An Intelligent Information System for Maritime Navigation Information

Raphael Malyankar
Dept. of Computer Science
Box 875406, Arizona State University
Tempe, AZ 85287-5406, USA.
Tel: (480) 965-2804
E-mail: rmm@acm.org

Abstract

This paper describes aspects of an information system being constructed for intelligent information retrieval in maritime navigation applications. The system is based largely on digital nautical charts, and the construction of a taxonomy of chart features and concepts related to the maritime navigation domain is described. This taxonomy is used in formulating queries and sub-queries for disparate sources of information that supply information about navigation hazards, aids to navigation, and other relevant local information.

Introduction

Mariners use a wide variety of information collected from a number of sources, ranging from maps or charts to weather reports. Large quantities of information must thus be processed, especially for near-shore travel where shore effects (tides, currents, shore installations, etc.) must be considered. Further, any map or chart represents information only in context and for a specific purpose (e.g., physical and political maps of the same region). The context in which information is sought is therefore important for creating a responsive system.

The most sophisticated ENC (Electronic Navigational Chart) systems and digital cartographic systems currently available still deal with only very basic geospatial information, for example, routes and waypoints, currents, and graphic overlays of one kind or another. They are capable of only relatively basic geographic information retrieval and analysis. This paper describes a pilot project on creating an ontology (suitable for automated reasoning) for marine navigational knowledge and its use in a prototype retrieval system that answers queries relating to maritime navigation. The information included in the current version of the system includes mainly chart-based information, some procedurally represented information, and a subset of real-time information sources.

The aim of this effort is to allow automated systems to use geospatial information better, so that it is possible for a navigational program to "understand", for example, that shorelines can be crossed by aircraft but not by surface vessels, that tidal tables are important for near shore navigation, that routes may need to stay away from restricted waters, that cutting across shipping lanes should be minimized, etc.

This paper describes the ontological engineering involved in building such a system, use of some typical information sources, and prototype methods of query formulation and information retrieval within this problem domain.

Information Sources

Some of the information sources used are outlined below.

Digital Nautical Charts

Digital nautical charts are available in vector and raster formats. Vector format charts represent features by collections of connected points (individual point features are also included). A *feature* is an individual chart element denoting a geographic entity. Coastlines, buoys, ports, and lights are all examples of individual features in nautical charts. The digital charts used here are vector format charts containing information classified as harbor, approaches, and coastal and general charts. The database is organized as a directory hierarchy with subdirectories and files corresponding to *coverages* and *feature classes.* A *feature class* is a set of geospatial features of the same kind (e.g., buoys are one feature class; administrative boundaries are another). A *coverage* is a collection of feature classes all relating to one category of knowledge; for example, the 'navigation aids' coverage includes feature classes for different kinds of aids to navigation.

Tidal Computation

Information about tides must be represented either procedurally or in tabular form, because of the special constraints of tide information. Tide heights, currents and times are location-specific, only weakly correlated with nearby locations, and vary from day to day (cycling with a period of approximately 19 years). Algorithmic representation of tide information uses a number of harmonics (sinusoidal functions calculated from historical tide data) to predict high and low tides. Manuals of harmonic analysis and tide calculations have long been available, and it is a relatively simple problem to write a program to predict tide height at a specified time at any given station (places where regular tide observations are made). (Our system uses the XTide program written by David Flater and publicly available on the Internet). Tidal currents and heights can be calculated

using tide tables published by maritime authorities (NOAA, in the United States), which contain tide height, timing, and current information for a number of *reference stations*.

Real-time Weather

Weather information and current weather reports have been available on the Web for a few years. Weather data includes information on wind speed and direction, wave height, atmospheric pressure, etc. Prototype systems are available whereby it is possible to obtain relatively up-to-date data over the Internet. This data is also location specific, that is, it is necessary to specify location extents or a named 'station' in order to get the information for that place. Weather broadcasts by radio are the usual way mariners currently get this information. Deployment efforts for text broadcasts are also underway. For obvious reasons, this prototype has been restricted to sources available on the Internet.

Textual Information

Piloting guides and manuals, describing navigation and seamanship techniques are widely available. The "United States Coast Pilot" is a series of publications that provide information to sailors about local navigation hazards, currents, tides, etc. The US Coast Guard generates various documents with navigational information, some of which are online and can be directly accessed for retrieval of necessary information.

The next section discusses the ontological issues encountered in modeling this domain.

Ontological Issues

The nature of the domain and the requirement for a practical automated reasoning system necessitate dealing with the following issues:

- Multiple categories and representations of information: Processes, components, mechanisms, and fields, are all present in this domain, necessitating representation with appropriate ontologies. The relevant information is stored in tables, rules, hierarchical representations, as commonsense knowledge, flat files, structured data, etc. It is necessary to find describe these representations, their characteristics and contents; the system must be able to handle a variety of different representation methods. It is also necessary to represent information such as geographical features, flow processes such as currents, etc., and also fuzzy objects.

- Inaccurate and varying information: Data measurements in this domain are not all accurate to the same extent and will certainly vary with time, both predictably and unpredictably. Predictable variations include such items as tides and unpredictable variations include such items as scour (movement of material on the sea bottom) which will change depth values. Paper charts deal with this by adding warning notes. Allowing for inaccuracy and variations while still retaining practicality is one problem we will need to deal with; we anticipate that the simplest solution, specifying ranges and trends, might not suffice for this domain due to both the nature of the variations and cumulative uncertainty effects for automated inference.

- Large volumes for some categories of data: Charts are densely packed with information; even a single chart covering a small region may contain hundreds of soundings alone. Tagging this data, and devising representations that are tractable for inferencing with large amounts of data is necessary. For example, a sample digital chart covering the approaches to San Diego Harbor, covering an area of 0.25 degrees of latitude and 0.5 degrees of longitude, contains 6400 sounding values and more than 3000 other features. Admittedly, harbor approaches are surveyed in far more detail than other regions of the coastline, and mariners need only a limited number of charts under most circumstances, but these volumes of data are still enough to stretch the capabilities of retrieval systems, particularly since many of the data items are referenced to geographical locations, may be graphical in nature, and, in addition, relationships and proximity determinations are very important.

- Contexts: Categories of information that are useful (or even available) will be context-dependent. For example, tidal information is unlikely to be needed except near shorelines. Map and chart databases are already subdivided by *coverage* and *feature class*. For example, in the Vector Product Format database standard, the *hydrography* coverage in a map database includes *soundings* and *bottom characteristics* as feature classes under it. ('Soundings' are depth measurements; 'bottom characteristic' tells whether the bottom is sandy, rocky, mud, etc.) This subdivision is useful for database modeling and data management purposes. However, the mapping of feature classes to practical usages (i.e., application classes) of this information is many-to-many — one might be interested in some features from the *Hydrography* coverage, most feature classes from the *Navigation Aids* coverage, etc., etc. That is, some application classes will need selected feature classes from selected coverages, and some feature classes will be useful for more than one application class. In other words, the contexts set by database standards do not converge with contexts required for practical use. For paper charts, of course, this problem is solved through human selection of the coverages and feature classes to include, which is done by the chart publisher, using accepted standards for what must be displayed on nautical charts.

- Scaling: Details important at one map scale might be irrelevant at another. This is not necessarily the same issue of context. Means of capturing scaling effects and information abstraction must be explored. In the nautical chart domain, scaling is done by using charts of different scales (1:80,000, 1:600,000, etc.), again selecting, excluding, and combining features through human intervention.

Reuse of Ontologies

Geographical information appears to have been dealt with only cursorily in the ontological engineering community to date, though CYC (Lenat 1995) apparently does deal with geographical information to some extent and does contain theories on related information such as terrain, maps, graphic elements for maps, and weather. There has also been some research reported on the representation of fuzzy objects in the knowledge representation and spatial information representation literature (Burrough 1996; Coucelis 1996; Schneider 1996; Cohn & Gotts 1996a; 1996b). However, navigation information and knowledge are ill-covered in the released CYC upper-level ontology. Apparently most of the knowledge engineering work in this area still remains to be done.

The next section describes our approach to some of the ontological issues.

Ontology Construction

The ontology is based on the coverage and feature classes in the sample digital chart obtained from the National Imagery and Mapping Agency (NIMA). The taxonomy inherent in the organization of this chart was closely followed. This stage resulted in hierarchical coverages and feature classes such as 'breakers', 'coastline', etc. In addition to the taxonomic relationships expressed in the organization of the chart files, National Oceanic and Atmospheric Administration (NOAA) and Federal Geographic Data Committee (FGDC) publications (FGDC 1998) were used to add knowledge about some types of features, their roles, and their interrelationships (the organization of the digital chart is purely tree-like). The NOAA definitions and those in the Spatial Data Transfer Standard (SDTS) are more detailed than the file organization of the sample chart (for example, the chart places all buoys in the same file). Another hierarchy superimposed upon the first was a classification into point, line, and area features. Further, process and field 'features' such as tide height computation were added to the ontology. Source information (whether the feature was available from a digital chart, used a program, etc.) was also included.

Construction used the Ontolingua server at Stanford (ontolingua.stanford.edu) and was done largely by hand using that interface, with some automation to facilitate repetitive tasks. Individual features were not added to the knowledge base. The ontology was translated using the translation facilities on the Ontolingua server and processed locally with special-purpose routines as part of the information system. The pilot version of this ontology is mostly taxonomic, and it is being used mainly for taxonomic knowledge in the prototype system.

Figure 1 shows some of the classes and their hierarchy.

```
Coverage
  Coverage-Aids-To-Navigation
  Coverage-Hydrography
    Bottom-Characteristic-Point ...
    Sounding
  ...
Feature
  Field-Feature
    Tide-Current-Direction-Field ...
  Line-Feature
    Breakers
    Coastline ...
  Obstruction
    Breakers
    Rock
    Wreck ...
  Point-Feature
    Bottom-Characteristic-Point
    Marker
      Beacon
        Light
          Direction-Light
          Special-Light ...
        Radar-Beacon
        Buoy
          Light-Float-Buoy
          Special-Purpose-Buoy ...
      Fairway-Marker
        Leading-Lights
        Lights-In-Line ...
      Fog-Signal
    Rock
    Sounding
    Wreck ...
  Process-Feature
    Location-Dependent-Process
      Tide-Height-Location-Process
```

Figure 1: Class hierarchy

Query Types

Though the data are voluminous, this particular problem has only a limited number of question categories, and queries are general spatially oriented. In particular, one common form of use is an examination of navigational features, hazards and aids in a specified area, for example, along a specified route. Software to 'zoom' in and out on specified areas is already available, and there are more sophisticated tools under development by various entities that can produce overlays of different feature classes and allow point-and-click information retrieval of data about specific features from digital chart files. Overlays of weather information and tide data and predictions are also available from various sources, mainly as prototypes at present, but also in a few commercial systems. The map-based systems on the Web (for example street location systems) are cases in point, though they are mainly for land and especially street maps, though some sources provide weather observations and predictions. The prototype system described here is an attempt to provide more flexible and intelligent retrieval than currently available software.

Queries fall into the following categories:

1. Spatial: Spatial queries are requests for retrieval of all features meeting specified criteria within a specified spa-

tial extent, i.e., "show me the hazards in Long Island Sound". They fall into three main subcategories: extent-based, route-based, and named-location-based queries. Extent based queries specify a rectangular area delimited by latitude and longitude boundaries. Route-based queries request information relevant to a specified route (a sequence of point coordinates). Queries based on named locations request information about features near to, or at, a specified location (e.g., "Long Island Sound").

2. Feature-specific: This category is used to cover questions such as "what landing facilities are available at San Diego" where the expected answer would include descriptions about pier facilities, cargo unloading capacity, perhaps information about pier occupancy and delays, etc.

3. Summaries: These requests entail selection of features, extraction of feature information, then minor analysis of the extracted features. For example, the question "What is the bottom like off Wellfleet?" entails selection of the appropriate bottom characteristic point features and analysis of the bottom type values of the selected features (if any).

4. Inferential analyses: These involve some level of interpretation of the questions; for example, the question "Can I anchor near Wellfleet?" requires translation into queries that retrieve the bottom characteristics for the specified spatial reference, and match the retrieved values to those that allow anchoring (or to those where anchoring is discouraged).

5. Complex Queries: These are combinations of the above types. Requests for weather information, for example, are often couched in spatial terms (e.g., "what is the weather forecast for Cape Cod"), but involve some feature-specific information (e.g., the data recorded by weather buoys in the area of interest).

The next section discusses how these queries are formulated in the prototype.

Query Formulation

Simple queries (from the users' point of view) in this domain can lead to extremely complicated internal queries. For example, responding to a request to supply a course in a near-shore environment will need to take into consideration the following items and more: depth, draft of the vessel in question, currents, the type of vessel (e.g., sail/power), time of day, tide conditions, traffic lanes, hazards, and so on. Question answering in this domain is thus more than information retrieval, but requiring a user to formulate exact or nearly exact queries listing all the kinds of information needed will result in an unwieldy system. An intermediate choice was made: the system accepts a range of queries either by selection from a repertoire of known choices, or by pattern-based analysis of a text question that uses a restricted vocabulary and restricted set of phrases.

The word and phrase vocabularies are a collection of pre-specified words and question patterns encoded into the system, and are used to formalize the questions put to it. A form-filling interface accessible with a Web browser is utilized for interaction. Queries have two elements: a spatial reference and a 'text' section. The spatial reference is supplied in the form of an extent specification (specification of the corners of a box in latitude and longitude) or a path (a sequence of waypoints, also as latitude/longitude pairs). The question itself is entered in another section of the screen. Questions may be entered as text or by selection from a list of choices. Text questions must be posed using a restricted subset of English phraseology. Text questions are matched against the aforementioned internal phrase library, which is a set of phrases that have meaning in the context of the domain. (No attempt to create a sophisticated natural language interface is made.) Once these items are entered, the following sequence of actions is taken:

1. The question is translated into a high-level intermediate format consisting of the tuple:
   `[<context>] <head> {<req>} <spatial reference>`
   where `<context>` is a context for the query, which is used to specify background information that may be necessary for qualifying lower-level queries or for later filtering;
   `<head>` is the type of query, for example, any quantification information such as *forall* or *some*;
   `<req>` is the requirement or subject of the question, for example, *hazards* or *navigation aids*;
   `<spatial reference>` is the spatial reference accompanying the query, consisting of the type of spatial reference (e.g., extent or location) and the value of the reference (e.g., the corner points for an extent box).

   The brackets `[]` and {} are used to denote optional and repeatable fields respectively.
   The `<context>` portion is used for filtering the results; the `<head>` portion is converted into an existential or universal quantifier, or into operators that must be applied to the retrieved results (e.g., *max()* and *min()*). This intermediate form is then converted into conjunctive normal form if needed, using the same fields.

2. The high-level intermediate form is translated into low-level intermediate form, using the taxonomic knowledge from the ontology. This step performs two functions: separation of complex requirements (the `<req>` field above) into simpler components based on the structural components into which the system's knowledge representation systems are organized, and translation of the requirements into the terminology of the underlying KRS or database (without converting the query completely into KRS- or database-specific format).

3. The low-level intermediate form is translated into KRS- or database-specific queries using the appropriate query formats, which are values of the source descriptions obtained from the aforementioned ontology.

For example, given the (user-supplied) extent

$$[\ (-117.5, 32.5)\quad (-117.0, 32.75)\ ]$$

which are the south-west and north-east corners of a box bounded by the meridians $117^o W$ and $117^o 30' W$, and the parallels $32^o 30' N$ and $32^o 45' N$ and the query "What are the dangers?", the following process is gone through:

The text query is translated into a high level intermediate; the sample given would translate to

```
:context nil :head extract :req *dangers*
:ref extent -117.5 32.5 -117.0 32.75
```

Translation into lower-level form gives

```
:context nil :head extract :req
obstructions :ref extent -117.5 32.5 -117.0
32.75
```

This particular problem is one of extraction of the relevant features from the chart file, and so it is translated into command-line arguments for the feature extraction utility: `-feat '`*<chart location>*`, obs, Obstructions, reefa.aft, ...' -ext -117.5 32.5 -117 32.75` *<other arguments>*`'`
to extract all the reef area features, and similar sets of arguments for other features.

The use of OKBC-friendly format will be investigated later, as this will necessitate layering an OKBC interface over charts, etc.

Pilot System

A pilot system has been constructed that uses digital chart data files, tide prediction software, and weather information sources to answer a limited set of questions. The sample chart used was obtained from the National Imagery and Mapping Agency, and covers the approaches to San Diego harbor. Approximately 10,000 entities are included. Features are classified as point, line, area, text, and complex features. There are 12 coverages (excluding the reference coverages) and 43 feature classes (not all feature classes are populated). As mentioned, the ontology constructed includes this organization as one of its hierarchies.

Tide heights were calculated as required, using the XTide software mentioned elsewhere in this paper. This utility calculates tide heights at specified tide stations at given times (and for a period after the given time). While the results are not intended for actual navigation, they suffice for the purposes of this pilot system. Weather information was obtained from historical files downloaded from the National Data Buoy Center (NDBC) Web site and from queries to this site. There are other sources of weather, but again the weather buoy data was considered sufficient for feasibility investigations. (While these shortcuts are, we believe, acceptable in an initial prototype, a deployable information system would need significant investments of time and labor in locating and integrating better sources, and would need to integrate sources from elsewhere than the Internet too).

The approach to ontological issues and query formulation has been discussed in earlier sections and will not be repeated here, except to remark that in its present form, the use of the ontology is mainly for translation and taxonomic reasoning; other significant inferential procedures have not been investigated at this time.

Feature retrieval is done using adaptations of software available from NIMA that reads and processes charts in the format they use. This program accepts an extent (a latitude/longitude box) and feature class as command-line arguments and extracts features of the specified class that are located within the specified extent. This means that features involved in spatial queries based on extent specifications can be extracted relatively easily; other forms of spatial reference, such as route-based queries entail conversion of their spatial specifications into combinations of extents, with a slowdown in performance.

Identification of sources, and formatting of queries for different source, is done using the relationships between feature classes and sources that were included in the ontology developed earlier; that is, one function of the ontology is as a data model of the three information sources mentioned.

Presentation of Results

Presentation of results is extremely important for this domain; the most important information must be present and should be clearly apparent. Results may be viewed under conditions of time or other stress. Different kinds of response presentations are needed: for example, plain-text responses, course plots, graphical elements of different kinds. For example, in search and rescue contexts, a likely query is to plot a search pattern given information about the target (small boat, raft, person in lifejacket, etc.) and search resources available. For this particular application, wind speed and direction, searchers' capabilities (speed, endurance, etc.), currents, and weather conditions must all be taken into account. Presentation schemes for the pilot system are currently under exploration.

Future Work

Plans for the future include the addition of spatial representations and more sophisticated spatial reasoning methods, expansion of data by one or two orders of magnitude and evaluation of performance with this higher volume, integration of other sources of information, particularly information in text form, and development of a presentable user interface. The use of a KRS for serving chart entities (that is, loading the features into a KRS) and evaluation of performance are also planned. Evaluation of the system is expected to be undertaken by comparing the responsiveness, flexibility and range of information retrieval with those of commercial and under-development systems that fulfill similar mariner needs.

Limitations

The most important limitation is the rudimentary nature of spatial reasoning in the current incarnation of the system. Relationships between feature location, for example the adjacency of point features is neither encoded not inferable (except by direct computation of distances using latitude and longitude values). Route planning (in the sense of determining a sequence of waypoints) is currently naïve. An initial attempt to add spatial relationships to the knowledge base was commenced, but the large number of items even in

a small area makes it unlikely that exhaustive inclusion of spatial relationships will be useful or feasible. Limiting explicit spatial relationships to certain elements, that is, a subset of the area features and selected line and point features, while depending on numeric computations for determining adjacency and closeness when necessary for other classes of features is now being explored. We do not anticipate creating rigorous route planning procedures (that is, procedures that produce a sequence of waypoints demarcating a feasible path between a specified starting point and destination) based on formal logic or similar representations, mainly because there are already good (though computation-intensive) methods for solving this particular problem, using vector or area representations of space. Naïve versions of route planning will be used where appropriate. Another limitation is the inability to reason about unpredictable changes in information, for example discrepancies in navigation aids and sounding depth values (all that can be done is to keep the chart database as up-to-date as possible). Note that predictable and time-dependent changes, can be retrieved and announced to users (and the tide computation program does this).

Related Work

Formal ontology development for the geospatial domain is beginning to attract interest. There has been some discussion of this in "specialist meetings" of the National Center for Geographic Information and Analysis (NCGIA) (Mark, Egenhofer, & Hornsby 1997; Peuquet, Smith, & Brogaard 1998). Spatial representations and reasoning have been a continuing interest in the areas of spatial information theory and geographic information systems. The latter work has investigated some fields of interest to ontology research, particularly the representation of object extent and boundaries, for example (Coucelis 1996; Frank 1996; Smith & Varzi 1997), and the representation of fields and flows (Kavouras 1996; Peuquet, Smith, & Brogaard 1998).

The CYC project (Guha & Lenat 1990; Lenat 1995) (of MCC and later Cycorp) contains some geospatial information in its "Geography" collection and there is apparently a certain amount of current interest in using Cyc in GIS. An examination of the information made public by Cycorp (the "Upper Ontology") indicates that the geospatial relations may be limited to high-level information retrieval rather than navigation. For example, the handling of cartographic units and directions does not appear to have been designed for navigation and chart applications. Neither the list of applications for CYC mentioned by Cycorp, nor the projects currently listed as being part of the DARPA High-Performance Knowledge Base (HPKB) project appear to include detailed geospatial knowledge of the kind described in this paper. The work of Ferguson, Donlon, and Forbus on spatial reasoning (Ferguson & Forbus 1999) and trafficability (Donlon & Forbus 1999) may be the closest-related work in AI, but though techniques and representations may carry over, significant extensions (at the very least) will be needed before this system can be used in this domain.

Commercial information systems dealing with maritime information are, at present, generally confined to charts, tide, and port information. These are GUI-based systems requiring point-and-click interaction with either a chart or a menu of overlays containing pre-specified features. Combinations of overlays (e.g, coastlines and navigation aids) is possible. Systems displaying nautical charts on personal computers are already commercially available, and under-development applications are planned for real-time distribution of data about port and cargo loading facilities. These systems do not include any internal reasoning about what data should be presented or any internal analytic processing of data.

Conclusion

This paper has described the creation of an ontology of nautical chart elements and its use in an information system for maritime navigation information. An exploration of query formulation for certain classes of related queries was also described. The ontology is currently being used largely for taxonomic reasoning and as a data model of a representative collection of information source types. The system described herein is intended as a platform for exploring the development, representation, and retrieval issues involved in adding intelligence to geospatial information systems, particularly in the mapping and charting domain, and an investigation of problems and areas for future effort in relation to these is currently underway. An evaluation of the utility of 'well-known' knowledge representation and reasoning techniques in this particular problem domain is also being undertaken.

Acknowledgements

The ontology was constructed using the Ontolingua system at the Stanford Knowledge Systems Laboratory Ontolingua server. Tide height predictions were computed using David Flater's *XTide* program. Real-time and historical weather buoy data were obtained from the NOAA National Data Buoy Center's publicly accessible Web site. Sample chart data files were obtained from the National Imagery and Mapping Agency. Use has been made of certain software routines and utilities that are part of NIMAMUSE, a NIMA software product. The use of the software does not indicate endorsement or approval of the work described in this paper by the Secretary of Defense or the National Imagery and Mapping Agency. This material is based on work partially supported by the National Science Foundation under Grant No. EIA-9876604. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation.

References

Burrough, P. A. 1996. Natural objects with indeterminate boundaries. In *Geographic Objects with Indeterminate Boundaries (GISDATA2)*. London: Taylor and Francis Ltd. chapter 1, 3–28.

Cohn, A. G., and Gotts, N. M. 1996a. The egg-yolk representation of regions with indeterminate boundaries. In *Geographic Objects with Indeterminate Boundaries*. London: Taylor and Francis Ltd. chapter 12, 171–188.

Cohn, A. G., and Gotts, N. M. 1996b. Representing spatial vagueness: A mereological approach. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Fifth International Conference (KR '96)*, 230–241.

Coucelis, H. 1996. Towards an operational typology of geographic entities with ill-defined boundaries. In *Geographic Objects with Indeterminate Boundaries*. London: Taylor and Francis Ltd. chapter 3, 45–56.

Donlon, J., and Forbus, K. 1999. Using a geographic information system for qualitative spatial reasoning about trafficability. In Proceedings of the Qualitative Reasoning Workshop, Loch Awe, Scotland. www.qrg.ils.nwu.edu/papers/papers.htm.

Ferguson, R., and Forbus, K. 1999. GeoRep: A flexible tool for spatial representation of line drawings. In Proceedings of the Qualitative Reasoning Workshop, Loch Awe, Scotland. URL: http://www.qrg.ils.nwu.edu/papers/papers.htm.

FGDC. 1998. Spatial data transfer standard. Federal Geographic Data Committee, U. S. Geological Survey. Proposed standard.

Frank, A. U. 1996. The prevalence of objects with sharp boundaries in GIS. In *Geographic Objects with Indeterminate Boundaries*. London: Taylor and Francis Ltd. chapter 2, 29–40.

Guha, R. V., and Lenat, D. B. 1990. *Building large knowledge-based systems : representation and inference in the Cyc project*. Reading, MA: Addison Wesley.

Kavouras, M. 1996. Geoscience modelling: From continuous fields to entities. In *Geographic Objects with Indeterminate Boundaries*. London: Taylor and Francis Ltd. chapter 21, 313–324.

Lenat, D. B. 1995. Cyc: A large-scale investment in knowledge infrastructure. *Communications of the ACM* 38(11):32–38.

Mark, D. M.; Egenhofer, M. J.; and Hornsby, K. 1997. Formal models of commonsense geographic worlds: Report on the specialist meeting of research initiative 21. Technical Report 97-2, National Center for Geographic Information and Analysis, San Marcos, TX.

Peuquet, D.; Smith, B.; and Brogaard, B. 1998. The ontology of fields: Report of a specialist meeting held under the auspices of the Varenius project: Panel on computational implementations of geographic concepts. Technical report, National Center for Geographic Information and Analysis, Bar Harbor, ME.

Schneider, M. 1996. Modelling spatial objects with undetermined boundaries using the Realm/ROSE approach. In *Geographic Objects with Indeterminate Boundaries*. London: Taylor and Francis Ltd. chapter 10, 141–152.

Smith, B., and Varzi, A. C. 1997. Fiat and bona fide boundaries. In *Spatial Information Theory, A Theoretical Basis for GIS: COSIT '97, Laurel Highlands, Pennsylvania*, 103–119. Springer-Verlag.