PRODUCING ANSWERS TO SPATIAL QUESTIONS

Gail Langran
Intergraph Corporation
2051 Mercator Drive
Reston, VA 22091

ABSTRACT

This discussion explores alternatives to standard GIS command procedures. The goal is for the user to describe the information he or she seeks rather than the data manipulations that should be performed so the system can provide appropriate display content and format. If the same information model were used for both user input and display generation, then spatial, temporal, or thematic questions could be matched to tabular or graphic answers. Using a model of geographic information, the potential for an artificial language that would permit a user to phrase geographic questions using English-like grammar and language is examined.

INTRODUCTION

Since their inception, GISs have become increasingly more sophisticated in terms of standardization, data structuring, error control, and analytical options. However, ease of use is still a major obstacle to the full exploitation of GIS technology. Many systems force new users to enter the world of GIS by navigating a maze of command-line interfaces, voluminous documentation, and user-built displays. This initial investment in learning a new tool may be beyond the means of busy analysts with other options for performing their analyses. It also speaks ill of the GIS discipline; if the main responsibility of an information system is communication of information, it follows that no mattter how sophisticated the storage and analytical capabilities, these systems somehow fail unless communicative powers are developed (Webber 1986b).

The harshness of the GIS user environment is gradually softening. On-line help sequences and self-explanatory point-and-click input forms are changing the book-on-the-knees posture so common to users of command-line GISs. Special-purpose systems that address the specific needs of a particular application (for example, oil exploration) tend to be simpler to use because they are geared to users whose expertise is expected to be in areas other than GIS. But the difficulty of formulating a set of commands in a multi-purpose GIS and the tedium of selecting an appropriate display format for responding to queries to the system remains.

Special-purpose cartographic query languages do exist. Nyerges (1980) developed a query language geared specifically to cartographic purposes that permitted users to request information via a grammar and keywords that the system could decode. Frank (1980) developed a query language that could automatically produce a display of the data selected if relational logic alone were used. Broekhuysen and Dutton (1983) describe the design of the Odyssey command language and their efforts to achieve the effect of a dialog with the computer. Morse (1987) describes an expert system for forest resource management that accepts if-then statements from a user, translates those into GIS commands, and produces a standard report.

Shortcomings of the present approach

A universal problem among current approaches is the intuitive barriers that users face when combining relational and spatial logic to describe a course of action to the computer. Egenhofer et al. (1989) have addressed this problem by augmenting the query language style with "point and click" methods of indicating the objects involved. While this approach will ease the plight of the user, it does not change the fundamental fact that so-called "query languages" are essentially ways of defining a subset of data to retrieve without indicating the purpose of the retrieval. In contrast, the ideal mode of discourse with a GIS is for a user to describe what information is sought rather than how the system should manipulate the data to produce the information. Compare a high-level programming language to assembly language.

High-level: *Add 2 + 3 and store the results in A.*

Assembler: *Place the values 2 and 3 in registers, add the registers, place the result in another register, whose number the user must track for subsequent manipulations of the result.*

Today's GISs users express their information needs in the semantic equivalent of assembly language. A preferable mode of discourse with the system would provide the user with a means to express his or her information needs, as opposed to data-manipulation steps. Compare the following dialogs.

Ideal: *Tell me where cow pastures border this stream.*

Current: *Build a narrow buffer around the stream, overlay the buffer with the data, select cow pastures from within the buffer, create a map whose extent covers the stream, shade the agricultural areas in a selected pattern or tint.*

The fact that the logic of data retrieval and spatial overlay is slippery to many GIS users does not imply that GIS users are slow when compared to users of other information systems. Rather, the methods we presently use to interact with computers are designed to be straightforward to computers rather than to humans. Katzeff (1989) tested novice database users on their ability to construct a query and predict a response. Most of the users did correctly predict the response, but only one-fifth were able to construct a correct query. A question and a desired answer may be clearly in mind, but translating that question into a command sequence that produces the desired answer is a challenge.

Given that questions and answers are clear, while the precise data manipulations needed to obtain an answer are not, it is worthwhile to investigate whether an information system could be designed to receive "raw" questions and produce comprehendable answers without involving the user in the specification of data manipulations. Information needs are easily expressed as questions to the system. Where are all the slopes steeper than 10%? Who owns this parcel? When did this parcel last change hands and who was its previous owner? If a user could inform a system of the information being sought, the system would have the raw materials to provide more helpful help sequences, on-line documentation, and default displays. A move toward this ideal requires two major enhancements to the current approach: users need a more straightforward way to express their information needs and GISs need an automated method of choosing a display format.

Natural language is one route to facilitating human/computer dialogs. But practical considerations demand that we proceed to investigate a higher level of human-machine discourse without waiting for natural language processors, since work in that area is still in its infancy (see Quillian 1985, Schank and Rieger 1985, and Webber 1986a for surveys of the natural language approach; see Morse 1987 and McGranaghan 1989 for geographic applications of natural language). A simple artificial language that operates at a similarly high semantic level could produce results comparable to a natural language, since it appears that syntactical constraints do not impede users if the underlying logic of the discourse is clear. Borenstein (1986) compared the learning abilities of users equipped with natural-language help sequences to those who used a verb-noun artificial command language and found no significant difference. Thus, this work investigates the design of an artificial language to describe geographic questions.

The next section describes the problem in more detail and is followed by a presentation of a theoretical basis for designing a high-level question-answer mode of discourse for users of GISs. Later sections examine the elements of an artificial language to express geographic questions in such a way that the computer can answer with default displays, and summarize the goals and findings of this study.

STUDY GOALS

A high-level method of human/computer dialog that describes information needs rather than data-processing instructions is evidently desirable. The issues that must be addressed for such an improvement to become a reality are

What syntax would the artificial language use?

How would questions be linked to data-manipulation commands?

How would the system choose appropriate display formats?

The first and last of the three issues are open questions. The second issue, while equally challenging, has been treated to some degree. Wu et al. (1989)

describe a frame-based GIS that can receive an expression in a high-level formal language, translate the expression into primitive data manipulation procedures, order the primitives by predefined optimization rules, and execute the commands. The query language developed by Nyerges (1980) also had a multi-level structure comprised of a query language, query decoder, and query processor. While enhancements are always useful, the work of Wu and Nyerges demonstrates the premise that GISs can be informed with sufficient intelligence to match high-level commands to low-level procedures (Table 1).

Table 1. Producing answers to geographic questions. An event sequence includes intermediate dialog between the human and the computer to verify the treatment of a question.

Human	**Computer**
Frame a question or questions.	Match the questions to a set of commands.
Verify the command structure.	Propose a default display format.
Verify display format.	Perform requested data manipulations.
	Produce display format.

To address the questions of syntax and display selection, it is useful to form some preliminary requirements. The query syntax should implicitly embed the information needed to select the appropriate content and format of a response display. If that were the case, standardized answers to questions could be produced in a range of formats without burdening a user with cartographic decisions, either at the micro (e.g., color or gray scale, pattern, shape, generalization, font) or macro (e.g., map type, geographic window, scale selection, entities depicted) level. Several interesting attempts have been made to automate the type of decisions referred to here as "micro" (see Robinson and Jackson 1985, Muller 1986, Mackeness 1987, and Weibel and Buttenfield 1988). It is the "macro" decisions that remain unaddressed. Macro decisions are linked less to legibility than to semantic integrity; in other words, a person may respond to a question in well-modulated tones and correct grammar, but if the answer is off target, the effort is in vain.

If questions are straightforward (e.g., "Where are all the forests in this area?") a display is relatively easy to produce automatically. Frank (1982) describes a GIS query system that could match simple queries with displays by automatically determining a window and scale, then selecting the requested relational entities that would appear. But as demands on GISs become more intricate and GISs themselves become more complex (e.g., by incorporating temporal information), the selection of default displays becomes correspondingly more complex. Entities may be mentioned in a question that are not included in the display that answers it. More complex questions require more complex syntax and a broader system vocabulary. And many output formats become possible and needed (Table 2).

Table 2. Different questions require different formats for answering.

Question	Format	Content
Where are the pastures in this watershed?	Map	Highlight pastures
What are the different types of agricultural land use in this watershed?	Map Listing	Color or list by agricultural type
Who owns the pastures that border the stream?	Map Listing	Color or list by owner
What is the grazing density of the pastures beside the stream?	Map Listing	Shade or list by density
Where does the stream border pastures?	Map	Highlight stream segments
What pastures have changed from nonagricultural use?	Map Listing	Highlight or list pastures
When did they change?	Map Listing	Shade or list by years
How many acres changed?	Value	Number of acres

The problem of default displays becomes more complex when a GIS includes temporal data, since, in addition to spanning space, the queries can span time and space/time. For example, the simple question of where a given feature or attribute type occurs becomes more complex when the question of occurrence concerns a time in the past (i.e., where it occurred ten years ago); a timespan (i.e., where it has occurred any time during the past ten years); flows, motion, or trends (which imply a timespan when the movement occurred); or a change over time (i.e., where it has changed from one feature or attribute to another over a given timespan). Vasiliev (1990) discusses different forms of temporal maps and provides examples of the many methods available for expressing spatial change in graphic terms.

Current GISs place the onus for specifying how the reply should look on the user (Figure 1). Will it be graphic or tabular? If graphic, will it be an outline map dotted with point symbols or a choropleth map? What colors, shapes, and patterns should be used? If the display is tabular, what are the rows and columns? How wide should they be? Ideally, the next generation of GISs will shoulder this responsibility unless a user specifically asks to share it.

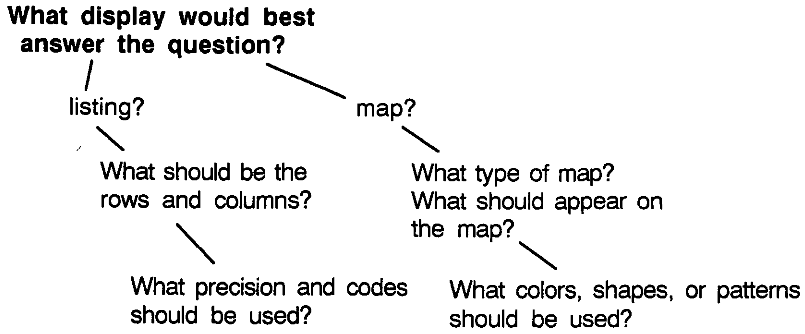

Figure 1. Display options for answering a question posed to a spatial database.

THEORETICAL BASIS

Evidently, some mechanism must exist for mapping questions to formats for answers. Two general approaches are possible: enumerate high-level information requests and map each to a default display, or develop a classification scheme for questions and answers and a means of recognizing what class of question has been posed.

Enumerating GIS operations

Many different attempts have been made to enumerate GIS operations. The capabilities that different authors describe as useful to a GIS can be divided into two groups: information desired, and functions available. Table 3 is an aggregated listing of these two classes that was collected from Nystuen (1968), Salmen (1977), Honeycutt et al. (1980), White (1984), Wu et al. (1989), and Guptill (1989).

Table 3. Types of GIS information and functions.

information desired	**functions available**
multi-scale analysis	windowing
multi-map compositing	rotate, shift, scale
spatial clustering and aggregation	attribute aggregation
edge detection	map overlay - union
direction of flow	map overlay - intersection
comparison	map overlay - negation
precedence	calculate area, length, volume
coincidence	calculate aximuth, bearing, coordinates
proximity	calculate statistics from tabular data
adjacency	buffer zoning: erode and dilate
interpolation	search (locate all)
corridor delineation	line smoothing or simplification
slope and aspect	point in polygon
optimum path	point in line
feature recognition from geometry	line in polygon
weighting	
intervisibility	

The problems of enumeration become apparent when one investigates the enumerations that exist: no two enumerators have arrived at identical listings, the completeness of any given enumerations is debatable, and none include a temporal dimension. Given the immaturity of GISs in general, and temporal GISs in particular, few would argue that any enumeration of information needs could be considered exhaustive. A question-to-answer mapping built on this approach would need to be adjusted or expanded continually.

Classification of operations

The alternative to "hardwiring" questions to answers is to adopt a classification scheme for both that permits questions to be mapped to answers of the same class. Logically, a classification scheme must be built on a reasoned understanding of the information involved. This follows the thinking of Booth (1989) and Jarke and Vassiliou (1985), who argue that establishing a mutual conceptual framework aids in arriving at a mutual understanding of the topics being discussed. Even in the very restrictive setting of a human/computer dialog, a common view of the information being treated seems fundamental. Various authors offer conceptual frameworks concerning the nature of geographic information and operations. None was designed to be used as a semantic basis for a human-computer dialog; however, each has merit.

The first framework considered was that of Tomlin (1983), who names three types of geographical modeling operations.

- The output value is a function of a point.
- The output value is a function of neighborhood or adjacency. Neighborhoods can be immediate, extended, or indeterminate (i.e., the neighborhood must be computed or estimated after the process is underway)
- The output value is a function of a vicinity or region.

These three types of operations would be multiplied in a temporal database, since each type of data could exist at a point in time or in a trajectory through time, and each modeling operation could consider a point in time or a trajectory through time.

A second prospective framework is that of Rucker (1987), who describes a mathematical treatment of reality that could serve as a basis for a model of geographic information. As stated by Rucker, the five archetypical patterns of mathematics are number, space, logic, infinity, and information. A geographical region can be used to illustrate these concepts.

- Number. A region contains a certain number of buildings and a certain number of wetlands. The buildings have a certain height and length that can be measured numerically; the wetlands have measurable moisture, animal populations, and acreage. The area of the region itself can be measured, as can its population.

- Space. A region is not flat, as it might appear on a map. It exists in four-dimensional space/time. It has no holes (as defined here), and it connects to other regions. It follows the curvature of the earth, its stream network branches in a one-dimensional pattern, and the earth's relief forms roughly conical bulges in three dimensions. The region has subregions, which might by accessed by White's (1984) windowing, buffering, boundary, and endpoint operators. Subregions also may intersect, coincide, or be included within one another.

- Logic. The subregions within a region are interconnected. The region also reacts to external changes. A dammed stream floods land upstream. A zoning change alters land use. Tax disparities between regions cause outmigration or unpredictable settlement. Interconnections and reactions also exist that are unknown or poorly understood.

- Infinity. By zooming out from the region, one might see that it forms a pattern with other regions. By zooming in, one might notice greater detail and apparent structure within that detail and cells, and then atoms come into focus. What meets the eye when examining a region relates closely to the scale at which the region is examined.

- Information. Over time, a region is subject to random influences that leave their mark. Rucker suggests two ways to measure the information content of an entity: by the number of questions required to build a replica of the entity, or as the length of the shortest computer program required to answer any possible question about it.

As one might expect, this mathematical model addresses conceptual units and measurements, which could be useful in selecting map formats. However, it is not linked to the components of geographic information and hence would need to be extended considerably to meet the needs expressed in this discussion.

A more promising framework is that of Sinton, whose 1978 work on spatial data representation is useful for organizing spatiotemporal information because it addresses all three components of spatial information: attribute, location, and time. Sinton argues that traditional representation methods can measure only one of these attributes. A second is fixed to a constant value, and the third is controlled to a range of values or a set of categories (Table 4).

Sinton starts with a map and classes it according to how the map treats the various components. If it were possible to start with a question and class it according to how the question treated each component, then Sinton's framework provides a likely starting point for addressing issues one and three because questions could be matched to the appropriate formats for answers.

Table 4. The representation of geographic data in various formats (extended from Sinton 1978).

	Fixed	**Controlled**	**Measured**
Soils data	time	attribute	location
Topographic map	time	attribute	location
Census data	time	location	attribute
Raster imagery	time	location	attribute
Weather station reports	location	time	attribute
Flood tables	location	time	attribute
Tide tables	location	attribute	time
Airline schedules	location	attribute	time
Moving objects	attribute	location	time

Sinton's framework revolves around the components of the information itself, rather than functions or measures alone. In addition, the framework is tied naturally to graphic forms, so an artificial language that is structured to express what components are fixed, controlled, or measured could also indicate what output form is appropriate. Returning to the goal of a dialog based on questions and answers, we can see that Sinton's framework can be mapped readily to question words.

Attributes: what, who, how many, how much

Location: where

Time: when, how long

Using that mapping, a tie between question and answer is already evident, since the "question word" indicates the measured variable.

AN ARTIFICIAL LANGUAGE FOR POSING GEOGRAPHIC QUESTIONS

A high-level artificial language built upon questions designed to elicit answers would alleviate the difficulty of using a GIS. Today's GISs force the user at the helm to express information needs using a combination of relational and spatial logic. Questions are not asked directly, so the system has few options for providing meaningful help sequences or display formats.

How it would work

The high-level language would be designed to sit atop of the attribute and spatial command languages that dictate data manipulations. High-level commands would call lower-level commands, just as a high-level programming language decomposes to machine language before functions are performed. This concept is similar to the methods adopted by Nyerges (1980) and Wu et al. (1989), although in those two implementations, the top-level query language did not disassociate the user from data-manipulation commands, as supported here. As demonstrated by Wu et al., the effect of the high-level layer on performance is negligible; the translation of question-to-command occurs before commands are executed, and a language can be designed for direct translation into commands.

The high-level language would be comprised of a limited set of verbs, nouns, and modifiers and a grammar to describe sequence within an expression. The system would parse the expression and select the appropriate data retrieval and manipulation commands to execute. At the same time, the sequence in which entities are mentioned would indicate which elements of the expression were fixed, controlled, or measured. That information, and identifying whether the data types are point, line, or area, would indicate what output format to select.

Although this study has not defined a query language that meets the goals listed above, certain patterns in the sequence of words in questions indicate a possible syntax. It is useful to note the natural syntax of English questions to ensure that an artificial language is truly "English-like" and easy to learn. Using the questions from Table 2 (above), Table 5 describes how natural English grammar contains clues concerning measured components, content desired, and window.

Table 5. Common English grammar used to frame questions gives a basic indication of the format required to answer. Questions are taken from the examples of Table 2. In all cases, the region of interest in "this watershed."

Measured component	**Primary subject**	**Attribute modifier**	**Relative location**	**Period of interest**
Where are	pastures			now
What is	agriculture	by type		now
Who is the	owner	of pastures	by stream	now
What is	grazing density	of pastures	by stream	now
Where does	stream		border pastures	now
What	pastures	changed from nonagriculture		past ten years
When did	pastures	change from nonagriculture		ever
How many	acres	changed from agriculture to nonagriculture		past ten years

To map from a question to an output format is not entirely straightforward, even when the natural-language version appears so easily parsed. Aside from the question "where...," which is answered most naturally via a map, most "question words" can be answered in either map or tabular format. Accordingly, a user should be able to select either a map or a listing as output unless one is patently inappropriate. Ideally, the system can recognize when only one ouput format will do. Use of the word "where" is one clue. Others also exist; the single question listed in Table 5 that requires a listing has a unit of measurement as its primary subject.

A second problem of selecting a map design by default is informing the computer of the data types involved. Attributes can use nominal, ordinal, interval, or ratio measures, and the mapped entities can be points, lines, areas, or surfaces. Each combination has a set of appropriate mapping techniques. The logic required to select an output format would follow the lines offered in Table 6. Once again, criteria for each decision will be collected from different parts of the question. The question word combined with the subject indicates the level of measurement. For example, a "where" question requires only nominal symbols to answer, but a "what" can require an ordinal, interval, or ratio answer depending on the subject. *What is the agriculture by type?* requires a different map response than *What is the grazing density?* A reasonable approach to automating this decision is to include the level of measurement for each data attribute in the data dictionary.

Table 6. Deducing an appropriate output format given a question to answer.

Select entities to appear on the map using the region and period of interest, the primary subject and attribute modifiers, and constraints on relative location.
Determine whether the primary subject is point, line, or area.
Determine whether the measure is nominal, ordinal, interval, or ratio.
Select format.

Although Table 5 decomposes the questions only to a coarse level of detail, more information exists in the natural expressions that help describe the desired manipulations and output format to the computer. For example, the verbs include tense, which indicate a temporal map and an excursion into the past-tense database. In addition, certain words (e.g., "border..." and "change...") are keywords that describe a buffering and a temporal operation, respectively.

Possible enhancements

Several measures are possible to move the dialog to yet a higher level of discourse. Specialized views or objects could be developed to express to the system combinations of data that have special meaning to the user. Using Table 5's example, the watershed could be stored as a named MBR in the database so it could be referenced by name and a window automatically selected. Alternately, a user could indicate the region of interest on the screen, as described in Egenhofer et al. (1989).

Other enhancements to a question-oriented query language could include using expert system techniques to shorten the instructions necessary to describe certain concepts to the computer. Two excellent examples of "mini" expert systems are found in the temporal information system literature. Overmyer and Stonebraker (1985) describe the development of a "time expert" within a relational database that permits the system to interpret such terms as "lunch," "today," incomplete dates, and ranges of dates. Kahn and Gorry (1977) describe a "time specialist" that interprets temporal constructs needed in problem solving. The Overmyer and Stonebraker system operates within an INGRES QUEL environment, while Kahn and Gorry's system is designed to interpret natural language expressions. Input method could equally well be query-by-example, graphics, or a fourth-generation language.

Additional bonuses

The focal purpose of developing a question-oriented query system is to relieve users of the task of specifying data-manipulation procedures to the computer, and to permit the computer to produce default displays automatically in response. The latter in particular is most effective if the computer understands the information sought.

If the human-computer dialog informed the computer of what information is sought by the user, conceivably the system could supply more useful help sequences. These would include model questions that the user could choose from and edit, and listings from the data dictionary. Incorrect syntax could be corrected through gentle feedback from the computer regarding correct options and the results they would produce.

Assuming that the high-level query language were built on the georelational architectures in common use today, the high-level question would be decomposed by the system into a series of SQL and spatial commands. One option for experienced users would be to list these commands and permit them to be edited for greater control over output. New users also could employ such a listing as a learning tool.

SUMMARY AND CONCLUSIONS

This discussion does not introduce a finished high-level query language. It does, however, supply a basis for continuing investigation into the topic. It argues the following.

- The next generation of GISs should relieve users of the burden of specifying data manipulations and output formats by permitting them to specify their information needs (rather than data-processing needs) in an English-like artificial language supplemented by point-and-click inputs.

- The design of any cartographic query language should be linked to the problem of how a query's response should be formatted. At present, systems permit users to design and save an output format, and reference that format as the desired output for a query. Arguably, this level of automation can be increased and system designers can relieve users of the burden of map specification altogether.

- Sinton's theory of geographic representation provides a possible starting point for developing a conceptual framework for questions and answers that the human and computer can share.

Continuing research in this area can take two parallel tracks by attacking the problem from its two ends: the query language and the output format. Development of a query language requires an investigation of the natural phrasing of geographic questions and how their elements decompose and map to data manipulations, and further examination of how Sinton's theory applies to geographic questions. Development of default output formats requires a better understanding of what formats best answer what questions, and how the components of the question indicate the composition of the format.

Many have argued for the importance of cartographers in the GIS discipline. Given the current state of human/GIS communication, however, one might wonder at how involved cartographers have been to date; the cartographer's purported interest in communication has not apparently improved the state of affairs for GIS users. Until now, much of the harshness of the GIS environment could be attributed to the immaturity of hardware and software systems. But the raw materials do exist now to better the lot of GIS users considerably; it is our understanding of geographic questions and answers that lags behind.

REFERENCES

Biskuo, Joachim; Räsch, Uwe; and Stieteling, Holger. 1990. "An Extension of SQL for Querying Graph Relations." *Computer Languages* 15, 2, 65-82.

Booth, Paul. 1989. *An Introduction to Human-Computer Interaction.* East Sussex: Lawrence Erlbaum Associates Ltd.

Borenstein, Nathaniel S. 1986. "Is English a Natural Language?" *Foundation for Human-Computer Communication,* K. Hopper and I. A. Newman, editors. North-Holland: Elsevier Science Publishers BV, 60-72.

Broekhuysen, Martin and Dutton, Geoffrey. 1983. "Conversations with Odyssey." *Proceedings of Auto-Carto 6,* Ottawa, November, Vol. 2, 15-24.

Egenhofer, Max J.; Frank, Andrew U.; and Hudson, Douglas L. 1989. "User Interfaces for Geographic Information Systems." Abstracts of the 14th World Conference of the International Cartographic Association, August, Budapest, 162.

Embley, David W. 1989. "NFQL: The Natural Forms Query Language." *ACM Transactions on Database Systems* 14, 2, 168-211.

Frank, André. 1982. "MAPQUERY: Data Base Query Language for Retrieval of Geometric Data and their Graphical Representation." *Computer Graphics* 16, 3, July, 199-207.

Guptill, Stephen C., editor. 1988. "A Process for Evaluating Geographic Information Systems." Federal Interagency Coordinating Committee on Digital Cartography, U.S. Geological Survey Open-File Report 88-105.

Honeycutt, Dale; Brooks, Kristina; Kimerling, Jon. 1980. "GIS: a review of selected operational and functional capabilities." Department of Geography, Oregon State University, January.

Hoppe, Heinz Ulrich. 1990. "A Grammar-Based Approach to Unifying Task-Oriented and System-Oriented Interface Descriptions." *Mental Models and Human-Computer Interaction*, D. Ackermann and M. J. Tauber, editors. North-Holland: Elsevier Science Publishers BV, 353-373.

Jarke, Matthias and Vassiliou, Yannis. 1985. "A Framework for Choosing a Database Query Language." *Computing Surveys* 17, 3, 313-340.

Jarke, Matthias. 1986. "Current Trends in Database Query Processing." *On Knowledge Base Management Systems*, Michael L. Brodie and John Mylopoulos, editors. New York: Springer-Verlag, 111-119.

Kahn, Kenneth and Gorry, G. Anthony. 1977. "Mechanizing Temporal Knowledge." *Artificial Intelligence* 9, 87-108.

Katzeff, Cecilia. 1989. "Strategies for Testing Hypotheses in Database Query Writing." *Man-Computer Interaction Research*, F. Klix, N. A. Strietz, Y. Waern, and H. Wandke, editors. North-Holland: Elsevier Science Publishers B.V., 125-147.

Mackaness, William and Fisher, P. F. 1987. "Automatic Recognition and Resolution of Spatial Conflicts in Cartographic Symbolization." Proceedings of Auto-Carto 8, Baltimore, 709-718.

McGranaghan, M. 1989. "Context-Free Recursive-Descent Parsing of Location-Descriptive Text." *Proceedings of Auto-Carto 9*, Baltimore, April, 580-587.

Morse, Bruce W. 1987. "Expert System Interface to a Geographic Information System." *Proceedings of Auto-Carto 8*, Baltimore, March, 535-541.

Muller, J.-C., Johnson, R. D., and Vanzella, L. R. 1986. "A Knowledge-Based Approach for Developing Cartographic Expertise." Proceedings of the Second International Symposium on Spatial Data Handling, Seattle, July, 557-571.

Nyerges, Timothy L. 1980. "Modeling the Structure of Cartographic Information for Query Processing." Unpublished PhD dissertation, Ohio State University, Columbus.

Nystuen, John D. "Identification of Some Fundamental Spatial Concepts." In *Spatial Analysis*, edited by Brian J. L. Berry and Duane F. Marble. Englewood Cliffs, NJ: Prentice-Hall Inc., 1968.

Overmyer, Ricky and Stonebraker, Michael. 1982. "Implementation of a Time Expert in a Database System." *SIGMOD Record* 12, 3, 51-59.

Robinson, G. and Jackson, M. 1985. "Expert Systems in Map Design." Proceedings of Auto-Carto 7, Baltimore, 430-439.

Schank, Roger C. and Rieger, Charles J. 1985. "Inference and the Computer Understanding of Natural Language." *Readings in Knowledge Representation*, Ronald J. Brachman and Hector J. Levesque, editors. Los Altos, California: Morgan Kaufmann Publishers, Inc., 119-140.

Plumb, Gergory A. 1988. "Displaying GIS Data Sets Using Cartographic Classification Techniques." Proceedings of GIS/LIS '88, San Antonio, December, Vol. 1, 340-349.

Quillian, M. Ross. 1985. "Word Concepts: a Theory and Simulation of some Basic Semantic Capabilities." *Readings in Knowledge Representation*, Ronald J. Brachman and Hector J. Levesque, editors. Los Altos, California: Morgan Kaufmann Publishers, Inc., 98-118.

Rucker, Rudy. 1987. *Mind Tools: the Five Levels of Mathematical Reality.* Boston: Houghton Mifflin Company.

Salmen, Larry J. 1978. "Natural Resource Analysis Techniques - Their Impact for DBMSs." *Harvard Papers on GIS,* Vol. 2, edited by G. Dutton. Reading, Massachusetts: Addison-Wesley.

Sinton, David. 1978. "The Inherent Structure of Information as a Constraint to Analysis: Mapped Thematic Data as a Case Study." *Harvard Papers on GIS,* Vol. 7, edited by G. Dutton. Reading, Massachusetts: Addison-Wesley.

Tomlin, Data. 1983. "Digital Cartographic Modeling Techniques in Environmental Planning." Unpublished PhD dissertation, Yale University.

Vasiliev, Irina. 1990. "Examples of the Treatment of Time as a Variable on Maps." Paper presented at the 1990 Conference of the Association of American Geographers.

Walshe, A. 1989. "Formal Methods of Database Language Design and Constraint Handlings." *Software Engineering Journal,* January, 15-24.

Webber, Bonnie Lynn. 1986a. "Natural Language Processing: A Survey." *On Knowledge Base Management Systems,* Michael L. Brodie and John Mylopoulos, editors. New York: Springer-Verlag, 353-364.

Webber, Bonnie Lynn. 1986b. "Questions, Answers, and Responses." *On Knowledge Base Management Systems,* Michael L. Brodie and John Mylopoulos, editors. New York: Springer-Verlag, 365-402.

Weibel, Robert and Buttenfield, Barbara P. 1988. "Map Design for Geographic Information Systems." Proceedings of GIS/LIS '88, San Antonio, December, Vol. 1, 350-359.

White, Marvin S. 1984. Technical Requirements and Standards for a Multipurpose Geographic Data System. *American Cartographer* 11, 1, 15-26.

Wu, Jian-Kang; Chen, Tao; and Yang, Li (1989). "A Versatile Query Language for a Knowledge-Based GIS." *International Journal of GIS* 3, 1, 51-57.