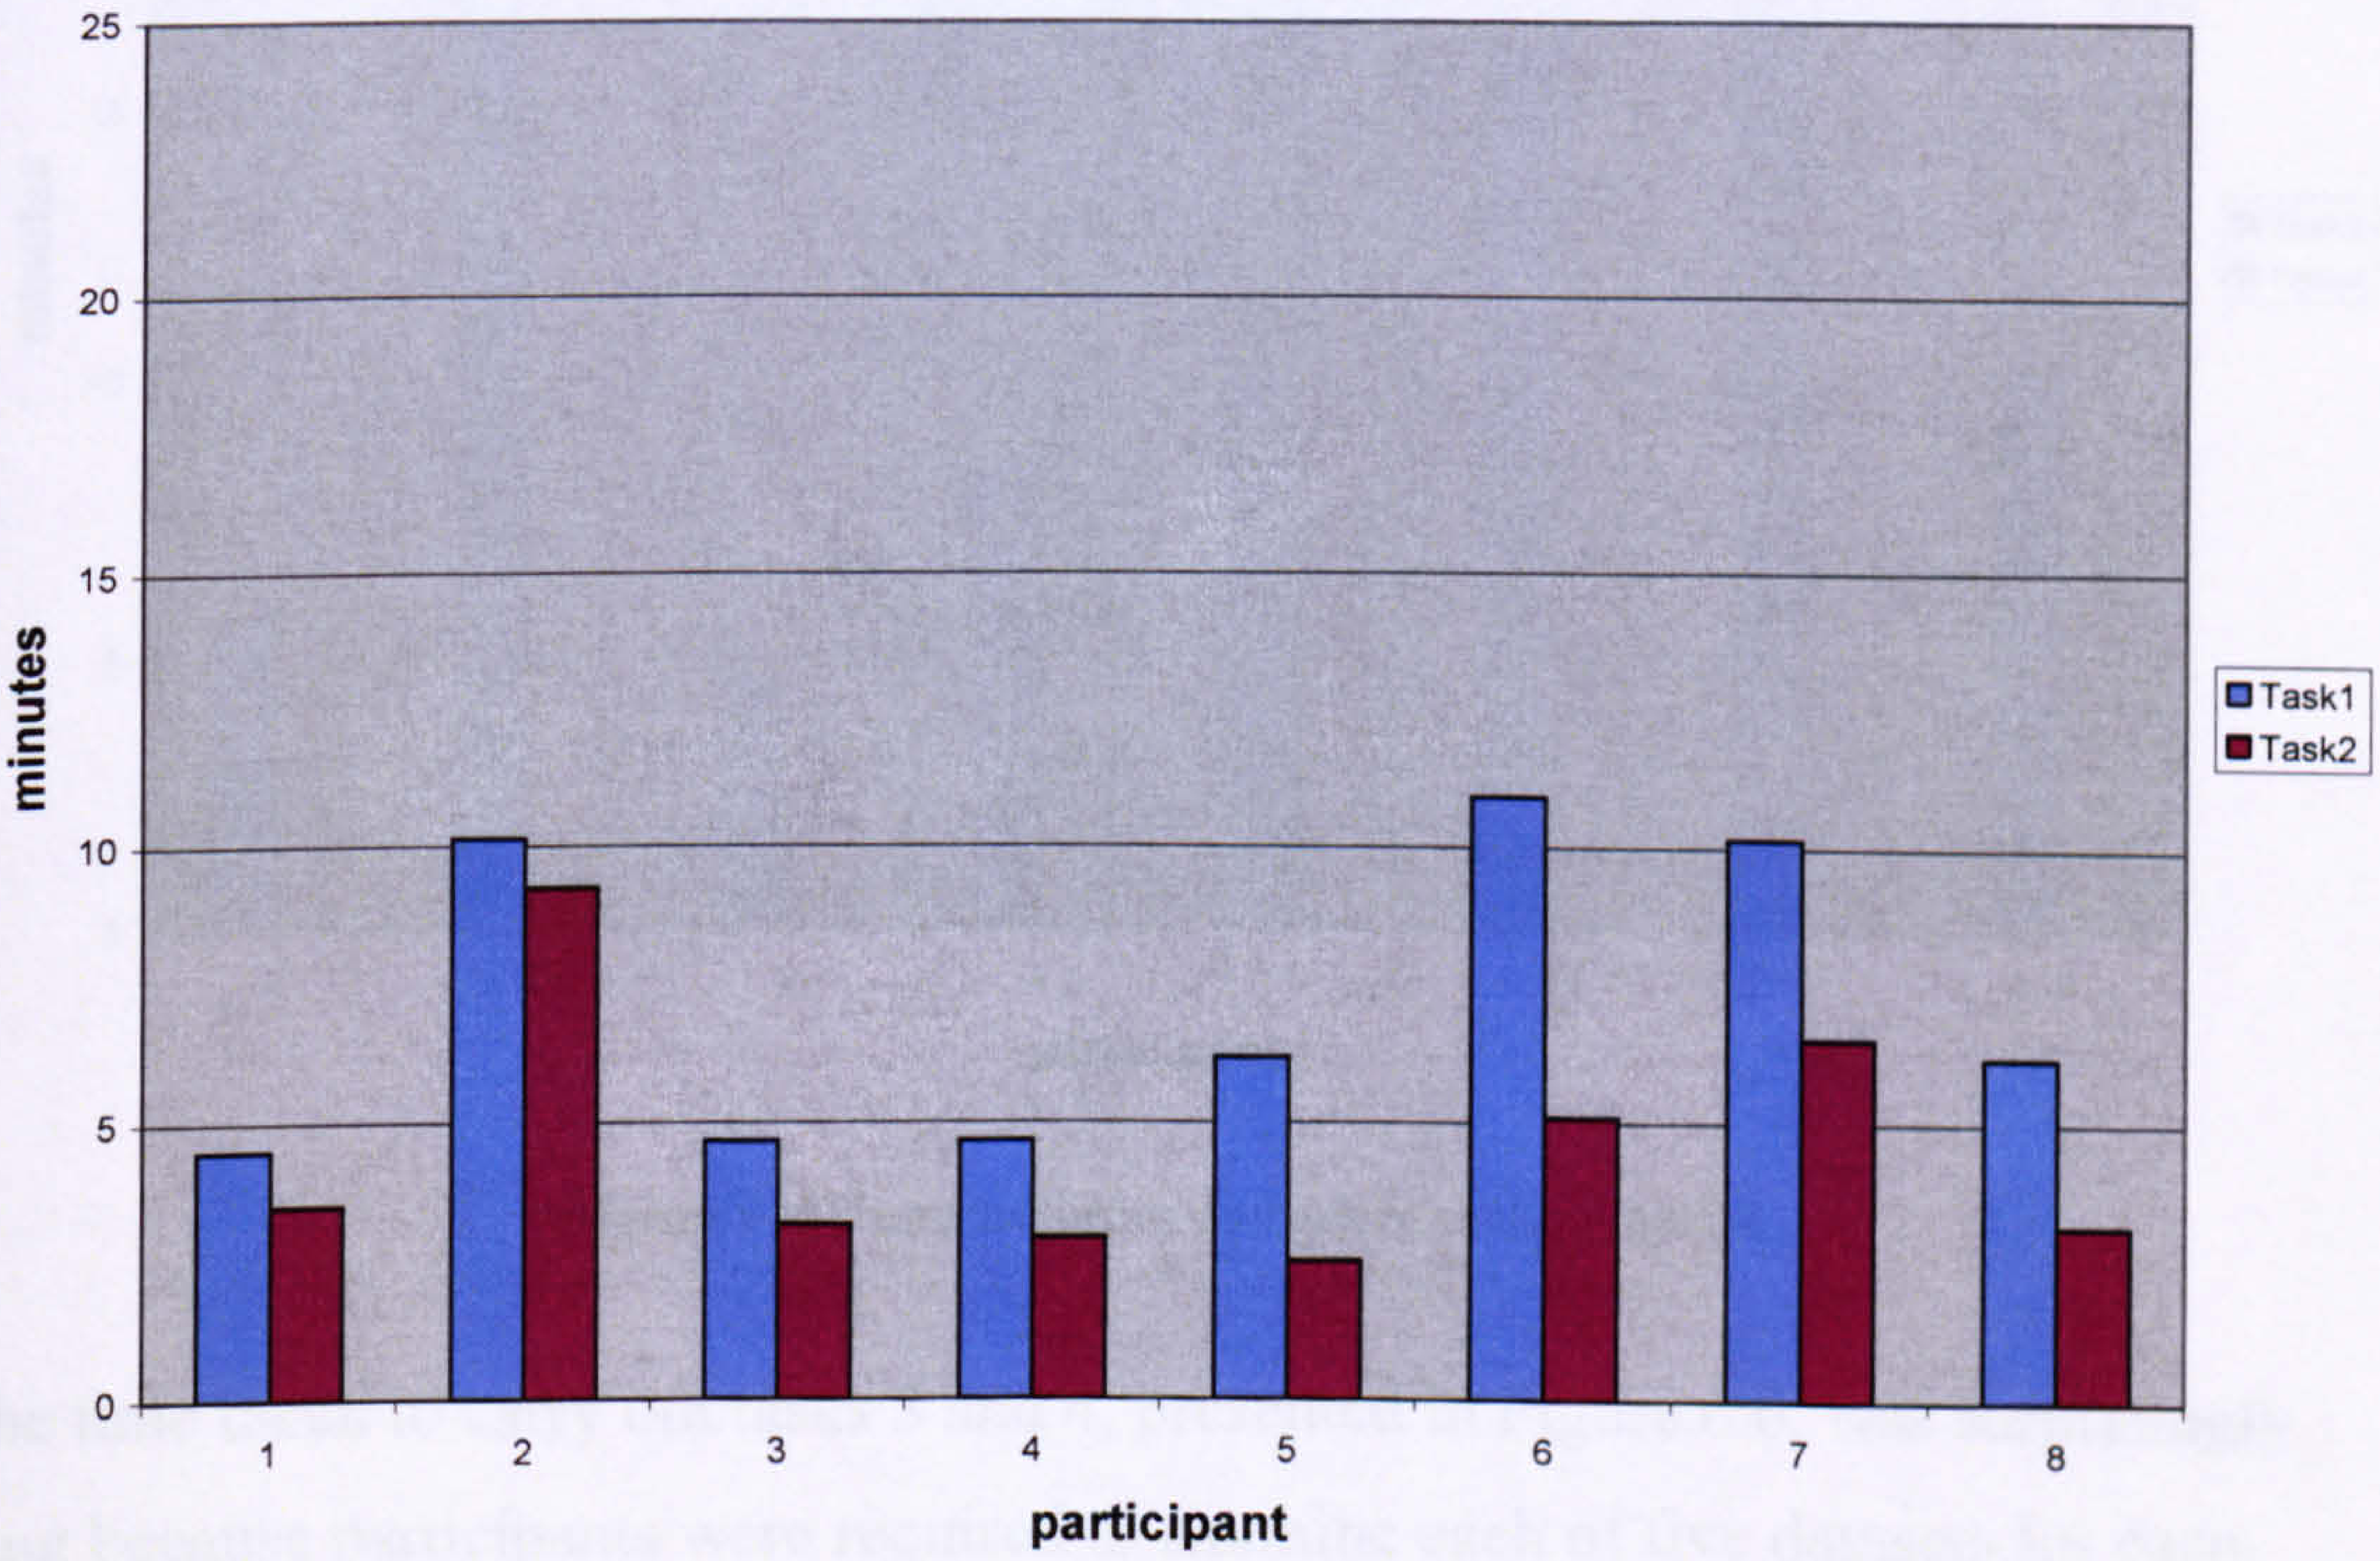

Figure 6.5 Time taken to complete tasks 1 and 2

The times taken to complete tasks 1 and 2 are presented in Figure 6.5. It should be noted that the observed time includes time spent on familiarisation and on simple exploration of the system. The results however, exclude unexpected interruptions such as loss of internet connection. For example, participant 2, who holds a bachelor's degree in geography, spent a significantly long time on tasks 1 and 2. His think-aloud session produced a high amount of data. In contrast, participant 3 spent relatively shorter time carrying out the tasks and consequently producing a relatively smaller amount of data. Fabrikant (2000: pp. 97) observed similar trends, with expert users spending more time on tasks than non-expert users, and suggested that "they were not only intrigued by the representation, but also seem to particularly enjoy to be able to explore and manipulate objects in 3D". This thesis supports the suggestion and further attributes the longer durations to a higher understanding of geographic concepts leading to more issues being noticed by experts.

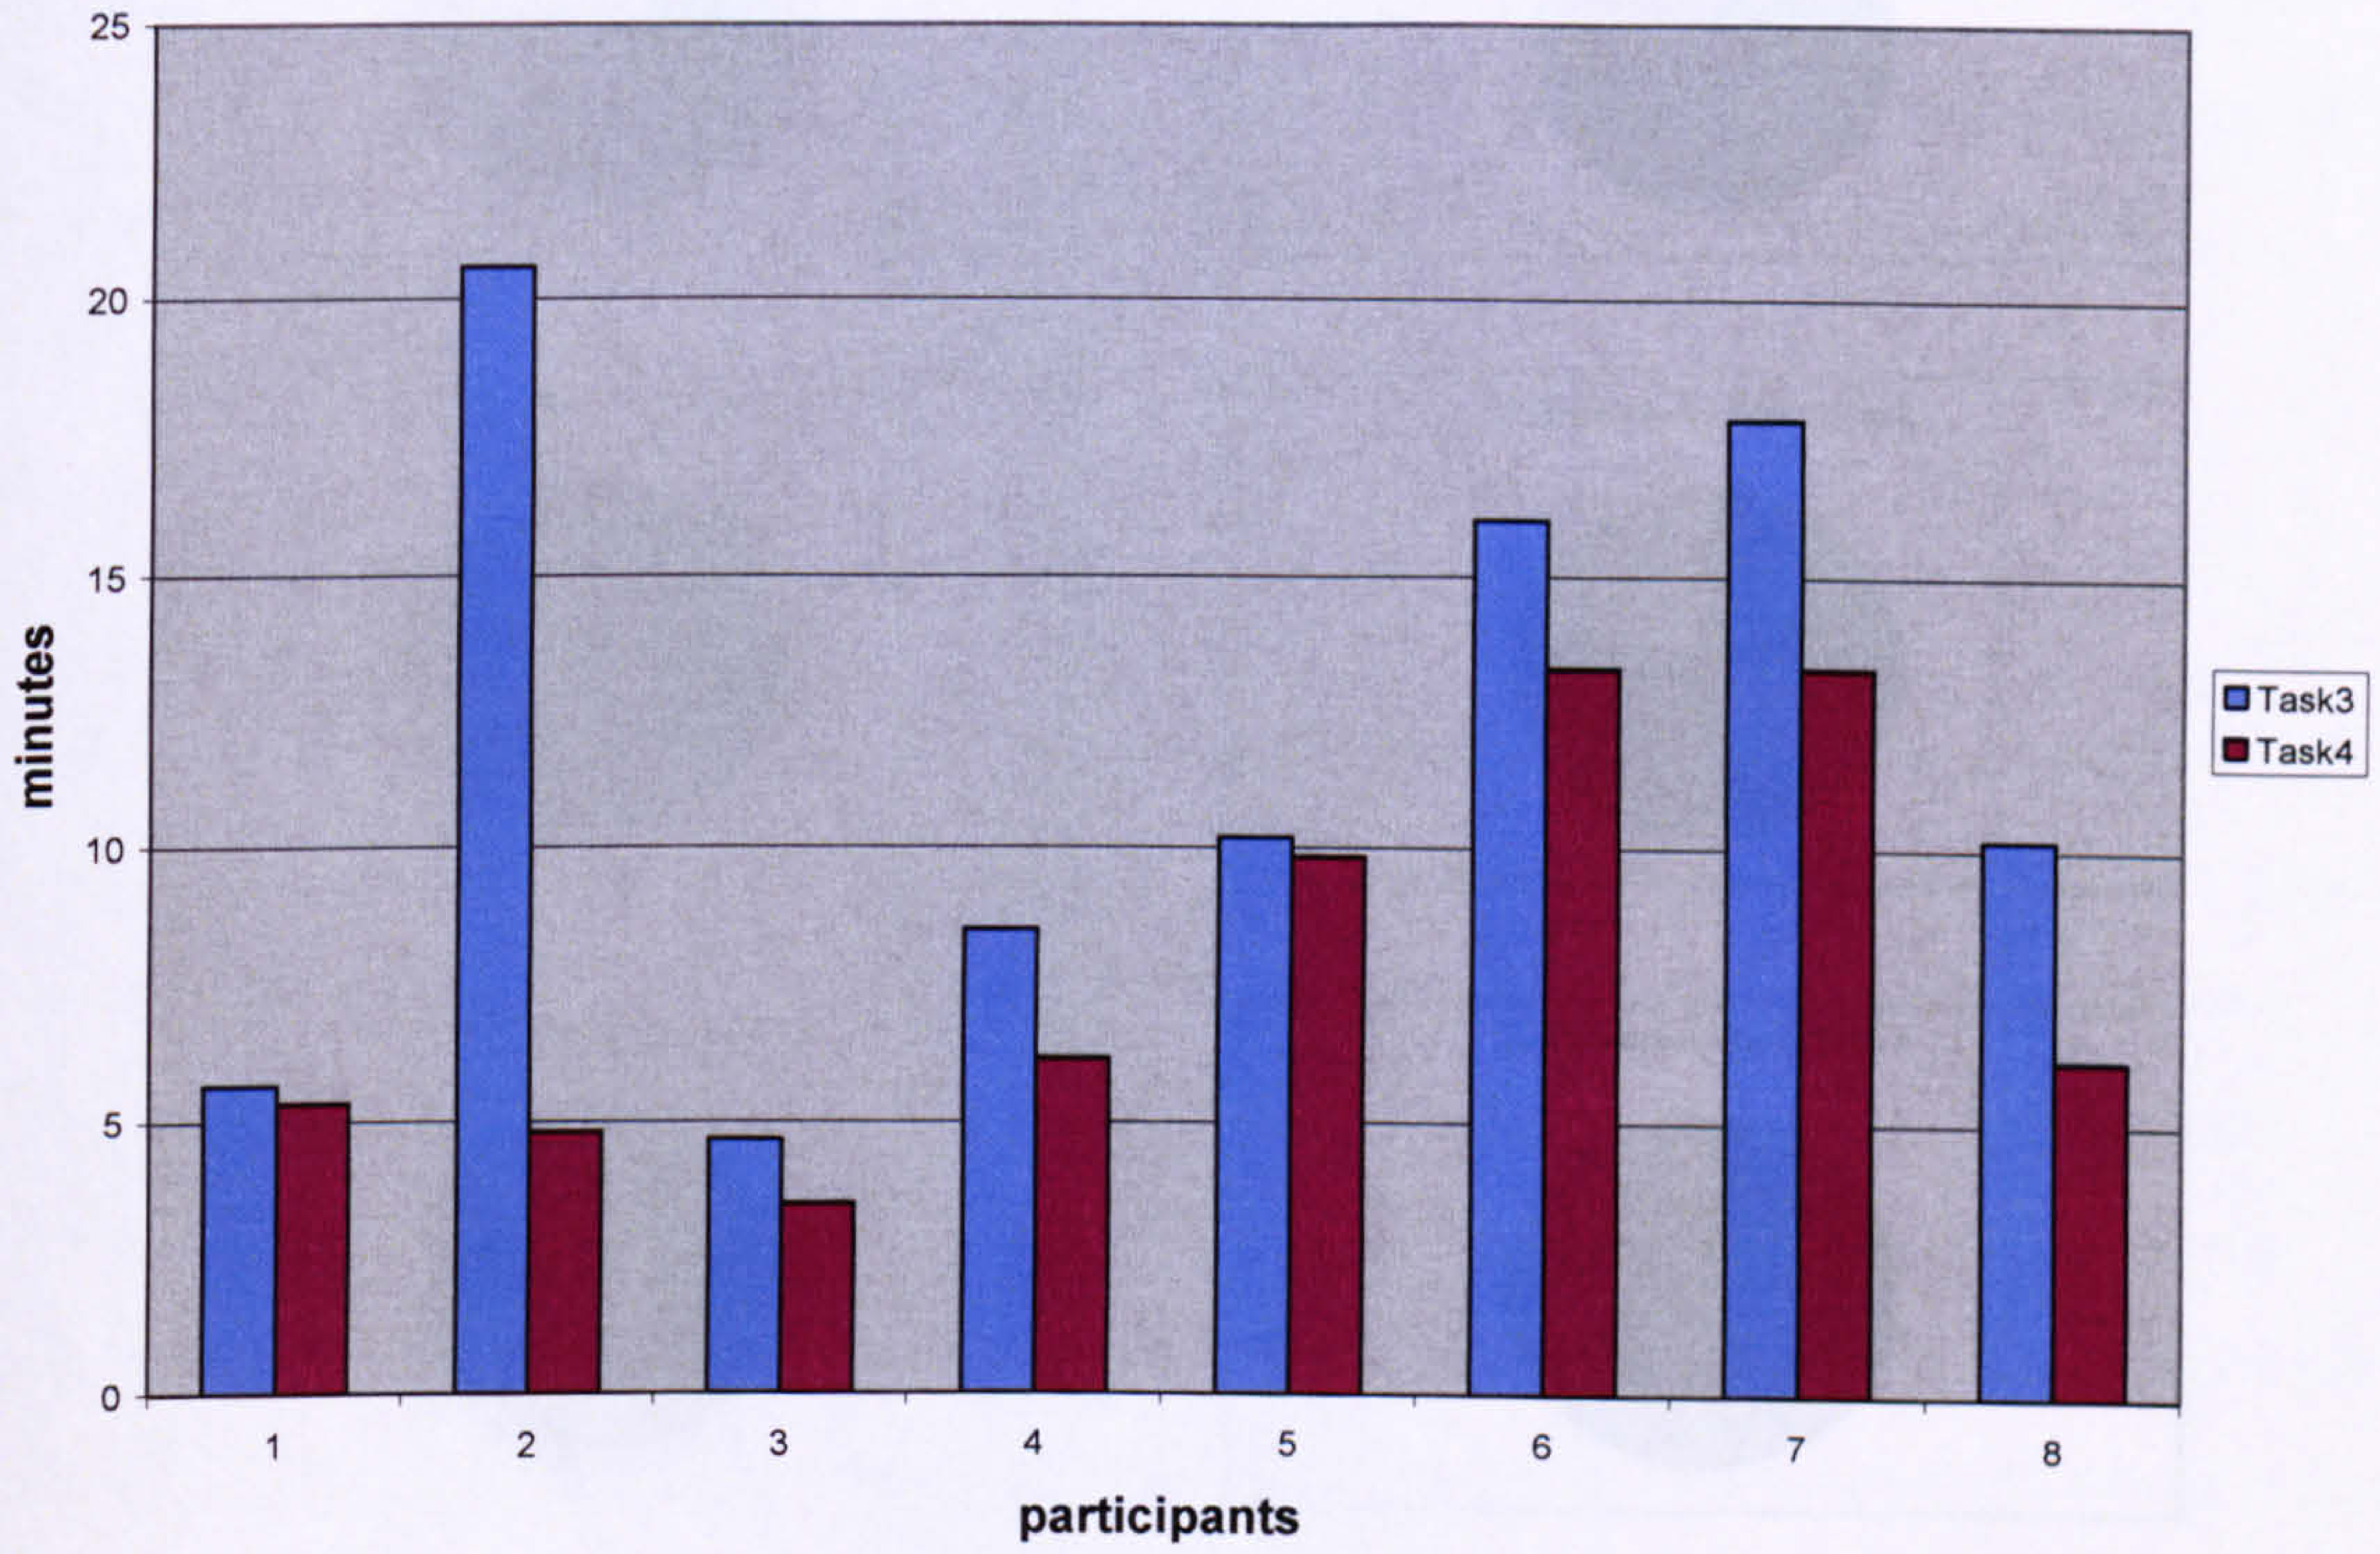

Figure 6.6 Time taken to complete tasks 3 and 4

The time taken to carry out tasks 3 and 4, presented in Figure 6.6, was surprisingly long because participants were required to examine each of five datasets for each task and then to navigate back to the search results to select the next dataset to examine. Two of the participants consequently suggested that it be made possible to upload all search results into the applet in order to avoid having to exit the applet each time a new dataset had to be viewed. Further observations made by the researcher for each session are presented in Appendix D. The text refers to coordinate (1,1,1) of the STORM visualisation as the 'optimum position'.

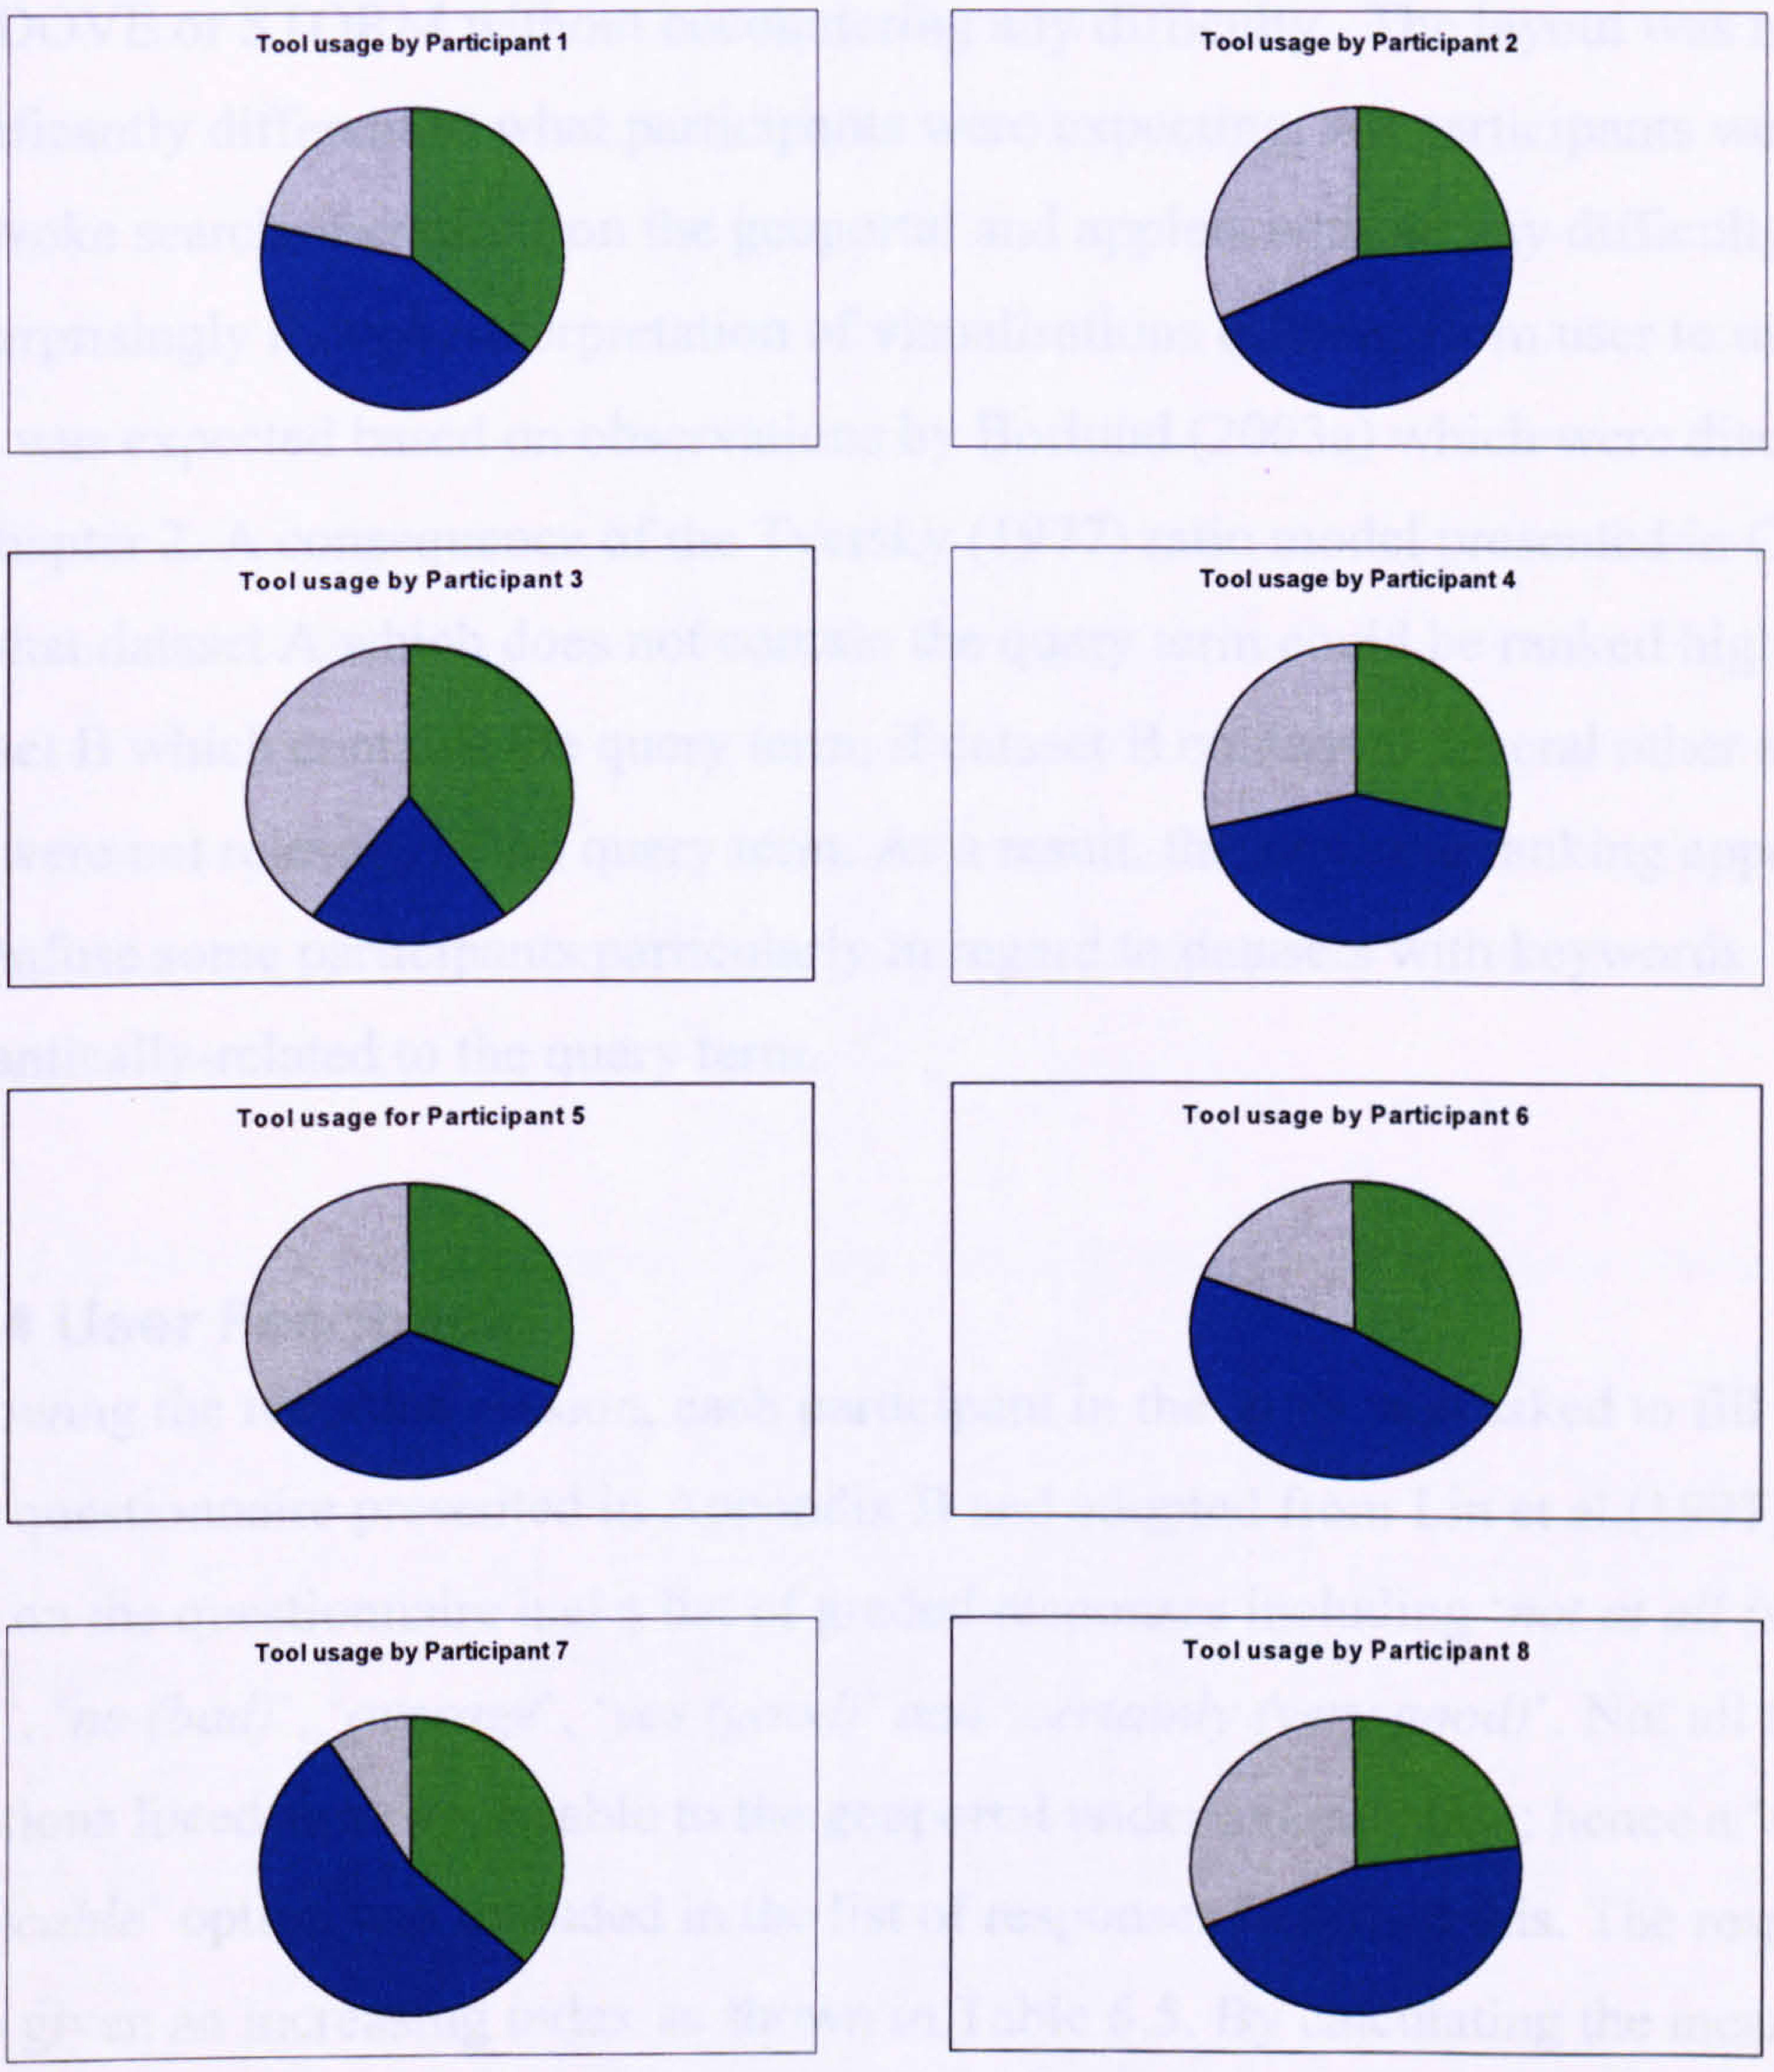

Figure 6.7 Samples of tool usage by the participants whilst carrying out tasks 1 and 2 (green:ranked list, grey: metadata panel, blue: 3D visualisation panel)

The number of times tools were invoked when carrying out task 1 and 2 was recorded and the proportion of tool usage is presented in Figure 6.7. This included instances when the mouse pointer was hovering over a user interface control, coupled with verbalisations about that control. It was envisaged that this would provide insight into user preference. The results suggest that the 3D visualisation on STORM drew most of participants' attention. However, all three user interface components (the 3D visualisation, the ranked list and the metadata panel) appeared to contribute to each participant's final judgement of relevance.

In summary, all participants were able to navigate from the main geoportal to either GeoDOVE or STORM without encountering any difficulty. The layout was not significantly different to what participants were expecting. All participants were able to invoke search operations on the geoportal and applets without any difficulty. Unsurprisingly though, interpretation of visualisations differed from user to user. This was expected based on observations by Borlund (2003a) which were discussed in Chapter 2. A consequence of the Tversky (1977) ratio model presented in Chapter 2 is that dataset A which does not contain the query term could be ranked higher than dataset B which contains the query term, if dataset B contained several other terms that were not relevant to the query term. As a result, the semantic ranking appeared to confuse some participants particularly in regard to datasets with keywords semantically-related to the query term.

6.4.4 User Feedback

Following the recorded session, each participant in the study was asked to fill a 100-item questionnaire presented in Appendix B and adopted from Lin et al.(1997). Each item on the questionnaire had a list of graded responses including '*not at all (very bad)*', '*no (bad)*', '*average*', '*yes (good)*' and '*certainly (very good)*'. Not all the questions listed were applicable to the geoportal under examination; hence a '*not applicable*' option was included in the list of responses to reflect this. The responses were given an increasing index as shown in Table 6.5. By calculating the mean index for each question, it was possible to identify positive, negative and average responses; positive responses defined as those that resulted in a mean of greater than three and negative responses as those with a mean less than three. A question was excluded from the analysis if two or more users had marked it as being 'not applicable'.

response	*not applicable*	*not at all (very bad)*	*no (bad)*	*average*	*yes (good)*	*certainly (very good)*
Index	0	1	2	3	4	5

Table 6.5 Scaled responses on the questionnaire

Most of the questions regarding compatibility of the system to users' experience received highly supportive responses. Similarly supportive responses were received for questions regarding consistency of the user interface. This suggests that the geoportal (including STORM and GeoDOVE) did not overwhelm the users and was thus consistent with their expectations and experience. Further, the user interfaces were intentionally made consistent with that of existing geoportals and geovisualisation applications so as to reduce the time required for familiarisation. Responses regarding the potential of the participant to learn how to use the system suggested that there was clarity of wording and the ordering and grouping of the menu options was logical. These responses are consistent with the video recordings which show all participants managing to navigate from the NEEDS home page to each visualisation application.

It should be noted that several of the questions that received low responses were recorded as being not applicable to the system. However, some of the questions were concerned with the automation of minimal action, for example the provision of default values or function keys for frequent control entries. The responses are particularly relevant for the visualisation applications, as the provision of function keys for viewing from specific directions could effectively improve usability. Other features for minimising actions could be the automated suggestion of keywords as a user types in a search term. Such a facility requires the ability of the browser to retrieve keywords from a server-side database for each key-press. Responses to questions regarding error feedback suggest that more information is required in the event of errors occurring.

6.4.5 Reflective Summary

With both STORM and GeoDOVE incorporated into the NEEDS geoportal, a usability study of the complete system was carried out and presented in this subsection. Whereas the evaluation of the geoportal does not directly address the hypothesis of this study, it examines the outcome of integrating the proposed visualisation approaches within a geoportal. The think-aloud method (protocol analysis) was adopted for the usability study and augmented with a questionnaire. Observations made during the think-aloud sessions showed that all the participants managed to navigate from the home page of the geoportal to the visualisation components (STORM and GeoDOVE). The participants were able to comfortably navigate between the visualisation components. Most of the participants preferred to view the visualisations from oblique angles, however one of the participants preferred 2D orthographic views of the data on both STORM and GeoDOVE. The time taken to achieve the tasks involving GeoDOVE was disappointingly long. Observations indicate that this was due to navigating between the search results and the applet for each possibly relevant dataset. Some of the feedback from users, suggested passing references to all datasets from the search results to the visualisation applet, to avoid having to restart the applet.

6.6 Conclusion

This chapter has presented evaluations of two visualisation applications (STORM and GeoDOVE) and a geoportal(NEEDS). The two visualisation applications are first evaluated independently of the geoportal to determine system performance, then evaluated as part of the complete geoportal to determine if a geoportal that hosts the applications would be acceptable to users. The evaluation of STORM addressed issues of retrieval precision and usability whereas the evaluation of GeoDOVE employed a determination of response time for two sets of raster and vector geographic datasets. The evaluation of the NEEDS geoportal was concerned with the usability of the geoportal once STORM and GeoDOVE had been deployed on the geoportal.

The results presented in this chapter suggest that the STORM browser enhances the retrieval performance of a GIR system and is acceptable to users for supporting spatial data discovery. However, the small improvement in precision suggests that it is more appropriate as a complement rather than as a replacement to existing approaches. The results also suggest that GeoDOVE is an acceptable application for presenting 3D geovisualisations during spatial data discovery; however it inherits limitations of speed characteristic of Java applets. Regarding the NEEDS geoportal, the results suggest that usability is maintained even when the visualisation applications are deployed on the geoportal. However, improvements could be made to how search results are passed from NEEDS to GeoDOVE, so that users do not need to frequently alternate between the two when examining several datasets.

Chapter 7 Conclusions

7.1 Thesis Overview

This thesis has considered the design and development of visualisation applications for improving the discovery and dissemination of multidimensional geographic information. An overview of geographic information retrieval was offered in Chapter 2 in order to provide a background of considerations in spatial data discovery. A discussion of visualisation principles with respect to spatial data discovery was presented in Chapter 3. An approach for presenting the results of a geographic search was proposed in Chapter 3. An approach for visualising geographic data on web-deployable Java applets was also proposed in Chapter 3. Chapter 4 examined methodologies for disseminating multidimensional geographic data through web-based technologies. The technologies include relational database management systems (RDBMS) and geographic web services.

In Chapter 5, the design and development of prototypes based on the visualisation approaches proposed in Chapter 3 was discussed. The development of a geoportal that consumes data disseminated using the approaches discussed in Chapter 4 was also discussed. An evaluation of the visualisation prototypes and the geoportal was presented in Chapter 6. First, the prototypes were evaluated independently of the geoportal, and then evaluated as part of the geoportal to determine usability and performance.

The results of this study have presented an application of ontology-assisted multidimensional visualisation that performs better than the traditional ranked list approach (in precision). However, visual metaphors can take a variety of forms and hence we cannot generalise to cover all ontology-assisted visualisations. Therefore, we qualify the hypothesis by concluding that *it is possible* for geoportals offering ontology-assisted multidimensional visualisation services to enable users to discover and retrieve more relevant geographic information resources than those that do not offer these services.

7.2 Revisiting the Research Questions

In Chapter 1, five research questions were presented. This subsection revisits the research questions to determine if they have been answered.

- **What are the limitations of existing approaches in the visualisation of candidate datasets during geographic data discovery?**

The thesis examined geographic data discovery from the perspective of SDI, spatial data clearinghouses and geoportals in Chapter 1. The role of geoportals as gateways to spatial data clearinghouses was highlighted. A detailed discussion of advances in Geographic Information Retrieval (GIR) was given in Chapter 2. Limitations to the presentation of the results of a geographic search using textual ranked lists were highlighted in Chapter 3; the main limitation being the reduction of multidimensional geographic relevance into a single dimensional relevance value. The thesis bases its argument on the fundamental definition of geographic data as linking a place, time and theme (Longley et al., 2001).

Limitations of existing approaches for web-based 3D geovisualisation using scene description languages, such as VRML, were highlighted in Chapter 3. Particular emphases were made on the lack of support for double-precision floating point coordinates in VRML. Support for a very limited set of coordinate sytems was highlighted as another limitation of VRML (specifically, its geographic profile GeoVRML) and its successor X3D.

- **How can ontology be used to support visualisations of geographic search results?**

An approach for visualising the results of a geographic search within a three-dimensional visualisation environment was proposed. Each axis within 3D space is assigned to present one of spatial, temporal or semantic relevance measures. The approach was named the Spatio-Temporal Ontological Relevance Model (STORM) based on its use of ontologies for calculating semantic relevance. The design and implementation of a prototype system was discussed in Chapter 5. The system includes a STORM browser, ontological query expansion tool and a metadata server. The prototype is able to extract metadata from OGC catalogue services based on the Z39.50 protocol.

- **From search results, how can a multidimensional geographic dataset be visualised in greater detail to determine its relevance?**

An approach for web-based 3D geovisualisation was proposed and implemented in an applet called GeoDOVE. The applet is called the Geospatial Database Online Visualisation Environment (GeoDOVE) and implements the Simple Features geometry model by the OGC. By retrieving data directly from a relational database server, the applet is able to offer support for several more coordinate systems than traditional geovisualisation approaches based on VRML. As the applet is implemented using Java3D, it is also able to support double-precision floating numbers, thereby overcoming the aforementioned limitation of VRML. Web-based methodologies for serving 3D geographic data using web services and database servers were discussed in Chapter 4, these serve as data sources for GeoDOVE. The proposed approach therefore addresses several of the limitations of VRML/X3D identified in this thesis.

- **Could the suggested approaches be incorporated into a conventional geoportal?**

The NEEDS geoportal was implemented using conventional technologies such as HTML, JSP and servlets. The traditional search component uses the PHP/Yaz toolkit to retrieve metadata records from external servers. By implementing a basic geoportal through these technologies, it is envisaged that the prototype visualisation applications will be transferable to any of the several geoportals that use these traditional web technologies. The STORM browser integrated into the geoportal through the inclusion of a hyperlink that invokes the Java webstart program. Therefore, the STORM browser offers an alternative search facility to the conventional one offered by the geoportal. As the GeoDOVE applet is specifically for geovisualisation during discovery time, it was embedded within a web page that is accessed through the metadata viewing page. Parameters are automatically passed from the geoportal to the applet such that a data source is automatically uploaded when the applet is launched.

- **What effect would the suggested approaches have on the performance and usability of a geoportal?**

The visualisation components were evaluated independently of the geoportal to isolate issues specific to the prototypes. By virtue of being a GIR application, the STORM system was evaluated using IR principles to determine the retrieval performance in terms of the precision. Based on suggestions in related studies by Nielsen (2000, 1994), a group of ten users took part in the study. The results suggest that the STORM browser improves retrieval performance by approximately 9.9%. Similarly, the GeoDOVE applet was also evaluated independently of the geoportal. The response time with respect to the upload and 3D modelling of vector and raster geographic datasets was monitored. The results suggest that the applet offers acceptable performance for vector datasets of up to approximately 11 megabytes, after which the response time exceeds an adopted ten-second benchmark. Raster upload exceeds the benchmark for datasets larger than 2100 x 2100 pixels. Surprisingly, the creation of DEMs from raster datasets does not appear to exceed the benchmark for all the dataset sizes observed.

Having determined the performance of the prototype visualisation applications, their usability from within the geoportal was evaluated next. The think-aloud method (also known as protocol analysis) was employed for the usability study. This generated a large amount of data, though most of it produced qualitative results. However, the think-aloud sessions revealed trends in user behaviour that would have been impossible to discover using other approaches. For example, processes that took a long time to complete due to repetitive actions were identified. The think-aloud study was augmented with a questionnaire to provide more targeted questions regarding the geoportal.

7.3 Major Findings

This thesis proposed two approaches for visualising the results of a geographic search and for presenting geographic data during data discovery. The STORM system displays documents returned by a geographic search according to their inferred spatial, temporal and semantic relevance. The relevance scores are plotted in a three-dimensional graphics environment to allow the searcher to simultaneously evaluate all three scores. The GeoDOVE system allows the searcher to preview candidate datasets in a 3D graphics environment as Digital Elevation Models or as textured 3D objects. The prototypes have been implemented using a variety of programming platforms, for example Java and PHP, and incorporated into a geoportal named the North East Environmental Data Server (NEEDS).

Evaluation experiments of the STORM browser support the hypothesis that geoportals that offer ontology-assisted multidimensional visualisations potentially can retrieve more relevant geographic datasets than those that do not offer these facilities. The observed ability of users to discern similarity from the multidimensional visualisations is consistent with conclusions by Leuski and Allan (2004: pp. 282) who concluded that "users have no difficulty grasping the idea of spatial proximity as the metaphor for inter-object similarity". The degree of improvement in precision observed during the evaluation of the STORM browser suggests that the approach is better suited as a complement rather than a replacement of the existing ranked list.

The GeoDOVE applet supports remote invocations of spatially enabled database servers. This capability allows an applet to use some of the computational resources available on the server-side and thereby reducing the load on a client. The ability to support database servers, offers the applet access to several more coordinate systems than those supported by existing geovisualisation approaches based on scene description languages such as GeoVRML. The performance-based experiments of the GeoDOVE applet suggest that Java3D applets can support multidimensional geovisualisation for data discovery purposes.

Results of the evaluation of the complete geoportal, after integration with STORM and GeoDOVE, suggests that the visualisation prototypes continue to be usable after being deployed on a geoportal. This demonstrates that the visualisation prototypes can be integrated with existing geoportals, thereby providing further support for the stated hypothesis. Results of the think-aloud study carried out during the evaluation of the geoportal, showed that all participants were able to navigate GeoDOVE geovisualisations comfortably, although with varying levels of proficiency. This observation contradicts suggestions by Nielsen (2000) that 3D visualisation is inappropriate for web-based applications because 3D graphics are inherently difficult to navigate.

7.4 Future Work

This thesis has addressed issues concerning geovisualisation support for the discovery and delivery of multidimensional geographic data. Although the evaluation of the STORM browser demonstrated an improvement in precision and the deployment of both STORM and GeoDOVE on the NEEDS geoportal demonstrates the feasibility of use of the proposed methods within existing geoportals, there are some issues that need to be addressed by future research. This section presents some of these and thus addresses the final research objective identified in Chapter 1.

7.4.1 Multidimensional Visualisation-Support for GIR

Several approaches for multidimensional visualisation support for GIR have been proposed by various studies including this one, as discussed in Chapter 3. Although each approach offers its own unique metaphor, most of them are presented within a three dimensional visual environment. Further, some of the visualisation approaches such as STORM and GeoCrystal are based on searches on geographic data collections. This suggests that it may be possible to integrate all or most of the approaches into a single application, allowing users to switch between visual metaphors at will. Such an application could for example, allow users to switch between the STORM, GeoCrystal or the Lighthouse visualisations depending on the information need or their understanding of the visualisation metaphor.

7.4.2 Extended Use of Ontology in GIR

As was noted in Chapter 2, ontologies have been used in GIR for deriving semantics from keywords and for determining the locations of places from keywords. The chapter emphasized the value of ontologies over historic approaches such as gazetteers. Although several studies have examined the design and development of ontologies for GIR and GIS applications, few of those studies have examined the use of ontology-supported inference engines for the determination of spatial relationships. Further research in this area should consider the inference of location based on known spatial relationships, for example, given three entities A, B and C, if A and C both intersect B and B is 100 metres long then it suggests that C is within a 100 meters of A. Although, it is evident that such inferences may not result in accurate geographic coordinates, the density of points in gazetteers suggests that a group of such inferences may result in reasonable inferences of location from ontology. A possible inference engine that could be used for this purpose is Cyc, which was introduced in Chapter 2. In addition to hypernym and hyponym relations, Cyc also allows for the addition of user-defined predicates. Predicates are used to construct sentences which can include at least two concepts and a relation between the concepts. For example, the sentence *((capitalOf, London, UnitedKingdom))* could mean London is the capital of the United Kingdom. In this case, the relation is neither a hypernym nor a hyponym but a predicate. Cyc contains several of such sentences. Future studies should consider which of these relations, in addition to ISA relations, would benefit a GIR system.

7.4.3 Evaluation of GIR Systems

In Chapter 6, the evaluation of GIR systems based on the measures of precision and recall was discussed. The existing methods of measuring precision and recall do not take into consideration the variations of perceptions of relevance between users and the changing perception of relevance as a user examines candidate datasets (Borlund, 2003a). As the purpose of geovisualisation is to make apparent hidden geographic information through visual means, it means that revised approaches are necessary for the evaluation of visualisation-supported interactive GIR systems. The revised approaches should consider the changing perception of relevance during data discovery. Other issues specific to evaluation of GIR systems, have been discussed in Chapter 6; these include the development of a geographic track for IR evaluation platforms such as CLEF. Initial efforts have already been started by researchers from the field of GIR, through the GeoCLEF initiative discussed in Chapter 6.

7.4.4 Incorporation into SDI

One of the main reasons for adopting the use of Z3950 metadata servers was to ensure that the deliverables of this study could be easily deployed within an SDI-based on OGC catalogue service standards. As was noted in Chapter 1, over 400 metadata servers listed on the FGDC clearinghouse registry adopt the Z39.50 metadata standard. This suggests that the prototypes implemented in this study already have a potential user community. Each node on the FGDC clearinghouse registry has a specific application domain, this provides a well defined test area such that it could be determined precisely which disciplines have difficulty with the proposed visualisation approaches. Such findings could be used to modify the proposed approaches to address issues specific to a particular domain.

7.5 Concluding Remarks

The development of the NEEDS geoportal was based on service-oriented geospatial infrastructure as was described in Chapters 1 and 5. This has allowed different visualisation applications to be 'attached' to the main geoportal. The successful deployment and usage of the visualisation prototypes on geoportals demonstrate the value of visual methods in spatial data discovery. Additionally they highlight the importance of web-based data dissemination in enabling those visualisation applications. This thesis therefore concludes that geoportals offering visualisation support, allow for the discovery of more relevant geographic datasets than those that do not offer such support. As research into GIR continues, it is envisaged that the NEEDS geoportal will continue to act as a platform for future development in the area of visualisation and ontology-assisted GIR.

References

Abdelguerfi, M., Wynne, C. and Cooper, E. (1998) 'Representation Of 3-D Elevation In Terrain Databases Using Hierarchical Triangulated Irregular Networks:A Comparative Analysis', *International Journal of Geographic Information Science,* 12, (8), pp. 853-873.

Abdelmoty, A. I., Smart, P. D., Jones, C. B., Fu, G. and Finch, D. (2005) 'A critical evaluation of ontology languages for geographic information retrieval on the Internet', *Journal of Visual Languages & Computing,* 16, (4), pp. 269-382.

Agarwal, P. (2005) 'Ontological considerations in GIScience', *International Journal of Geographical Information Science,* 19, (5), pp. 501-536.

AGI (2006) *GIGateway.* Available at: http://www.gigateway.co.uk (Accessed: 03/02/2006).

Alameh, N. (2003) 'Chaining Geographic Information Web Services', *IEEE Internet Computing,* 7, (5), pp. 22-29.

Albertoni, R., Bertone, A. and de Martino, M. (2005) 'Visualization and semantic analysis of geographic metadata', *In proceedings of the 2005 workshop on geographic information retrieval.* Bremen, Germany, 4th November 2005. Visualization and semantic analysis of geographic metadata: ACM Press, pp.

Andrienko, G. L. and Andrienko, N. V. (1999) 'Interactive maps for visual data exploration', *International Journal of Geographic Information Science,* 13, (4), pp. 355-374.

Android Technologies Inc. (2006) *WordNet 2.0 data files in MySQL format.* Available at: http://www.androidtech.com/html/wordnet-mysql-20.php (Accessed: 01/12/2006).

Angel, E. (2006) *Interactive Computer Graphics.* Addison-Wesley: Boston,MA.

Antoniou, G. and van Harmelen, F. (2004) 'Web Ontology Language: OWL', in Staab, S. and Studer, R.(eds) *Handbook on Ontologies.* Springer: Berlin.

Arens, C., Stoter, J. and van Oosterom, P. J. M. (2005) 'Modelling 3D spatial objects in a geo-DBMS using a 3D primitive', *Computer & Geosciences,* 31, pp. 165-177.

Arpinar, I. B., Sheth, A. and Ramakrishnan, C. (2005) 'Geospatial Ontology Development and Semantic Analytics', in J.P., W. and Fotheringham, A. S.(eds) *Handbook of Geographic Information Science.* Blackwell Publishing.

Baeza-Yates, R. (1992) 'Text retrieval : Theory and practice', *In proceedings of the 12th IFIP World Computer Congress.* Madrid, Spain, 11-17 September 1992. Text retrieval : Theory and practice: pp. 465–476.

Beard, K. and Sharma, V. (1997) 'Multidimensional ranking for data in digital spatial libraries', *International Journal on Digital Libraries,* 1, (2), pp. 153 - 160.

Beckmann, N., Kriegel , H., Schneider, R. and Seeger, B. (1990) *International Conference on Management of Data.* Atlantic City, New Jersey, United States:ACM Press.

Benbunan-Fich, R. (2001) 'Using protocol analysis to evaluate the usability of a commercial web site', *Information & Management,* 39, pp. 151-163.

Bertin, J. (1983) *Semiology of Graphics.* University of Wisconsin: Madison,WI.

Blok, C. A. (2005) *Dynamic visualization variables in animation to support monitoring of spatial phenomena.* thesis. ITC-University of Utrecht.

Borlund, P. (2003a) 'The Concept of Relevance in IR', *Journal of the American Society for Information Science and Technology,* 54, (10), pp. 913-925.

Borlund, P. (2003b) 'The IIR evaluation model: a framework for evaluation of interactive information retrieval systems', *Information Research,* 8, (3), pp. paper no. 152.

Brin, S. and Page, L. (1998) *7th International World Wide Web Conference.*

Brink, T., Gergle, D. and Wood, S. D. (2002) *Usability for the web.* Morgan Kaufmann: San Francisco.

Brown, I. (1999) 'Developing a Virtual Reality User Interface (VRUI) for Geographic Information Retrieval on the Internet', *Transactions in GIS,* 3, (3), pp. 207-220.

Bucher, B., Clough, P., Joho, H., Purves, R. and Syed, A. K. (2005) *Proceedings of the 22nd International Cartographic Conference.* A Coruna, Spain: ICA.

Butler, D. (2006) 'Virtual globes: The web-wide world', *Nature,* 439, (7078), pp. 763-890.

Carvajal, A. (2005) *9th International Conference on Information Visualisation.* London, England, 6th-8th July 2005.

Chao, L. (2006) *Database Development and Management.* Auerbach: Boca Raton, FL.

Chen, C. (2004) *Information Visualization: Beyond the Horizon.* Springer-Verlag: London.

Chen, J. X. (2002) *Guide to graphics software tools.* Springer: New York.

Cockburn, A. and McKenzie, B. (2004) 'Evaluating Spatial Memory in Two and Three Dimensions', *International Journal of Human Computer Studies,* 61, pp. 359-373.

Commission of the European Communities (2004) *Proposal for a Directive on Establishing an Infrastructure for Spatial Information in the Community (INSPIRE).* Available at: http://www.ec-gis.org/inspire/ (Accessed: 03/03/2006).

Connolly, T., Begg, C. and Strachan, A. (1999) *Database Systems: A practical Approach to Design, Implementation, and Management.* Addison-Wesley: Harlow.

Crompvoets, J., Rajabifard, A., Bregt, A. and Williamson, I. P. (2004) 'Assessing the worldwide developments of national spatial data clearinghouses', *International Journal of Geographical Information Science,* 18, (7), pp. 1-25.

CycCorp (2003) *OpenCyc: The Project.* Available at: http://www.opencyc.org (Accessed: 02/02/2006).

de Oliveira, J. L., Gonçalves, M. A. and Medeiros, C. B. (1998) 'A framework for designing and implementing the user interface of a geographic digital library', *International Journal on Digital Libraries,* 2, (2-3), pp. 190 - 206.

de Oliveira, M. C. F. and Levkowit, H. (2003) 'From Visual Data Exploration to Visual Data Mining: A Survey', *IEEE Transactions On Visualization And Computer Graphics,* 9, (3), pp. 378-394.

Dix, A., Finlay, J., Abowd, G. and Beale, R. (2004) *Human-computer interaction.* Pearson/Prentice-Hall: Harlow, England.

Döllner, J. (2005) 'Geovisualization and Real-Time 3D Graphics', in Dykes, J., Maceachren, A. and Kraak, M.(eds) *Exploring Geovisualization.* Elseviers: Amsterdam.

Döllner, J. and Buchholz, H. (2005) *Proceedings of the 13th ACM International Workshop on Geographic Information Systems (ACM-GIS 2005).* Bremen, Germany, November 4-5, 2005.

download.com (2005) *Download Statistics.* Available at: http://www.download.com (Accessed: 13/03/2006).

Egenhofer, M. (2002) *Proceedings of the 10th ACM International Symposium on Advances in Geographic Information Systems.* McLean, Virginia: ACM Press.

ESRI (2004) *The GIS Portal Toolkit.* Available at: http://www.esri.com/software/arcgis/gisportal-toolkit/about/case-studies.html (Accessed: 02/02/2006).

ESRI. (2005a) *Raster Data in ArcSDE 9.1*

Fabrikant, S. I. (2000) *Spatial Metaphors for Browsing Large Data Archives.* thesis. University of Colorado.

Fabrikant, S. I. and Buttenfield, B. P. (2001) 'Formalizing semantic spaces for information access', *Annals of the Association of American Geographers,* 91, pp. 263-280.

Fabrikant, S. I. and Skupin, A. (2005) 'Cognitively Plausible Information Visualization', in Dykes, J., Maceachren, A. and Kraak, M.(eds) *Exploring Geovisualization.* Elsevier: Amsterdam.

FDGC (2005) *US Federal Geographic Data Committee.* Available at: http://www.fgdc.gov (Accessed: 03/02/2006).

Felkel, P. and Obdrzalek, S. (1998) *Proceedings of the 14th Spring Conference on Computer Graphics.* Budmerice, Slovakia, April 1998.

Fellbaum, C. (1998) *WordNet: An Electronic Lexical Database.* MIT Press: Cambrigde, MA.

FGDC (2004) *FGDC-ISO Crosswalk.* Available at: http://www.fgdc.gov/CrossWalk/ISO-FGDC-METADATA-CROSSWALK-V4.xls (Accessed:

Gaede, V. and Günther, O. (1998) 'Multidimensional access methods', *ACM Computing Surveys,* 30, (2), pp. 170 - 231.

Gayle, S. and Rainer, A. (2005) *Proceedings of the 11th IEEE International Software Metrics Symposium (METRICS 2005).* Como, Italy, 19th-22th September 2005.

GeoVRML WG (2002) *GeoVRML.* Available at: http://www.geovrml.org/ (Accessed: 23/01/2006).

Gey, F., Larson, R., Sanderson, M., Joho, H., Clough, P. and Petras, V. (2005) *GeoCLEF workshop.* Vienna, Austria, 21st-23rd September 2005.

Gilbert, E. W. (1958) 'Pioneer Maps of Health and Disease in England', *The Geographical Journal,* 124, (2), pp. 172-183.

Gilchrist, A. (2003) 'Thesauri, taxonomies and ontologies–an etymological note', *Journal of Documentation,* 59, (1), pp. 7-18.

Göbel, S., Haist, J., Reiterer, H. and Müller, F. (2002) 'INVISIP: Usage of Information Visualization Techniques to Access Geospatial Data Archives', *Lecture Notes In Computer Science,* 2453, pp. 371-380.

Google (2005) *Google Earth.* Available at: http://earth.google.com (Accessed: 17/05/2006).

Gore, A. (1998) *The Digital Earth : Understanding our Planet in the 21st Century.* Available at: http://www.digitalearth.gov/VP19980131.html (Accessed: 001/02/2006).

Güting , R. H. (1994) 'An Introduction to Spatial Database Systems', *VLDB Journal,* 3, (4), pp. 357-399.

Habing, T. (2002-2003) *VBZOOMC.* Available at: http://vb-zoom.sourceforge.net/ (Accessed: 13/05/2006).

Harter, S. (1996) 'Variations in relevance assessments and the measurement of retrieval effectiveness', *Journal of the American Society for Information Science and Technology,* 47, (1), pp. 37-49.

Hearst, M. (1999) 'User Interfaces and Visualization', in Baeza-Yates, R. and Ribeiro-Neto, B.(eds) *Modern Information Retrieval. Addison-Wesley-Longman Publishing co: Boston.* ACM Press/Addison-Wesley: New York.

Hodge, G. (2001) *Metadata Made Simpler.* NISO Press: Bethesda.

Horstmann, C. S. and Cornell, G. (2001) *Core Java.* Prentice Hall: Upper Sadle RIver, NJ.

Horstmann, C. S. and Cornell, G. (2004) *Core Java 2, Volume II: Advanced Features.* Prentice Hall.

HSQLDB Development Group (2005) *HSQLDB.* Available at: http://hsqldb.org/ (Accessed: 13/05/2006).

Huang, B. (2003) 'Web-based dynamic and interactive environmental visualization.' *Computers, Environment and Urban Systems,* 27, (6), pp. 62-636.

Hübner, S., Spittel, R., Visser, U. and Vögele, T. J. (2004) 'Ontology-based Search for Interactive Digital Maps', *IEEE Intelligent Systems,* 19, (3), pp. 80- 86.

Index Data (2006) *Yaz Toolkit.* Available at: http://www.indexdata.com/yaz/ (Accessed: 13/05/2006).

ISO (2003) *ISO19115:2003 Geographic information -- Metadata (International Standards Organization).* Available at: http://www.iso.org (Accessed: 03/03/2006).

ISO (2004) *ISO 19125:2004 Geographic information -- Simple feature access (International Standards Organisation).* Available at: http://www.iso.org (Accessed: 03/03/2006).

Jones, C. B., Abdelmoty, A. I., Finch, D., Fu, G. and Vaid, S. (2004) *Proceedings of the Third International Conference on Geographic Information Science (GIScience).* Maryland, USA:Springer.

Jones, C. B., Alani, H. and Tudhope, D. (2001) 'Geographical Information Retrieval with Ontologies of Place', in Montello, D. R.(ed), *Lecture Notes in Computer Science.* Vol. 2205 Springer: Berlin, pp. 323-335.

Jones, C. B., Purves, R., Ruas, A., Sanderson, M., Sester, M., van Kreveld, M. and Weibel, R. (2002) *Proceedings of the 25th annual international ACM SIGIR conference on Research and development in information retrieval.* ACM Press.

Kavouras, M., Kokla, M. and Tomai, E. (2005) 'Comparing Categories among Geographic Ontologies', *Computers & Geosciences,* 31, pp. 145-154.

Kersting, O. and Döllner, J. (2002) *Proceedings of the tenth ACM international symposium on Advances in geographic information systems.* McLean, Virginia, Klien, E., Lutz, M. and Kuhn, W. (2006) 'Ontology-based Discovery of Geographic Information Services - An Application in Disaster Management', *Computers, Environment and Urban Systems,* 30, (1), pp. 102-123.

Kobayashi, M. and Takeda, K. (2000) 'Information retrieval on the web', *ACM Computing Surveys,* 32, (2).

Kolas, D., Hebeler, J. and Dean, M. (2005) *Proceedings of the First International Conference on GeoSpatial Semantics (GeoS 2005).* Mexico City, Mexico:Springer.

Kraak, M. J. and Ormeling, F. (2003) *Cartography: Visulisation of Geospatial Data.* Pearson Education Limited: Essex.

Lake, R., Burggraf, D., Trninic, M. and Rae, L. (2004) *Geography Mark-Up Language: Foundation for the Geo-Web.* John Wiley & Sons.

Larson, R. and Frontiera, P. (2004) *Workshop on Geographic Information Retrieval SIGIR 2004.* Sheffield UK,

Larson, R. R. (1996) 'Geographic Information Retrieval and Spatial Browsing', in Smith, L. and Gluck, M.(eds) *GIS and Libraries: Patrons, Maps and Spatial Information.* University of Illinois.

Laycock, R. G. and Day, A. M. (2003) 'Automatically Generating Roof Models from Building Footprints', *Journal of WSCG,* 11, (1-3).

Lenat, D., Milller, G. and Yokoi, T. (1995) 'CYC, WordNet and EDR: Critiques and Responses', *Communications of the ACM,* 38, (11), pp. 45-48.

Leuski, A. and Allan, J. (2004) 'Interactive Information Retrieval Using Clustering And Spatial Proximity', *User Modeling And User Adapted Interaction,* 14, (2-3), pp. 259 - 288.

Lewis, C. (1982) *Using the 'thinking-aloud' method in cognitive interface design.* Research Report RC-9265: IBM T.J. Watson Research Center,Yorktown Heights, NY, USA

Lin, H. X., Choong, Y. Y. and Salvendy, G. (1997) 'A Proposed Index of Usability: A Method for Comparing the Relative Usability of Different Software Systems', *Behaviour & Information Technology,* 16, (4/5), pp. 267-278.

Longley, P. A., Goodchild, M. F. and Maguire, D. J. (2001) *Geographic Information Systems and Science.* Wiley: West Sussex England.

Lutz, M. and Klien, E. (2006) 'Ontology-based retrieval of geographic information', *International Journal of Geographic Information Science,* 20, (3), pp. 233-260.

Lynch, P., Luan, X., Prettyman, M., Mericle, L., Borkmann, E. and Schlaifer, J. (2004) *Proceedings of the International Conference on Information Technology: Coding and Computing.*

MacEachren, A. M. and Kraak, M. J. (2001) 'Research Challenges in Geovisualization', *Cartography and Geographic Information Systems,* 28, (1), pp. 3-12.

Maguire, D. J. and Longley, P. (2004) 'The emergence of geoportals and their role in spatial
data infrastructures', *Computers, Environment and Urban Systems,* 29, (1), pp. 3-14.

Maguire, D. J. and Longley, P. (2005) 'The emergence of geoportals and their role in spatial
data infrastructures', *Computers, Environment and Urban Systems,* 29, (1), pp. 3-14.

Martins, B., Silva, M. J. and Chaves, M. S. (2005) *Proceedings of the 2005 workshop on Geographic Information Retrieval.* Bremen, Germany:ACM Press.

Mihalcea, R. (2003) 'Turning WordNet into an Information Retrieval Resource: Systematic Polysemy and Conversion to Hierarchical Codes', *International Journal of Pattern Recognition & Artificial Intelligence,* 17, (5), pp. 689-704.

Miller, G. (1995) 'WordNet: a lexical database for English', *Communications of the ACM,* 38, (11), pp. 39 - 41.

Moore, K., Dykes, J. and Wood, J. (1999) 'Using Java to interact with geo-referenced VRML within a Virtual Field Course', *Computers and Geosciences,* 25, (10), pp. 1125-1136.

Nielsen, J. (1994) 'Estimating the number of subjects needed for a thinking aloud test', *International Journal of Human-Computer Studies,* 41, (3), pp. 385-397.

Nielsen, J. (2000) *Designing Web Usability.* New Riders Publishing.

Nogueras-Iso, J., Zarazaga-Soria, F. J., Bejar, R. and Alvarez, P. J. (2005) 'OGC Catalog Services: a key element for the development of Spatial Data Infrastructures', *Computer & Geosciences,* 31, pp. 199-209.

Nogueras-Iso, J., Zarazaga-Soria, F. J., Lacasta, J., Béjar, R. and Muro-Medrano, P. R. (2004) 'Metadata standard interoperability: application in the geographic information domain', *Computers, Environment and Urban Systems,* 28, (6), pp. 611-634.

OGC (2001a) *OpenGIS Service Architecture-ISO19119 (Open Geospatial Consortium).* Available at: http://www.opengeospatial.org (Accessed: 02/04/2006).

OGC (2001b) *Grid Coverage (Open Geospatial Consortium).* Available at: http://www.opengeospatial.org (Accessed: 02/04/2006).

OGC (2004) *Geography Markup Language (GML) Encoding Specification (Open Geospatial Consortium).* Available at: (Accessed: 03/03/2006).

Open Geospatial Consortium (2004) *Geography Markup Language (GML) Encoding Specification.* Available at: (Accessed:

Open Geospatial Consortium (2005) *Simple Features - SQL.* Available at: http://www.opengeospatial.org (Accessed: 24/01/2006).

Owen, S., Brereton, P. and Budgen, D. (2006) 'Protocol analysis: a neglected practice', *Communications of the ACM,* 49, (2), pp. 117-123.

Papadias, D. and Theodoridis, Y. (1997) 'Spatial relations, minimum bounding rectangles, and spatial data structures', *International Journal of Geographical Information Science,* 11, (2), pp. 111 - 138.

Podolak, I. and Demšar, U. (2004) *Proceedings of the 12th International Conference on Geoinformatics.* Sweden: University of Gävle.

PostgreSQL Global Development Group (2005) *PostgreSQL Documentation.* Available at: http://www.postgresql.org (Accessed: 24/01/2006).

Proctor, S. P., Gopal, S., Imai, A., Wolfe, J., Ozonoff, D. and White, R. F. (2005) 'Spatial Analysis of 1991 Gulf War Troop Locations in Relationship with Postwar Health Symptom Reports Using GIS Techniques', *Transactions in GIS,* 9, (3), pp. 381-396.

Purves, R., Syed, A. K., Yang, B. and Weibel, R. (2005) *Proceedings of the 13th International Cartographic Conference.* La Coruna, Spain, 9th-16th July 2005.ICA.

Rainio, P. (2005) *Visualization Of Geospatial Metadata For Selecting Geographic Datasets.* thesis. Helsinki University of Technology.

Raper, J., Dykes, J., Wood, J., Mountain, D., Krause, A. and Rhind, D. (2002) 'A framework for evaluating geographical information', *Journal of Information Science,* 28, (1), pp. 39–50.

Raskin, R. G. and Pan, M. J. (2005) 'Knowledge representation in the semantic web for Earth and environmental terminology (SWEET)', *Computer & Geosciences,* 31, (9), pp. 1119-1125.

Resnik, P. (1999) 'Semantic Similarity in a Taxonomy: An Information-Based Measure and its Application to Problems of Ambiguity in Natural Language', *Journal of Artificial Intelligence Research,* 11, pp. 95-130.

Rhyne, T. (1999) ' A commentary on GeoVRML: a tool for 3D representation of georeferenced data on the web', *International Journal of Geographic Information Science,* 13, (4), pp. 439-443.

Richardson, R. and Smeaton, A. (1995) *Using WordNet in a Knowledge-Based Approach to Information Retrieval.* School of Computer Applications, Trinity College Dublin, Working paper CA-0395, pp: 16.

Ripley, B. D. (2001) 'The R project in statistical computing', MSOR Connections, 1, 1, p.23-25.

Risden, K., Czerwinki, M., Munzer, T. and Cook, D. (2000) 'An initial examination of ease of use for 2D and 3D information visualizations of web content', *International Journal of Human Computer Studies,* 53, pp. 695-714.

Robinson, A. H., Morrison, J. L., Muehrcke, P. C., Kimerling, A. J. and Guptill, S. C. (1995) *Elements of Cartography.* John Wiley & Sons: New York.

Rodriguez, A. and Egenhofer, M. (2004) 'Comparing Geospatial Entity Classes: An Asymmetric and Context-Dependent Similarity Measure', *International Journal of Geographical Information Science,* 18, (3), pp. 229-256.

Rohlf, J. and Helman, J. (1994) *Proceedings of the 21st annual conference on Computer graphics and interactive techniques.* ACM Press.

Rohrer, R. M. and Swing, E. (1997) 'Web-Based Information Visualization', *IEEE Computer Graphics and Applications,* 17, (4), pp. 52-29.

Rose, L. (2004) *Geospatial Portal Reference Architecture.* Available at: http://www.opengeospatial.org/specs/?page=baseline (Accessed: 02/02/2006).

Selman, D. (2002) *Java 3D Programming.* Manning Publishers: Greenwich, United States.

Seron, F. J., Rodriguez, R., Cerezo, E. and Pina, A. (2002) 'Adding Support for High-Level Skeletal Animation', *IEEE Transactions On Visualization And Computer Graphics,* 8, (4), pp. 360 - 372.

Siegel, N., Goolsbey, K., Kahlert, R. and Matthews, G. (2004) *The Cyc® System: Notes on Architecture.* Available at: http://www.cyc.com/cyc/technology/whitepapers_dir/Cyc_Architecture_and_API.pdf (Accessed: 17/07/2006).

Sivakuma, R., Rao, M. L., DeLoatch, L., Garie, H. and Sideralis, K. (2004) *Proceedings of the 7th International Global Spatial Data Infrastructure Conference.* Bangalore, India, 2nd-4th February 2004.

Sjölinder, M., Höök, K., Nilsson, L.-G. and Andersson, G. (2005) 'Age differences and the acquisition of spatial knowledge in a three-dimensional environment: Evaluating the use of an overview map as a navigation aid', *International Journal of Human Computer Studies,* 63, (6), pp. 537-564.

Soanes, C. and Hawkes, S. (2005) *Compact Oxford English Dictionary of Current English.* Oxford University Press: Oxford.

Sutcliffe, A. G., Ennis, M. and Hu, J. (2000) 'Evaluating the effectiveness of visual user interfaces for information retrieval', *International Journal of Human Computer Studies,* 53, (5), pp. 741-763.

Tait, M. G. (2005) 'Implementing geoportals: applications of distributed GIS', *Computers, Environment and Urban Systems,* 29, (17), pp. 33-47.

Tang, W. and Selwood, J. (2005) *Spatial Portals: Gateways to Geographic Information.* ESRI Press: Redlands,CA.

Tversky, A. (1977) 'Features of similarity', *Psychological Review,* 84, pp. 327-357.

van Kreveld, M., Reinbacher, I., Arampatzis, A. and van Zwol, R. (2005) ' Multi-Dimensional Scattered Ranking Methods for Geographic Information Retrieval', *Geoinformatica,* 9, (1), pp. 61-84.

van Someren, M. V., Barnard, Y. F. and Sandberg, J. A. C. (1994) *The Think Aloud Method:A practical guide to modelling cognitive processes.* Academic Press: London.

Varelas, G., Voutsakis, E., Raftopoulou, P., Petrakis, E. G. M. and Milios, E. E. (2005) *Proceedings of the 7th annual ACM international workshop on Web information and data management.*

W3C (1999-2006) *World Wide Web Consortium Specifications.* Available at: http://www.w3.org (Accessed: 02/02/2006).

Wang, Y. D. and Forgionne, G. (2006) 'A decision-theoretic approach to the evaluation of information retrieval systems', *Information Processing and Management,* 42, (4), pp. 863-874.

Ware, C. and Frank , G. (1996) 'Evaluating Stereo and Motion Cues For Visualizing Information Nets in Three Dimensions', *ACM Transactions on Graphics,* 15, (2), pp. 121-140.

Ware, C. and Plumlee, M. (2005) '3D Geovisualization and the Structure of Visual Space', in Dykes, J., Maceachren, A. and Kraak, M.(eds) *Exploring Geovisualization.* Elsevier: Amsterdam.

Watt, A. and Policarpo, F. (2003) *3D Games : Animation and Real-time Rendering.* Addison-Wesley: London.

Web3D Consortium (2004) *X3D.* Available at: http://www.web3d.org (Accessed: 05/08/2004).

Weerawarana, S., Curbera, F., Leymann, F., Ferguson, D. F. and Storey, T. (2005) *Web Services Platform Architecture : SOAP, WSDL, WS-Policy, WS-Addressing, WS-BPEL, WS-Reliable Messaging, and More.* Prentice Hall.

West, L. and Schofield, L. (2004) *GigatewayTM 2004-2005: A Year In Review.* London: Association of Geographic Information (AGI)

Worboys, M. and Duckham, M. (2004) *GIS: A computing perspective.* CRC Press: Boca Ranton, Fl.

World Wide Web Consortium (2004) *OWL Web Ontology Langauge Overview.* Available at: http://www.w3.org/TR/2004/REC-owl-features-20040210/ (Accessed: 02/02/2006).

Yamuna, P. and Candan, K. S. (2000) *Proceedings of the 2000 ACM workshop on Multimedia.* Los Angeles,

Appendix A Think-Aloud Evaluation Handout

Scenario

The North East of England is undergoing a major regeneration exercise. As part of the exercise, information relating to the possible impact on the environment is required. In this exercise, your role is to acquire datasets that could help in the determination of the environmental impact of the regeneration exercise. It is envisaged that the datasets will help regional planners arrive at well-informed decisions regarding were to locate various infrastructure.

A colleague has recommended that you search the North East Environmental Data Server (NEEDS). It holds geographic data and associated metadata covering areas of the North East of England.

Data Collection

Datasets stored in the NEEDS Central Archive include large scale vector mapping of some buildings, roads, archaeological sites and other infrastructure. The Central Archive also holds some aerial imagery and photography covering parts of the North East of England. Those datasets that are stored in the Central Archive are downloadable; however those that are accessed from external archives are not.

Experimental Setup

For each of the tasks that follow please say your thoughts out load as you use the NEEDS system (that is, verbalise everything that comes to mind).

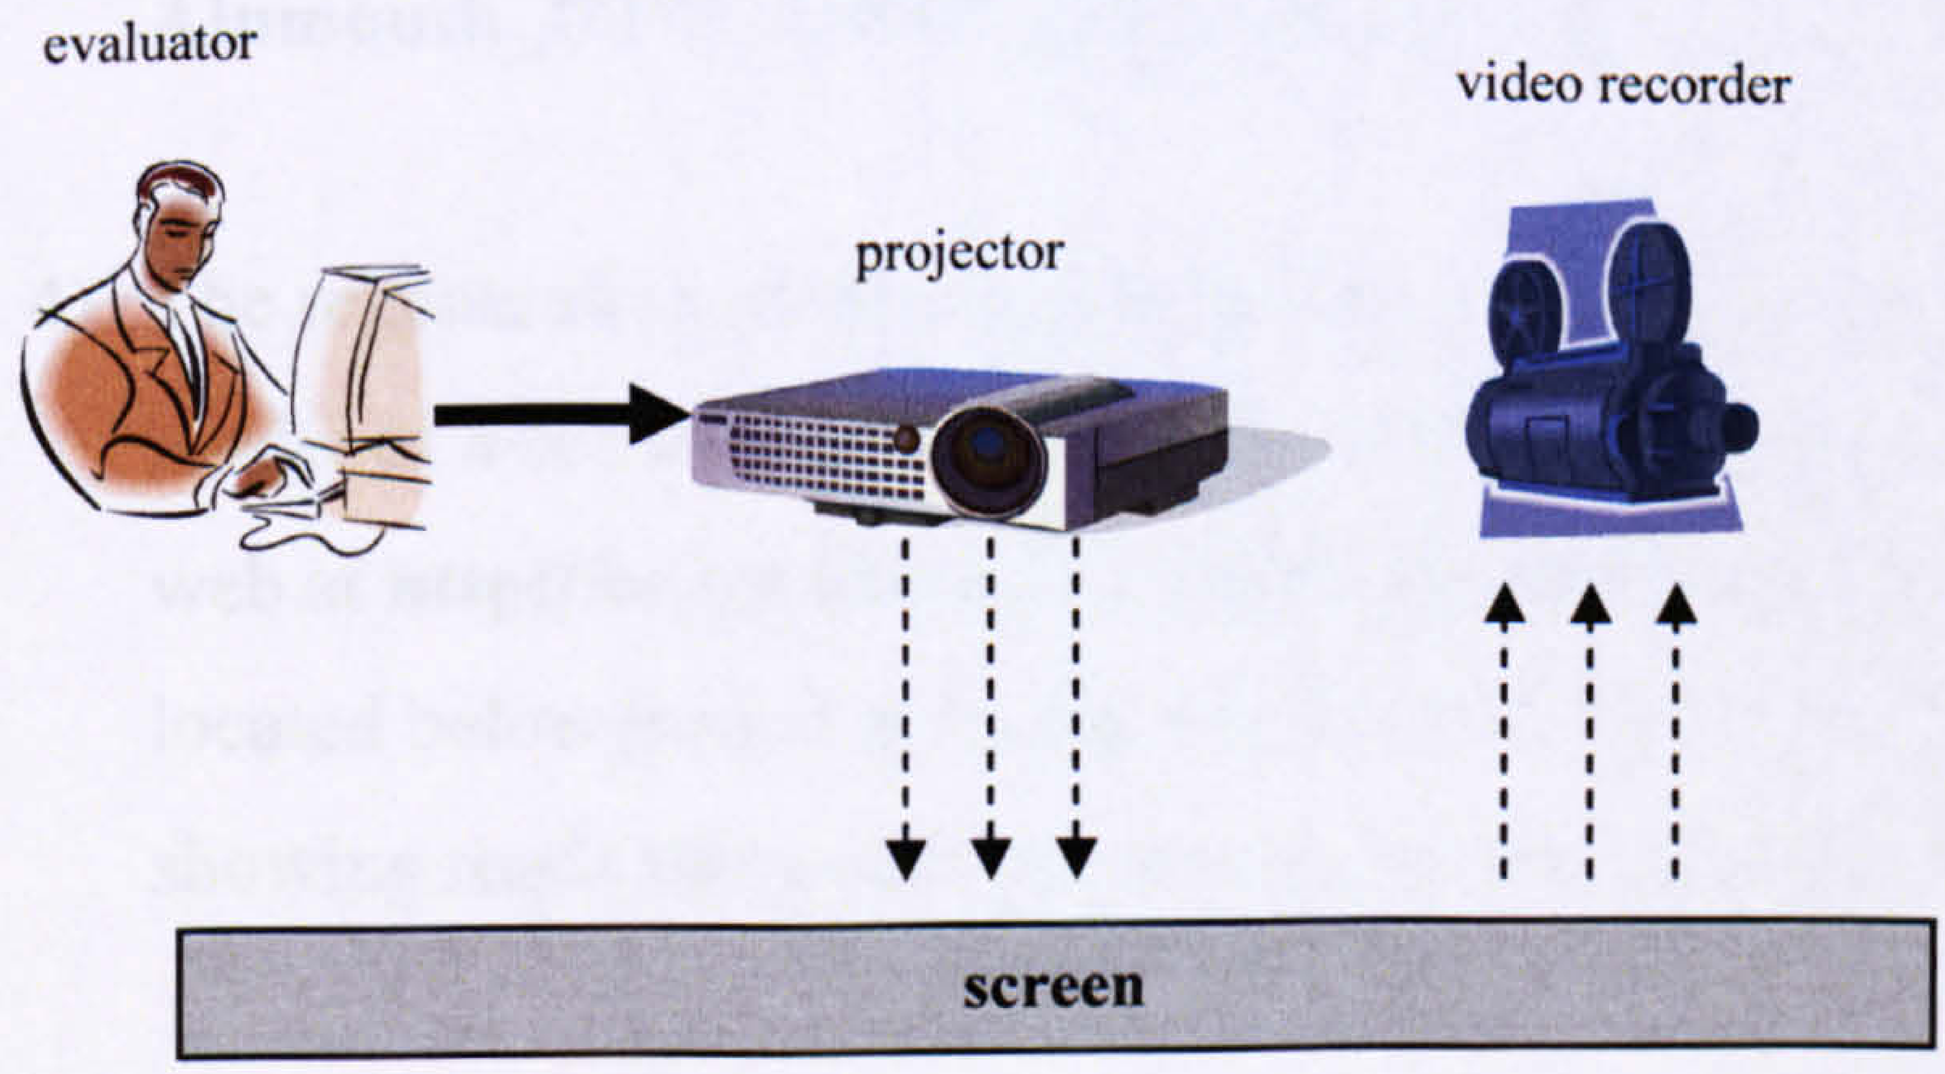

Tasks

The search terms for each task are in blue.

Additional datasets in GeoDOVE are in **black and bold.**

For tasks 1 and 2, search using the STORM browser.

1) Of major concern to regional planners is the impact the regeneration will have on current agricultural activity in the region. You have been asked to identify datasets containing agricultural information from the NEEDS archive, in the area bounded by N,S,W,E, = 56.609, 53.3, -4.146,-0.17697. For the period (02/01/2006—02/07/2006)

2) The National Heritage Trust is also concerned about the impact the regeneration will have on studies of fossils in the area bounded by N,S,W,E, = 59.32, 50.711, -7.642,2.201. Identify possibly relevant datasets. For the period (02/03/2003—02/05/2006)

For tasks 3 and 4, use the Fast Search tool on the NEEDS home page. Once the list of datasets has been found, use GeoDOVE to tackle the rest of the task.

3) Some of the regeneration could affect the flood defences of the town of Alnmouth. You are to identify vector datasets that show buildings of Alnmouth in danger of flooding. A Digital Elevation Model (DEM) of Alnmouth is available in the GeoDOVE Database Server and has been named **Alnmouth_DTM_21SW_grid.** Flood plains are coloured deep red.

4) The regeneration project may affect an ongoing geological study. A 3D model of a subsurface geological structure has been made available on the web at **http://ce-gw114.ncl.ac.uk/data/vrml/rock2.wrl** The structure is located belowground at coordinate 424560, 564620,-100. Find vector datasets showing roads that cover the area above the subsurface structure.

Questionnaire

After carrying out these tasks, please fill in the questionnaire at http://www.needs.ncl.ac.uk/needs2005/questionnaire.jsp

There are approximately 100 questions listed but they are not all applicable, so please feel free to ignore those questions that you deem not applicable.

Thank you for participating,

Gobe Hobona

3D Navigation

zoom in, zoom out - centre mouse button

rotate up or down – left mouse button

translate left, right, up or down – right mouse button

Change to multicoloured DEM

Double click on the raster layer's name in the Table of Contents, then click *Legend Symbol*

Adding a DEM

Click the *Add Data* button

Adding a 3D model

Click the *Add 3D Model* button

Appendix B NEEDS Usability Questionnaire

Adapted from Pelman, G.,Web-Based User Interface Evaluation with Questionnaires,available at http://www.acm.org/~perlman/question.cgi?form=PUTQ which is based on Lin, H.X. Choong, Y.-Y., and Salvendy, G. (1997) A Proposed Index of Usability: A Method for Comparing the Relative Usability of Different Software Systems. Behaviour & Information Technology, 16:4/5, pg267-278
Top of Form

1. COMPATIBILITY

1. Is the control of cursor compatible with movement?
2. Are the results of control entry compatible with user expectations?
3. Is the control matched to user skill?
4. Are the coding compatible with familiar conventions?
5. Is the wording familiar?

2. CONSISTENCY

6. Is the assignment of colour codes conventional?
7. Is the coding consistent across displays, menu options?
8. Is the cursor placement consistent?
9. Is the display format consistent?
10. Is the feedback consistent?
11. Is the format within data fields consistent?
12. Is the label format consistent?
13. Is the label location consistent?
14. Is the labelling itself consistent?
15. Is the display orientation consistent? -- panning vs. scrolling.
16. Are the user actions required consistent?
17. Is the wording consistent across displays?
18. Is the data display consistent with entry requirements?
19. Is the data display consistent with user conventions?
20. Are symbols for graphic data standard?
21. Is the option wording consistent with command language?

22. Is the wording consistent with user guidance?

3. FLEXIBILITY

23. Does it have by-passing menu selection with command entry?
24. Does it have direct manipulation capability?
25. Is the design for data entry flexible?
26. Can the display be controlled by user flexibly?
27. Does it provide flexible sequence control?
28. Does it provide flexible user guidance?
29. Are the menu options dependent on context?
30. Can user name displays and elements according to their needs?
31. Does it provide good training for different users?
32. Are users allowed to customize windows?
33. Can users assign command names?
34. Does it provide user selection of data for display?
35. Does it handle user-specified windows?
36. Does it provide zooming for display expansion?

4. LEARNABILITY

37. Does it provide clarity of wording?
38. Is the data grouping reasonable for easy learning?
39. Is the command language layered?
40. Is the grouping of menu options logical?
41. Is the ordering of menu options logical?
42. Are the command names meaningful?
43. Does it provide no-penalty learning?

5. MINIMAL ACTION

44. Does it provide combined entry of related data?
45. Will the required data be entered only once?
46. Does it provide default values?
47. Is the shifting among windows easy?
48. Does it provide function keys for frequent control entries?

49. Does it provide global search and replace capability?
50. Is the menu selection by pointing? -- primary means of sequence control.
51. Is the menu selection by keyed entry? -- secondary means of control entry.
52. Does it require minimal cursor positioning?
53. Does it require minimal steps in sequential menu selection?
54. Does it require minimal user control actions?
55. Is the return to higher-level menus required only one simple key action?
56. Is the return to general menu required only one simple key action?

6. MINIMAL MEMORY LOAD

57. How are abbreviations and acronyms used?
58. Does it provide aids for entering hierarchic data?
59. Is the guidance information always available?
60. Does it provide hierarchic menus for sequential selection?
61. Are selected data highlighted?
62. Does it provide index of commands?
63. Does it provide index of data?
64. Does it indicate current position in menu structure?
65. Are data items kept short?
66. Are the letter codes for menu selection designed carefully?
67. Are long data items partitioned?
68. Are prior answers recapitulated?
69. Are upper and lower case equivalent?
70. Does it use short codes rather than long ones?
71. Does it provide supplementary verbal labels for icons?

7. PERCEPTUAL LIMITATION

72. Does it provide coding by data category?
73. Is the abbreviation distinctive?
74. Is the cursor distinctive?
75. Are display elements distinctive?
76. Is the format for user guidance distinctive?

77. Do the commands have distinctive meanings?
78. Is the spelling distinctive for commands?
79. Does it provide easily distinguished colours?
80. Is the active window indicated?
81. Are items paired for direct comparison?
82. Is the number of spoken messages limited?
83. Does it provide lists for related items?
84. Are menus distinct from other displayed information?
85. Is the colour coding redundant?
86. Does it provide visually distinctive data fields?
87. Are groups of information demarcated?
88. Is the screen density reasonable?

8. USER GUIDANCE

89. System feedback: How helpful is the error message?
90. Does it provide CANCEL option?
91. Are erroneous entries displayed?
92. Does it provide explicit entry of corrections?
93. Does it provide feedback for control entries?
94. Is HELP provided?
95. Is completion of processing indicated?
96.Are repeated errors indicated?
97.Are error messages non-disruptive/informative?
98. Does it provide RESTART option?
99.Does it provide UNDO to reverse control actions?
100. Is the sequence control user initiated?

Appendix C STORM Evaluation Handout

Thank you for volunteering to participate in a user evaluation of the STORM (Spatial Temporal Ontological Relevance Model) browser. The STORM browser computes the strength of the relationship between a query, entered by the user, and the information attached to each spatial dataset returned by the query. The strength of the relationship is an indication of the relevance of the dataset to the user's task. The relevance is then mapped onto a three dimensional graph according to the degree of similarity in space (spatial), time (temporal) and theme (semantic) between the query and the dataset (shown in figure 1). A three dimensional ranked list is provided by the STORM browser to assist in interpretation of results. Both these components, the 3D graph and the 3D ranked list, make up the STORM browser.

Please launch the STORM browser and kindly test some of the queries in table 1. To enter parameters for each query, please click on the button circled in figure 1. Then after entering the query parameters, click on the button with the magnifying glass to run the query. Please fill in the number of retrieved datasets for those queries you tested in table 1 and kindly evaluate the system by filling in the questionnaire in table 2.

The browser requires Java version 1.5 and may be accessed through:

`http://ce-gw114.ncl.ac.uk/needz2/storm.jnlp`

You will be comparing the STORM browser (both the 3D graph and accompanying ranked list) to the traditional system offered by **GIgateway** at

`http://www.gigateway.co.uk/datalocator/default.html`

Please repeat the same queries on GIgateway, excluding the UNEP.NET ones as these cannot be accessed through GIgateway.

Things to notice:

- STORM's metadata panel highlights both the search term and terms related to that search term eg. **Agricultural** is related to **Agriculture, farming**, etc
- Thumbnails of similar documents appear clustered together in space
- The most relevant document is the one closest to coordinate (1,1,1)
- The least relevant is the one closest to (0,0,0)

Your participation in this evaluation is highly appreciated.

Kind regards,

Gobe Hobona
Research Student
School of Civil Engineering and Geosciences
University of Newcastle upon Tyne

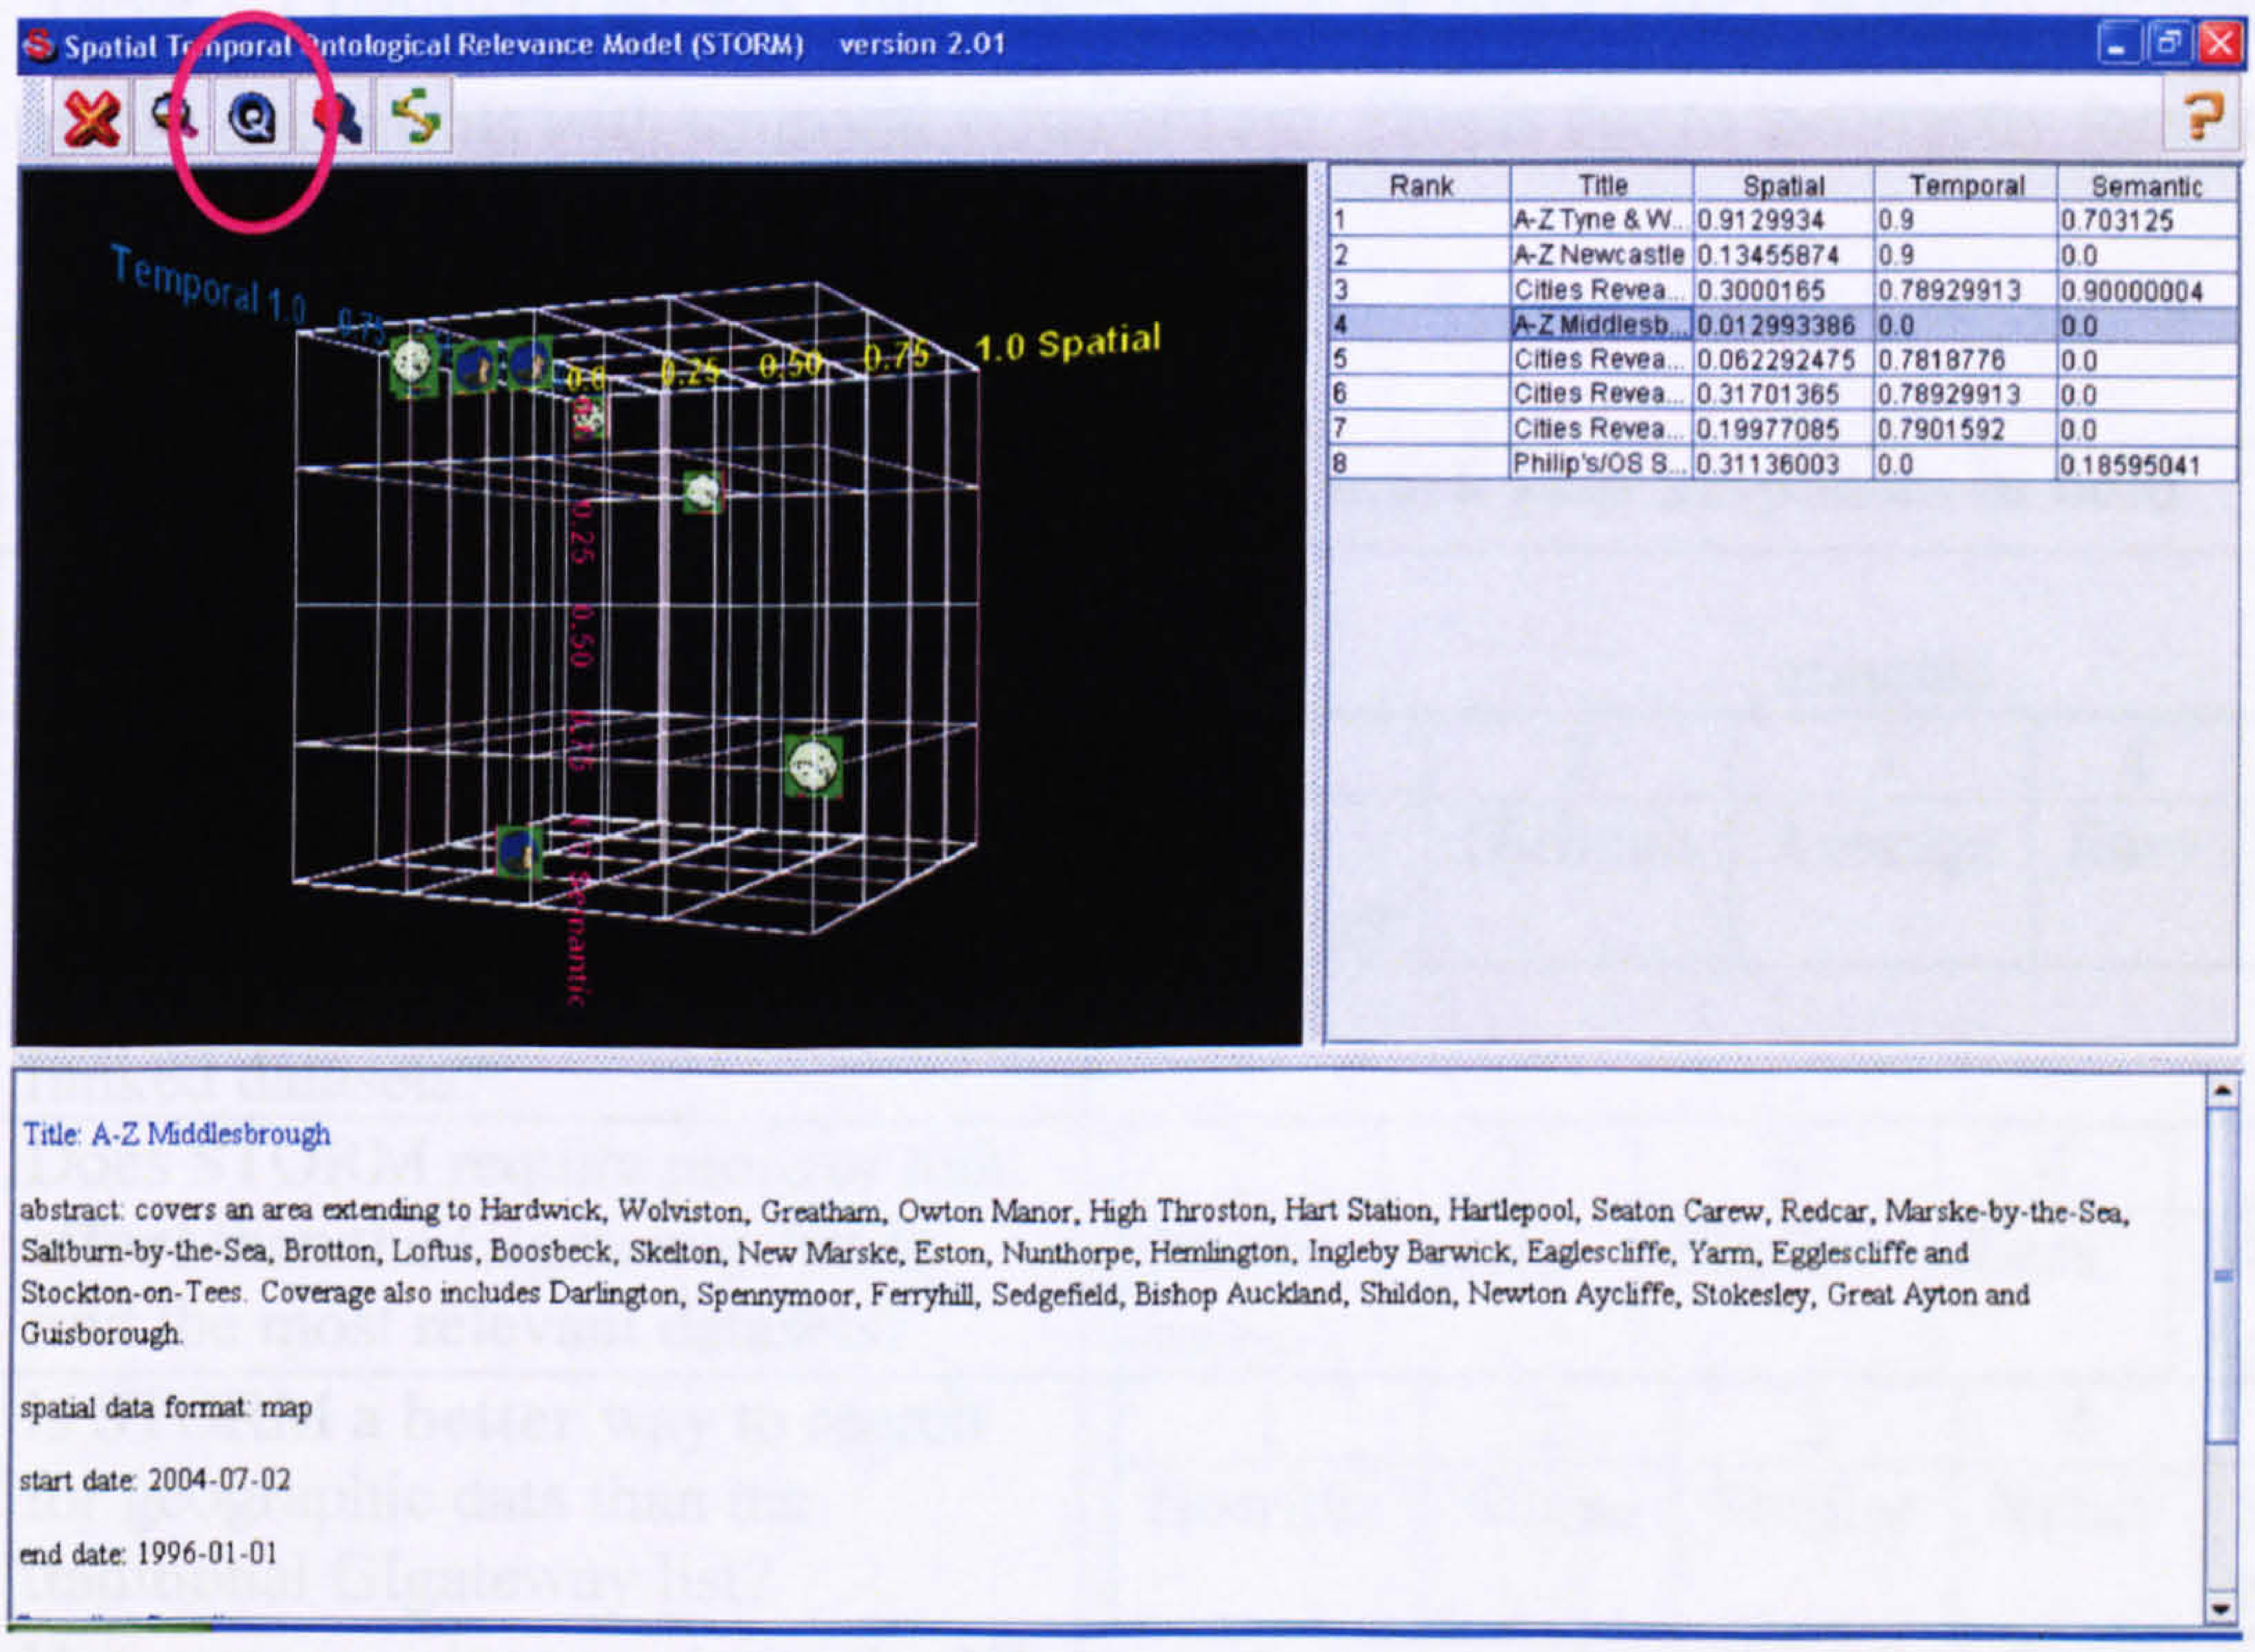

Figure 1. The STORM browser interface

Keyword	Spatial footprint	Collection	Number of retrieved datasets
borehole	north=54.955 south=53.024 west=-3.244 east=-0.87	British Geological Survey (BGS)	
agricultural	north=55.351 south=52.958 west= -3.354 east=-0.259	UNEP.Net	
soil	north=55 south=53 west=-3 east=-1	GIGateway	
census	north=56.624 south=54.67 west=-3.178 east=-1.51	Central Government (IGGI)	
seismic	north=55.592 south=52.628 west=-3.947 east=-0.895	British Geological Survey(BGS)	
erosion	north=54 south=51 west=-2 east=1	UNEP.Net	
macrofossils microfossils	north=55.329 south=52 .65 west=-3.508 east=0.597	British Geological Survey(BGS)	

Table 1. Controlled queries (all letters should be lower case). Some of the searches return documents with temporal value of zero. This is due to incorrectly formatted metadata.

Questions	Please mark your responses in bold				
How long did it take to **learn** how to use the system in minutes?	minutes				
How easy was it to **identify** the top 20 ranked datasets? i.e. was the system's organisation of thumbnails convenient for finding the top 20 ranked datasets?	1 Very Difficult	2 Difficult	3 Average	4 Easy	5 Very Easy
Does STORM require more or less **effort** than the GIgateway list to find the most relevant datasets?	1 Much more	2 More	3 Similar	4 Less	5 Much Less
Is STORM a **better** way to search for geographic data than the traditional GIgateway list?	1 Horrible	2 Worse	3 Similar	4 Better	5 Much better
How easy was it to **navigate** the 3D graph?	1 Very Difficult	2 Difficult	3 Average	4 Easy	5 Very Easy
Is STORM a more **attractive** and **interesting** way for finding	1 Much	2 Less	3 Similar	4 More	5 Much

geographic data on the web than the GIgateway list?	less | | | | more
How often do you play computer games?	1 Never | 2 Sometimes | 3 Always
Any other comments or suggestions?	

Table 2. Questionnaire

Appendix D Think Aloud Evaluation Observations

The following are observations made during the think aloud evaluation of the NEEDS geoportal.

Participant 1

When launching the STORM browser from the geoportal home page, the participant voiced his concern regarding the Java Security manager pop-up window. He appeared to prefer interactive selection of the query extent as opposed to directly typing into the textboxes. This participant exhibited the best navigation capabilities amongst all participants. He was able to comfortably 'fly through' the visualisation using both the keyboard and the mouse. He even discovered 'hidden functions' that had not been shown to him by the researcher. He made several references to the metadata panel and the text presented on the panel, stating the usefulness of the descriptions of data. Regarding the presentation metaphor, he complained that the visualisation does not show you where the data is located. During the session he mentioned that he was puzzled as to why a dataset without the search word had ranked higher than one with the search term. Despite this, he was able to correctly identify the most relevant datasets with respect to all three axes, even stating that certain datasets were "spatially good but temporally bad".

Once he had launched GeoDOVE he embarked on a process of familiarisation involving translation, zooming and rotation. After approximately 30 seconds of exploring the navigation features of the applet, he declared "the controls are easy". Although he was comfortable with the navigation, he however complained about there not being any indication of the level of zoom, tilt or pan of the current view. The participant's first action on GeoDOVE was to zoom out and view the terrain from an angle of about 30 degrees to the horizontal plane. For the rest of the session, the angle varied between 30 and 60 degrees. Similarly the orientation of the view changed from oblique to orthographic, without a clear indication of any preference between the two. After uploading the DEM he indicated that it was much easier to see the buildings when the DEM is in the background. The participant also complained that he could not see the 3D model of the rock after uploading it. This was because GeoDOVE does not automatically move the viewer to the location of the newly added layer or model.

Participant 2

Upon launching STORM from the NEEDS geoportal, the participant exclaimed that the Java security pop up was “scary”. When entering the search parameters he commented that the coordinates label was useful however the map was in an unusual projection (WGS84) instead of the British National Grid which is commonly used in the UK. When the first search is invoked his attention was drawn to the table of search results. He had noticed that the relevance scores were incomprehensible, prompting him to re-examine the search parameters. Indeed, he had entered erroneous coordinates. After results from the corrected search had been visualised, he suggested that the thumbnails showed that the spatial extents of the datasets cover areas close to one another but did not illustrate “how close”. He then rotated the grid such that the temporal axis was increasing away from the viewer, thereby making only the spatial and semantic axes visible. After rotating the grid, he acknowledged that he had ‘hidden’ the temporal axis because it was not relevant to his needs and therefore he was using the visualisation from a 2D perspective. He commented that he was “highly confused” by a dataset that appeared to be highly relevant judging by its metadata, but had been semantically ranked the lowest – this was because the dataset also had a several terms that were not relevant. Throughout the session, he never zoomed into the graph, preferring only to rotate it instead. He commented that after using the browser for a while it is quite irritating that the search window is not open by default and he has to click on a button to open it.

For the tasks involving GeoDOVE, the participant entered a query in the NEEDS geoportal and when the traditional textual list appeared he commented that the datasets required closer inspection to determine their relevance to the query. His first action once the DEM had been uploaded was to rotate and view the visualisation from directly over head (in orthographic projection). This suggested a preference for 2D visualisation and was consistent with the participant’s use of the STORM browser. In contrast to his earlier use of STORM, he navigated GeoDOVE by zooming, rotating and translating the visualisation.

Participant 3

The participant successfully invoked a search on the STORM browser. Upon presentation of the search results, he stated that he was puzzled by how a dataset with the keyword 'fossils' was ranked lower than one with the keyword 'macrofossils'. To investigate this issue further, he rotated the 3D grid such that all three axes were visible. The participant therefore, appeared to prefer observing the visualisation from an oblique angle with the coordinate of highest relevance (1, 1, 1) being the closest to the viewing camera. Although he could identify which datasets were closest to the coordinate of highest relevance, he commented that he did not fully comprehend how the application inferred relevance.

For task 3 involving use of GeoDOVE, the participant observed the visualisation from an angle of about 30 to 60 degrees from the horizontal plane. However, he changed to an angle of approximately 20 degrees for task 4; meaning he viewed the visualisation from as close to the DEM as possible. Throughout the session, he was comfortable navigating the 3D visualisation. In response to the requirements for task 3, the participant stated that he could not adequately determine which part of the DEM was the 'deepest red'. Technical difficulties were encountered during the session resulting in loss of some of the recorded material. However, notes compiled during the session were sufficient to overcome the problem.

Participant 4

The participant's first comment was that NEEDS has a very simple interface and it does not take long to identify the search tools. Once he had launched the STORM browser, his first action was then to click on the server listbox to confirm that the NEEDS metadata server had been selected for performing tasks 1 and 2 of the usability study. Instead of using the map to select the area of interest, he entered the coordinates directly into the textboxes. He was able to comfortably translate, zoom and rotate the visualisation. He appeared to prefer selecting datasets by clicking on the ranked list rather than by clicking on the thumbnails. He was able to identify clusters of datasets with similar spatial scores. He however, appeared to ignore semantic and temporal scores.

For the tasks involving GeoDOVE, the user first attempted to identify which datasets contained buildings data by examining the metadata presented by the NEEDS geoportal. He acknowledged that closer examination was necessary to fulfil the given tasks. Upon launching and exploring the navigation features of GeoDOVE, he commented that navigation in GeoDOVE was slightly more intuitive than the STORM browser, particularly the mouse-based scrolling (zooming in through mouse scrolling is disabled in STORM due to its single unit axes). After uploading the subsurface 3D model, he was puzzled because the model did not appear anywhere in the visible scene. He then navigated around the scene until he could see the 3D model. For most of the session he viewed the terrain from an angle of approximately 45 degrees to the horizontal plane.

Participant 5

The participant first confirmed that the browser was set to search the NEEDS metadata server. She then entered the query parameters directly into the textboxes and invoked a search. When the results appeared she rotated the 3D grid such that the optimum position (1, 1, 1) was closest to the viewing camera. Next, she clicked on some records on the results table and identified which thumbnails had been highlighted. Then she clicked on a thumbnail near the previously highlighted one and then examined the row on the results table, followed by the associated metadata panel. She then systematically clicked on each thumbnail or table row and examined the metadata of the selected dataset. For most of the session she rotated 3D visualisation without zooming in or translating the grid.

For the tasks involving the use of GeoDOVE, the participant was able to navigate to the applet from the NEEDS home page. Once the dataset had been uploaded she managed to add the DEM to the visualisation. She familiarised herself with the applet by zooming and rotating the visualisation to an angle approximately 60 degrees to the horizontal plane. After sometime however, she experienced some difficulty rotating the visualisation to a similar 60 degree angle. This difficulty in navigation caused her to abandon attempts to locate the uploaded subsurface model of a rock (for task 4).

Participant 6

The participant's first action was to launch the browser then familiarise himself with the navigation controls. He systematically tried to rotate both vertically and horizontally but then realised that the vertical rotation is disabled. Initially, he attempted to zoom towards a thumbnail but encountered some difficulty. This was because the thumbnails were offset from the centre of the screen and thus they disappeared as he zoomed in. Replay of the video protocols shows that by the time the participant had reached task 2, he was now able to zoom to specific thumbnails by translating them to the centre of the screen before zooming in. Whilst carrying out task 1, he observed that some of the documents had high semantic scores because they contained the search term 'agricultural', however they did not discuss any agricultural concepts (this was an example of erroneous entry of keywords). This shows an ability by the participant to cross-reference titles on the ranked table with thumbnails on the 3D visualisation.

Whilst familiarising himself with GeoDOVE, he commented that the navigation controls were better than those of the STORM browser, particularly the zooming tool which is controlled from the mouse scroll. For most of the session he rotated the terrain such that it was at an angle of approximately 45 degrees to the horizontal plane. Despite the availability of a graduated legend on the applet, the participant was not able to identify parts of low altitude on the DEM. His verbalisations suggest that he was not able to identify the parts of the DEM coloured in deep red (the symbol the legend assigned to areas likely to be flooded).

Participant 7

The participant commented on how the visualisation gives a clear indication of the relevance of each dataset. He rotated the graph to view only a single plane at a time and then rotated it to an oblique angle. He then stated that "the visualisation is good" because one can view relevance from any dimension. He further emphasized the importance of the metadata panel for describing datasets. However, he stated that he was confused by the ranking of some datasets, in task 1, as being highly relevant on the ranked list but not on the 3D visualisation (when viewed from position 1,1,1). For task 2 he observed that the most relevant datasets on the ranked list were also the closest to the optimum position of (1,1,1).

For task 3, the participant rotated the terrain and viewed it from an angle of 20-30 degrees. He was able to comfortably translate the viewing camera around the visualisation; however, at times he found vertical rotation difficult. He commented that the colours, particularly their blending and brightness, were good for analysis as it clearly highlighted areas that are susceptible to a flood. He also commented that he was able to identify areas that were at a higher altitude and thus relatively free from potential flooding. For task 4 he was able to correctly identify datasets that had the subsurface structure underneath the roads (as had been required of all participants).

Participant 8

For tasks 1 and 2, he viewed the visualisation from oblique angles. Initially, he encountered difficulty zooming into specific thumbnails but was able to find a methodology for doing so. He commented that he was confused by how datasets without the query term were ranked, particularly as some of them had been ranked higher than those with the query term. He was however, able to identify datasets that were highly relevant to the query and those that were not by observing the positions of their thumbnails relative to the optimum position (1,1,1).

For most of the session he viewed the terrain from an oblique angle of approximately 40 degrees to the horizontal plane. He comfortably navigated the visualisation; however, he commented that the vertical rotation was difficult to use. For tasks 3 and 4, he was able to identify datasets that were relevant to the given scenarios. He was able to navigate between the geoportal and the applet, in order to examine each dataset, but commented that it would be better if the applet could access all the datasets returned by the search on the geoportal.

Appendix E Features of Prototypes

The following are key features of the STORM browser

- 3D visualisation panel : This is the main panel for presenting search results through the STORM visual metaphor.
- Calibrated axes: The axes aid users in detecting distances of thumbnails from one another and from the principal coordinates (0,0,0) and (1,1,1).
- Metadata panel: Allows users to view attributes such as the title, abstract and geographic coordinates of each dataset.
- Title listing panel: Enables users to select datasets easily and presents numeric values for each similarity measure.
- Search by keyword: Allows for querying titles and abstracts.
- Search by range of dates: Allows for temporal constraints on a search.
- Search by geographic footprint: Allows for retrieving datasets with footprints that intersect the query footprint.
- Z39.50 Connectivity: Allows for searching several distributed Z39.50 metadata servers.

The following are key features of the GeoDOVE applet

- 3D visualisation panel: Panel from which geovisualisation is presented.
- Direct upload from geospatial databases: Allows for datasets to be transmitted through the internet directly to the applet.
- Vector data extrusion: Offering 2.5D visualisation of originally 2D vector datasets.
- Creation of 3D Digital Elevation Models from raster data: Allows for surface models to be added to geovisualisations.
- Ability to upload 3D VRML files: Allows for pre-built 3D models to be added to a visualisation, increasing content in a presentation.
- Facility for viewing attribute tables: Allows for attributes to be viewed together with the objects in the visualisation. Selection of records in the table highlights the associated objects in the visualisation, and vice versa.

- SQL query support for both the attribute table and 3D visualisation: Enables multi-object selection and selection through predicates and constraints.