International Geographic Information Systems (IGIS) Symposium: The Research Agenda

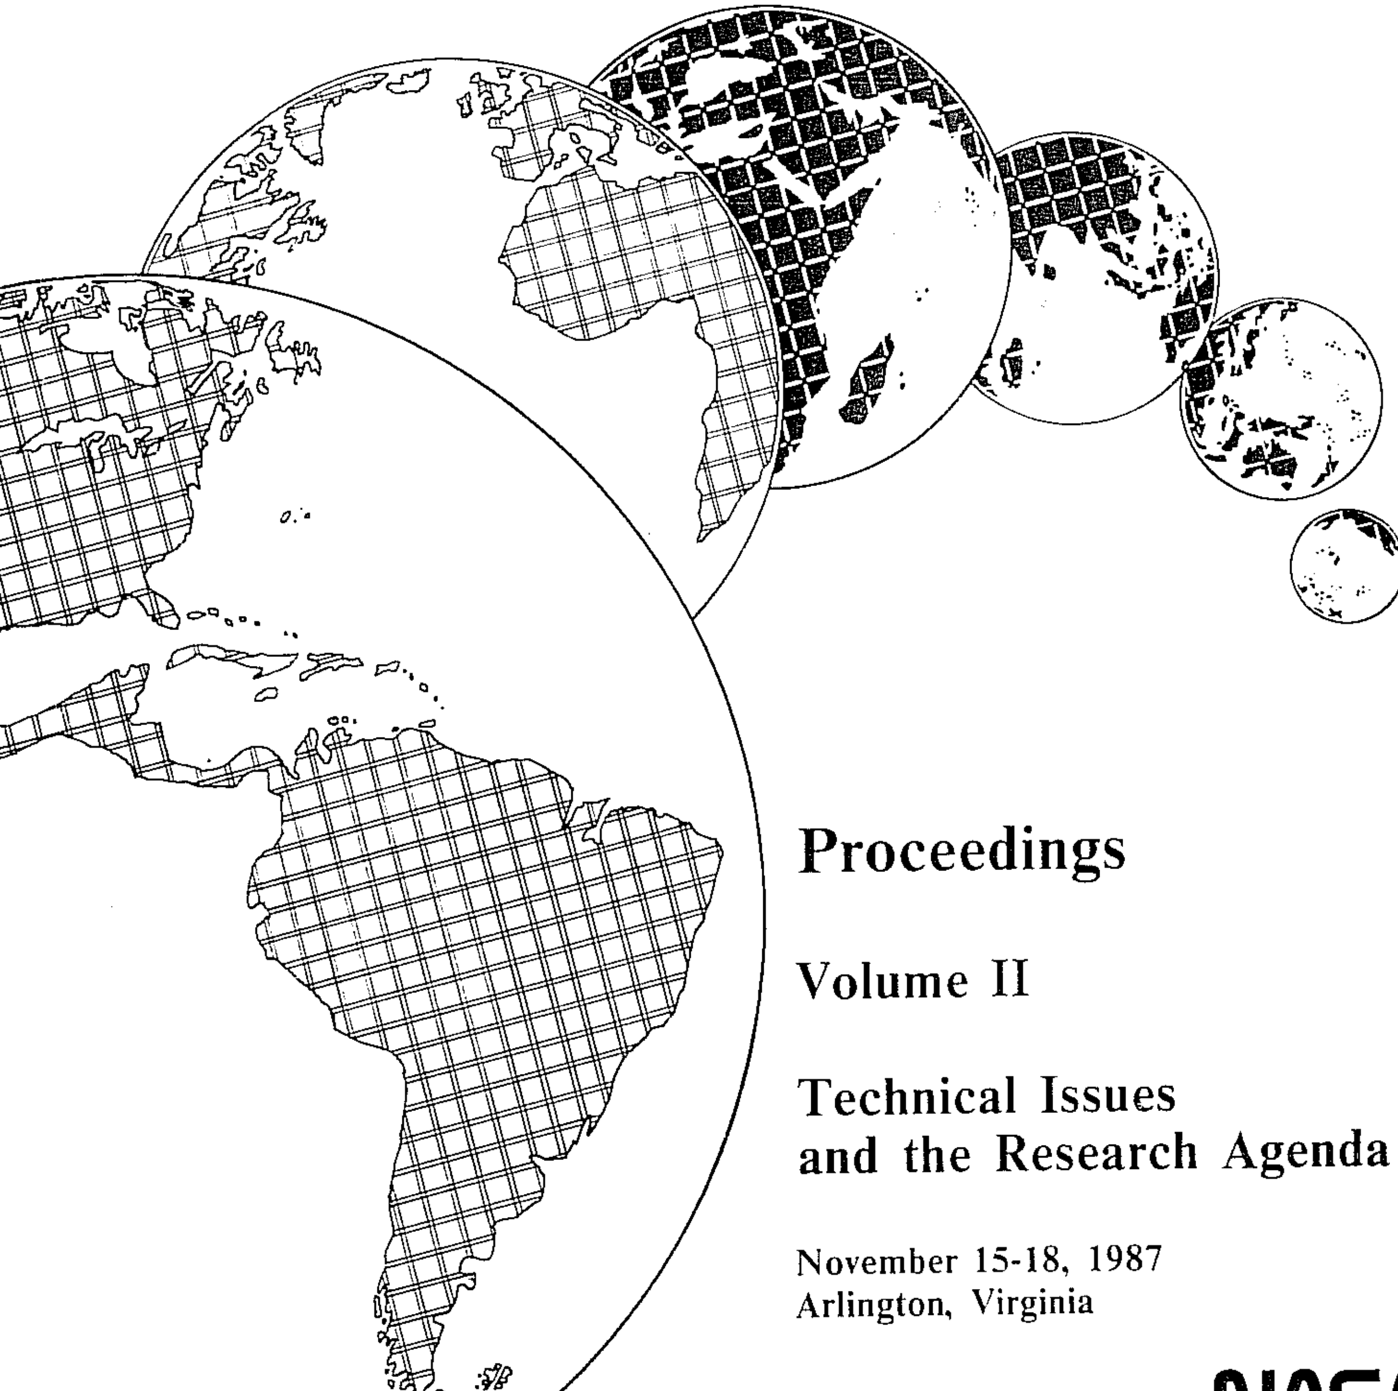

Proceedings

Volume II

Technical Issues and the Research Agenda

November 15-18, 1987
Arlington, Virginia

Artificial Intelligence Tools for GIS*

Andrew U. Frank
Douglas L. Hudson
Surveying Engineering Program
University of Maine, Orono, ME 04469, USA
FRANK@MECAN1.bitnet
HUDSON@MECAN1.bitnet

Vincent B. Robinson
Department of Surveying Engineering
University of Calgary, Calgary, Alberta T2N 1N4
VBR%SURVEY@UNCAEDU.bitnet

Abstract

The essential topics in Geographic Information Systems (GIS), spatial and temporal reasoning, are difficult and require powerful tools and a formal systems approach. Future GIS will include tools and methods developed in artificial intelligence research and will resemble today's expert systems. An expert system is defined here to be a computer-based mechanism that produces results comparable to those obtainable from a human expert in some field. Moreover, intelligent user interfaces and reasoning about data quality can also benefit from the power of expert systems.

1 Introduction

Future computer-based Geographic Information Systems (GIS) will employ methods and tools developed in artificial intelligence research and they will resemble what is now called an expert system. This will allow a GIS to provide great power and flexibility to a diverse user community.

This paper begins by discussing some of these methods and tools and their applicability for use in an 'intelligent' system. The assertion is made that the use of these methods and tools is indispensable for building an effective GIS.

*This work was partially funded by grants from NSF under No. IST-8609123 (principal investigator A. Frank) and No. IST-8412406 (principal investigator V. Robinson).

An intelligent systems paradigm is first introduced and is followed by some discussions concerning knowledge representation, search strategies and machine learning. Significant emphasis is placed on the use of expert systems which may assist the user in specific task areas such as geography, cartography, data quality, and spatial relations in general. The envisioned services of a GIS are so powerful that users will need some help in using them. The GIS itself is expected to provide that assistance as well.

Some of the techniques and methods discussed in this paper apply generally to any information system, but they are most necessary for spatial information systems given the additional complexity of geometrical data processing and graphical output. In this paper, therefore, we will use the terms *Spatial Information System* and *Geographical Information System* without intending to differentiate between them.

2 AI Tools and Techniques

Artificial Intelligence (AI) research can be segmented into specialty areas concentrating in such subject matters as natural language processing, intelligent data retrieval, expert consulting systems, theorem proving, robotics, automatic programming, combinatorial and scheduling problems, and perception [Nilsson 1980]. Most of these are also areas of concern in the development of a GIS. Therefore, the tools produced to assist in the solution of AI problems are often directly applicable to GIS implementations, e.g. efficient search strategies, heuristic control of processes, interesting and efficient schemes for knowledge representation, etc. Some of these tools will be discussed in this section.

2.1 Intelligent Systems Paradigm

One of the most productive ideas to come from studies in AI is that of maintaining a definite separation between knowledge and the mechanisms which operate upon and control the processing of knowledge. This separation is not usually found in programs written in procedural computer languages such as Pascal or Fortran; in these, the results are obtained by describing the processes by which the results may be obtained. In an AI system results are obtained by simply stating what knowledge is required, leaving it up to the control mechanism of the system to infer the results in the best way it can. For this reason these control mechanisms are often called *inference engines*.

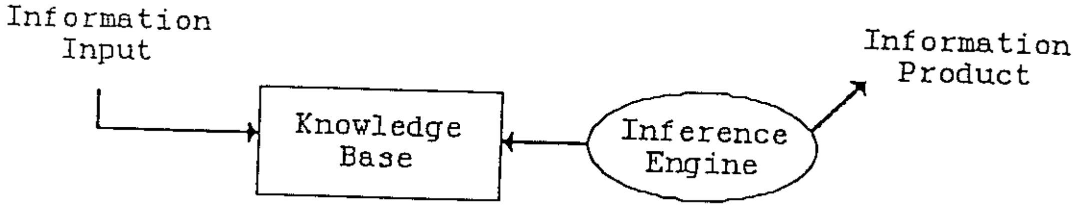

Separate Knowledge from its Control

Strict independence is not always easy (or sometimes even possible) to maintain. The advantage of this architecture, however, is that more powerful tools for operation or control can be developed independent of any specific knowledge domain. Moreover, new or modified knowledge domains can be constructed and applied to an existing inference engine to solve a variety of problems. Changes in the knowledge base (which are to be expected in any system which attempts to model reality) do not require corresponding changes to the general control machinery. More importantly from the user's view, it is possible to update the knowledge base without having to reprogram or even fully understand the internal operations of the AI system. Ideally the user only needs to understand a simpler model of the behaviour of the system, because inference engines are generally based upon a simple mathematical model describing very fundamental relations. It is usually not necessary for a user to comprehend the complexity of technical operations or to know the tricks-of-the-trade which are incorporated into the inference engine to enhance its performance. These latter concerns should be invisible to the user, freeing him or her to express relevant knowledge in a purer form.

Functional independence has obvious benefits to a GIS implementation. Geographic information is rarely static and, therefore, must undergo constant revision to remain current and accurate. This applies to factual aspects, e.g. that San Jose is currently the capital of Costa Rica, as well as to the rules of geography, e.g. that continental climates in temperate zones are characterized by hot summers and cold winters. Personnel more familiar with geography than with computer programming may modify the information component with relative ease if it is detached from programming concerns and if it can be formulated in familiar terminology. Complex programming problems which apply only to the inference engine can be addressed by people with the proper background in computer science. This partitioning allows each specialist to remain more productively in his or her own realm of expertise.

2.2 Knowledge Representation

Divorcing the knowledge base from the inferencing mechanism means that knowledge representation schemes are generally declarative in nature; it is only necessary to state what something is, not how to make one. Any procedural decisions are performed by the inference engine. The knowledge base is a collection of specific and general facts and descriptions of relations between them (e.g. Augusta is the capital of Maine; the area of Lot 6 is 2.45 acres; if the population of X is between 1000 and 25000 then call it a town; etc.). Most geographic information is of this kind. The format used for the conceptual representation of knowledge is very important as it determines what inference stategies can be used.

2.2.1 Logic

There is great interest in the use of logic for knowledge representation purposes. The first order predicate calculus (which is primarily declarative) easily lends itself to knowledge-based systems and there are readily available programming languages, such as Prolog and Lobster [Frank 1984], designed to work directly with the predicate calculus.

There are on-going discussions in the AI community concerning the adequacy and suitability of a first order logic for general knowledge representation [Winograd 1975] [Kowalski 1979] [Minsky 1981]. Many of the early arguments against it were guided by the perception that the predicate calculus is cumbersome and computationally inefficient [Newell 1981], but many of these problems can be overcome as more is learned about the controls that can be placed on the deductive process [Moore 1982]. Many other problems, however, are not yet satisfactorily resolved, for example those which concern time, uncertain or unknown values.

To resolve some of these questions multi-valued logics, those with more than just true or false as values, are being examined. Some interesting possibilities are: 3-valued logics, for instance using values TRUE, FALSE, and MAYBE, or TRUE, FALSE and UNKNOWN; probability logic in which a valid response may be that something has a .345 chance of being true; and Fuzzy Logic where values are determined according to interpretations of classification schemes.

2.2.2 Meta-level Knowledge

Meta-knowledge is knowledge about the organization of, or relationships among, other pieces of knowledge. There are meta-levels on top of meta-levels as well; no limits should be placed upon the levels of construction for meta-knowledge. Computational control can be based upon meta-level knowledge and it is imperative that meta-knowledge play an important role in the operation of an AI system. It is the only way in which the separation of knowledge base and inference engine can be maintained in non-trivial situations and it provides operational efficiencies the lack of which would make even small systems computationally intractable.

To incorporate meta-level knowledge it is important that the system provide support for basic knowledge abstraction. Procedures for classification, generalization, and aggregation need to be part of the system's design. Classification is the recognition that there is representative structure for individual objects, actions or concepts that share common properties; generalization is the meta-level extension of classification; aggregation is the recognition of the collective as an entity in its own right. Without these abstraction concepts the meta-knowledge will tend to be haphazard and difficult to comprehend and to use correctly.

2.3 Search Strategies and Heuristic Control

The database of information for a GIS is expected to be large and the mechanisms for accessing that information need to be powerful and efficient. Many of the early successes in AI research were developments of search techniques that make more efficient data retrieval possible and that can guide a search to optimal or near-optimal results. In general terms, a solution in an information system is found by searching through the database until all the relevant information has been found. Obviously the strategy of sequential search is only usable as long as the database is relatively small. For realistically large databases, more appropriate search stategies are required.

If all the possible solutions and intermediate steps to solutions of a problem are considered as states of that problem, the states can be visualized as a state-space or search space. By recognizing certain structuring possibilities in the state-space various search stategies can be advantageously employed. Often, if the state-space can be arranged hierarchically or as a graph, one of two basic search techniques, either forward or backward chaining, is used.

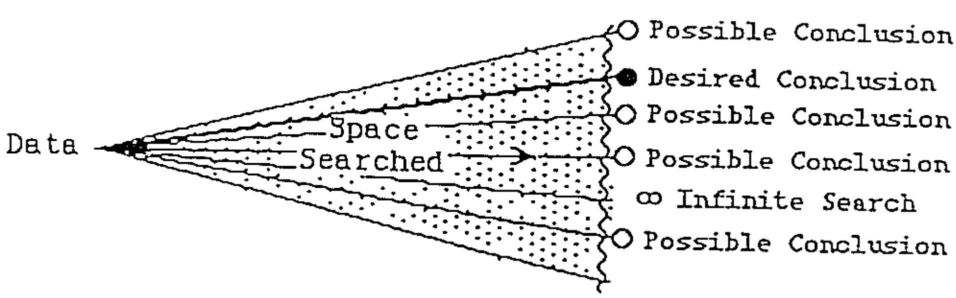

Forward chaining search space

The former technique starts with some facts and rules and then searches for interesting conclusions which can be derived from them. It is sometimes called a data-driven search. If there are many branches in the search space, this technique may investigate many unproductive trails before some success is met. If the space is infinitely large, the search algorithm is undecidable and may never meet with success.

With backward chaining, a conclusion is tentatively proposed and an attempt is made to show that the conclusion can be derived from facts and rules stored in the database. The search space in backward chaining is usually, but not necessarily, finite and is considered to be more easily directed because the search is made from a known goal to known data values.

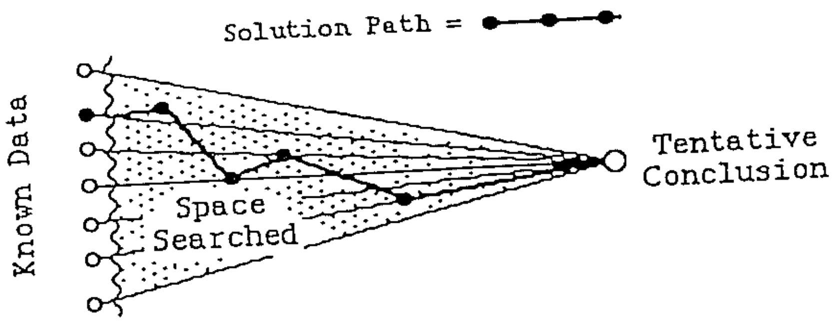

Backward chaining search space.

It may be possible to combine the two search methods to gain considerable efficiency. This is called a bi-directional search. The intersection of the two search spaces can be significantly smaller than with either technique alone.

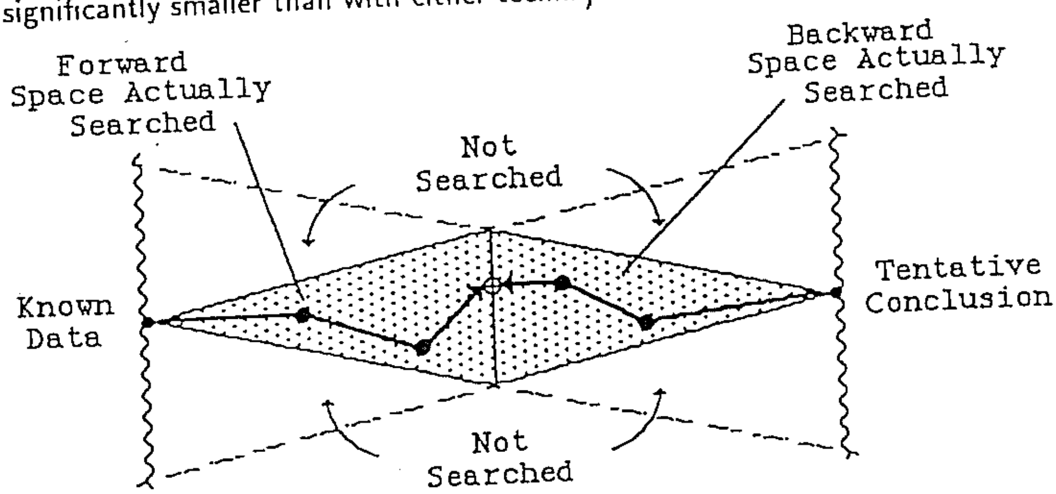

Bi-direction search may meet half way.

2.3.1 Heuristic Searches

In many instances the search directions can be controlled or assisted by additional information concerning the problem domain. Branch and bound algorithms, such as A* and AO* [Nilsson 1980], can often find optimal solutions without a full search of all those possible. Ocassionally heuristic methods are used to find solutions which may not be optimal. They may still be satisfactory, however, and can be obtained at costs far below more complete search methods.

A simple example of heuristic search control would be to cancel a request to search a long list of active military personnel for those who have reached the age of 150 years because there should be no values returned in this case. More complex examples can be found in chess playing programs where, from the many hundreds of moves that are possible, a particular one is selected based upon informed interpretation of current board positions and strengths.

Any information which will help to limit unproductive search, whether it is highly specific to the problem, a general rule-of-thumb, or just plain common sense, can be used as a heuristic. This technique has its disadvantages because there is often the chance that the best advice available turns out to be bad advice. For instance, in the case of a bi-directional search, heuristic control may cause the search paths to by-pass one another resulting in twice the work to find a solution.

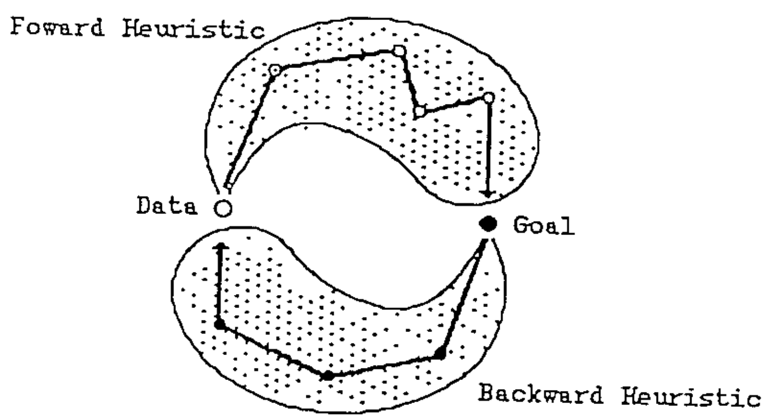

Bi-directional heuristic searches may take twice the work of a single direction heuristic search.

2.3.2 Machine Learning

One of the most interesting areas of AI research concerns machine learning. It is a continuing theme that computer programs contain only the intelligence that the programmer and the knowledge engineer program into them; however, there are a number of examples of programs that have the ability to learn on their own. Example problems or constructions of increasing complexity are presented to these programs. The programs search for common properties and relationships among data entities, possibly comparing the new information with knowledge previous acquired. The new information is then integrated with the old, increasing the knowledge base.

Most of the questions in this area are still unanswered; however, there are a few initial efforts in this area concerning GIS. One of these is SRAS [Robinson 1987], a system concerned primarily with the acquisition of *fuzzy* natural language concepts of spatial relations. The KBGIS programs [Smith 1984] [Smith 1986] [Glick 1985] claim to have learning capabilities as well. Since the knowledge acquistion techniques currently used to create intelligent systems require the employment of highly trained (and expensive) human knowledge engineers, machine learning is a field that has enormous potential for growth as the recognition of the utility of expert systems increases.

3 Expert Systems

The most useful product of AI research is, and will for a long time continue to be, the concept of the expert system: a computer-based mechanism that produces results comparable to those obtainable from a human competent in a particular field of interest. Although a strict definition of an expert system is elusive, one should try to include in any attempt a sense of flexibility in the use of the knowledge under its control.

As with human experts, expert systems can be found with different levels of expertise. A modest system such as one which can determine simple relationships family relationships can be built in an hour or so and can be an important pedigogical tool for aspiring knowledge engineers. More sophisticated systems attempt such tasks as medical diagnoses, spectral analyses, and combat situation evaluations; these can take many people-years to construct.

Expert systems incorporate many of the tools introduced above to provide an appropriate response in situations where human experts are too costly or are not readily available. The expert system can provide advice in the acquisition and integration of primary data, assistance in optimal data storage, guidance in the appropriate operations to perform upon data, and direction in the ways and means of retrieving either primary data or conclusions that can be drawn from the data. Moreover, the expert system can assist in the analyses and control the presentation of results as well.

Current expert systems, however, are characterized by rather narrow knowledge domains. Their expertise is generally very specialized to the task at hand. In order to accomplish the entire range of tasks mentioned in the previous paragraph a collection of systems would probably be required, each one specializing in one area. Perhaps another system could serve as an expert manager for the others.

3.1 Types of Expert Systems

There are a number of organizational schemes for expert systems. Most current expert systems are based upon a type of IF_THEN rule structure. It goes by various names, such as a situation-action or a production rule system. The general technique in all of these rule-based systems is that the recognition of one or more pre-condition states somehow triggers a transformation to a post-condition state. "If we have A, B, and C then we can get to D." Rule-based systems can be implemented in a variety of ways: large case statement structures, list-pattern matching, or as theories in logic which employ resolution to deduce results [Hayes-Roth 1985].

The concept of a knowledge frame [Minsky 1981] employs a data structure into which directly pertinent knowledge about some object or event is concentrated. The frame localizes information on one subject, and provides slots for default values and an interface to other related knowledge frames. For instance a frame could be constructed for a restaurant. Information about seating and the functions of cooking and serving would be included along with some default values like typical menu entries (e.g. hamburgers and fries). A frame for a different restaurant may have similar cooking and serving functions, but the values for seating could be increased and the default menu entries overridden (e.g. by fried rice and egg rolls).

The use of associative (or semantic) networks is another approach [Findler 1979]. Knowledge is arranged into a directed graph structure such that arcs represent relations between various pieces of knowledge which are stored at the graph nodes. The associative network concept may be combined with the rule-based approach if the rules for traversing the network are stored at the nodes.

At this time working with frames or networks is considered difficult because of the problems inherent in standardizing the knowledge structures to gain working efficiencies; however, there is a growing interest in these forms as some of the underlying problems are gradually resolved [Gevarter 1983].

4 Expert Systems Essential For GIS

Any future GIS will necessarily include techniques developed for use in expert systems. Expert systems are of such great utility that we assume that this is a common movement for all forms of database development and to refer to an "expert GIS" would be redundant; however, as mentioned above, currently expert systems have very narrow scopes of interest.

A GIS must be quite expansive in its knowledge domain to provide competent and useful information on geography, cartography, and fundamental spatial relationships.

The wide scope of expertise required for a GIS is too complex for current technologies to handle, however; a GIS as an expert system is, at this time, only a dream. Even a system to handle cartographic output requires knowledge that has not yet been successfully organized.

Despite the size of the problem, however, it is assumed that a comprehensive expert GIS will someday exist. All of the common definitions of a GIS imply that power. When the problems are eventually solved and a system is constructed, it will provide assistance to a diverse user group including private citizens, resource managers, highway and utility planners, and legal and governmental services. To accomplish its mission, a GIS will need to integrate a number of specialized knowledge domains or possibly combine a number of expert systems into one cohesive unit.

4.1 Spatial Organization Experts

A GIS is a spatial information system and it is, therefore, important to consider how best to represent spatial relationships among objects. Humans place special emphasis on the geometric properties of things; they want to know if two objects are near to one another (and usually how near) or if they are touching at a point or along a border line, etc. Perhaps they overlap, intersect, or penetrate. There is an enormous amount of information in even a simple graphic image and a GIS should have access to it all. Unfortunately no one has yet invented a single method for the computer representation of the spatial aspects of objects in a way that is entirely satisfactory. An important research effort is to develop a formal system for reasoning about spatial relations.

The two major technologies that have evolved in an attempt to treat this problem are raster and vector. Raster is concerned with the direct organization of space and vector is concerned with the organization of objects in space. The differences are not easily reconcilable at the technical level; however, they both address the same problem. Some operations are easier in one system than the other. It is assumed that future systems will find a way to accommodate both raster and vector, but it will probably take another expert system to decide which form or combination of forms is better suited to a particular solution.

4.2 Cartography Experts

Expert system technology will be crucial to the construction of systems that are to produce cartographic output on demand. Maps that are mass-produced and stocked for subsequent distribution must make compromises in their information content because their audience can not be narrowly defined. The major advantage of a spatial information system over current mapping efforts is that the spatial information system is intended to produce maps that closely match the specific needs of the individual user.

The graphic quality of a map produced by a GIS must remain at an acceptable level, although not necessarily as good as a human cartographer's effort. To accomplish this requires that rather sophisticated decisions be made as to map form and content. Some progress in this area has been made concerning the automatic placement of names of cities and towns, etc. such that overlaps and other conflicts are avoided [Ahn 1984] [Freeman 1984]. At least two levels of problems still remain, however. One is the proper selection of map entities to be displayed and the other is the proper selection and placement of representative symbols for those entities.

A user about to go fishing may ask for a map of all public boat ramps in Maine, but a map that only shows boat ramps will be of little use unless the sea coast and lakes and perhaps a few connecting roads are also shown. Choosing what to display as context for the primary information request requires expert intervention.

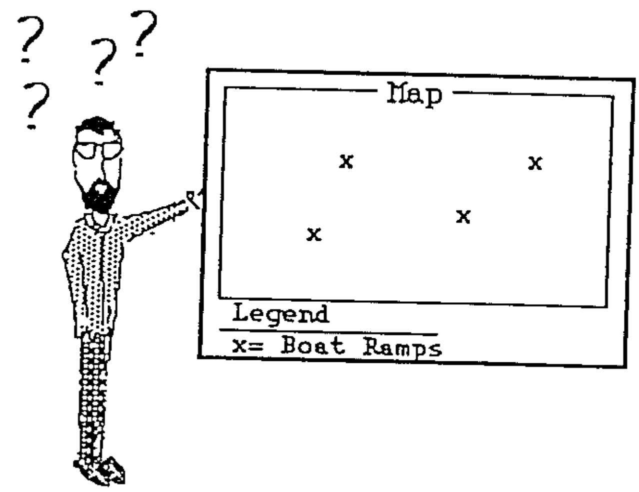

At some point during changes in scale a single detailed symbolic representation of some entity or a closely grouped collection of symbols must be either be eliminated or replaced by a simplified or generalized symbol. The entity may be a street map of a town which is changed to a plain colored area showing only major thoroughfares or which is replaced entirely by a small circle, the size of which is proportional to the town's population. The tightly grouped collection of symbols may be small squares representing houses along a street which may simply be eliminated altogether at larger scales. Deciding when and how to modify a symbol in these instances also requires some cartographic expertise.

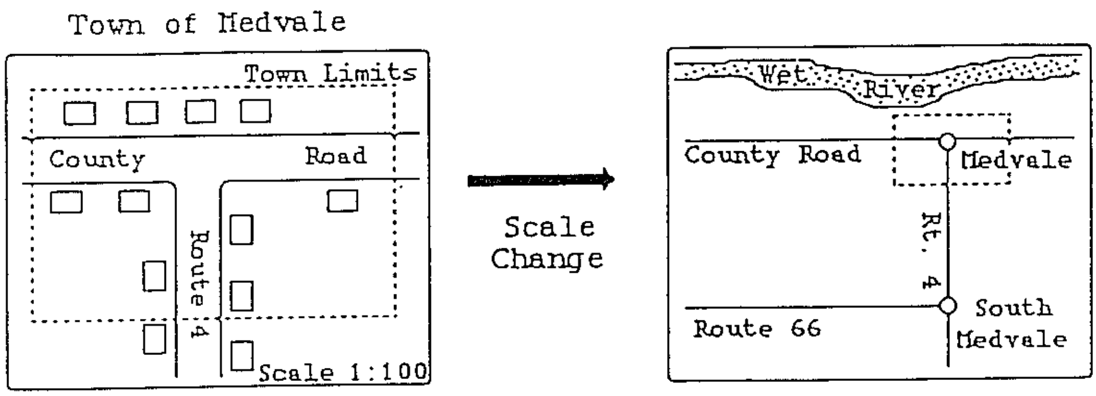

Scale changes cause symbol changes.

4.3 User Interface Experts

Users must be able to express their needs clearly and simply, in a way that does not require a deep understanding of the technology of the GIS. They should not have to write programs to extract needed results. A system should accept a description of the information required and find ways on its own to produce it with minimal further input. An expert system must make allowances for imprecise information requests and provide assistance in guiding a search for specific information.

In order for this to happen, the system must provide a language that both user and system can tolerate. This language must permit expressions of spatial relationships, such as inside, next-to, overlapping, as well as other, more conventional, queries. For instance, a request to a GIS might be to identify all of the spruce and pine forests within 15 miles of navigable rivers in Washington County, Maine, or determine that piece of land which is zoned as light industry and which is nearest to the town of Orono. It is preferable, but not absolutely necessary, that the language be close to a natural language in form and style.

To respond best to a user's request, a system needs some internal model of the user and his or her current focus. Human responses are somewhat adapted to the perceived audience and a GIS should have similar capabilities. Is the user interested in just an overview or some specific, detailed piece of information? Does he or she want an analysis of the results? Is the user interested in highway construction, utility planning or environmental issues? Does the user have extensive experience in this subject or does he or she need help in finding information or with understanding the information? Is this user color blind? Terminology, verbosity and help levels should all be variables in a GIS-user interface. Progress has been slow in this area [Friedell 1982] and much more work needs to be done.

4.4 Data Quality Experts

The quality of information product of a GIS is intimately associated with its reliability in decision-making. To make effective use of a GIS a user should know what the level of confidence to place in the results of the system. The expression "garbage in – garbage out" stills holds, but it is also possible to get garbage out even with a respectable input. An information system is a data processing system and all non-trivial computer transformations introduce errors which may compound as they propagate through the system [Hudson 1987]. As an expert in geographic information the system itself may be used to track data quality and provide the following:

- a listing on demand of the statistical confidence level of any piece of data displayed.
- warnings that the transformations employed have compounded errors to the point that the data requested is now useless.
- warnings against the inappropriate use of certain operations. The arithmetic mean of Portland and Bangor should not be computed.
- guidance in the proper collection and quality requirements of primary geographic data such as a suggestion that to obtain an output of a particular precision A requires that the primary input data be at least of precision B.

5 Conclusion

The work done by AI researchers to find general purpose solutions to the problems of getting a computer driven system to act intelligently can often be adapted successfully to the more specialized requirements of a Geographic Information System. Particular methods and tools developed in the field of AI, such as heuristically controlling searches, maintaining a separation between the representation of knowledge and the control of its processing, and giving the GIS the ability to learn in order to acquire new knowledge on its own, are indespensible in the construction of a viable GIS. The incorporation of expert systems is necessary because of the size and complexity of the problems expected to be encountered in a GIS. Ordinary procedurally oriented programming techniques will be insufficient to accomplish the task.

Acknowledgement

Thanks to Jeff Jackson who made the illustrations.

References

[Ahn 1984] J.K. Ahn. Automatic Map Name Placement. Technical Report IPL-TR-063, Image Processing Laboratory, Rensselaer Polytechnic Institute, Troy (NY), 1984.

[Friedell 1982] M. Friedell et al. Context-Sensitive, Graphic Presentation of Information. In: D. Bergeron, editor, Proceedings SIGGRAPH '82, ACM Computer Graphics, Boston (MA), July 1982.

[Freeman 1984] H. Freeman and J. Ahn. AUTONAP—An Expert System for Automatic Map Name Placement. In: International Symposium on Spatial Data Handling, Zurich (Switzerland), August 1984.

[Findler 1979] N.V. Findler. Associative Networks. Academic Press Inc., Orlando (FL), 1979.

[Frank 1984] A. Frank. Extending a Database with Prolog. In: L. Kerschberg, editor, Proceedings of the First International Workshop on Expert Database Systems, Kiawah Island (SC), October 1984.

[Glick 1985] B. Glick et al. Hybrid Knowledge Representation for a Geographic INformation System. In: Seventh International Symposium on Computer-Assisted Cartography, Washington (DC), 1985.

[Gevarter 1983] W.B. Gevarter. An Overview of Artificial Intelligence and Robotics. Technical Report, NASA Technical Memorandum 85836, Washington (DC), 1983.

[Hayes-Roth 1985] F. Hayes-Roth. Rule-Based Systems. Communications of the ACM, 28(9), September 1985.

[Hudson 1987] D. Hudson. Data Quality in GIS. Technical Report, Surveying Engineering Program, University of Maine, Orono (ME), 1987.

[Kowalski 1979] R. Kowalski. Logic for Problem Solving. Elsevier Science Publishing Co., New York (NY), 1979.

[Minsky 1981] M. Minsky. A Framework for Representing Knowledge. In: J. Haugeland, editor, Mind Design, MIT Press, Cambridge (MA), 1981.

[Moore 1982] R.C. Moore. The Role of Logic in Knowledge Representation and Commonsense Reasoning. In: AAAI-82, Pittsburgh (PA), 1982.

[Newell 1981] A. Newell. The Knowledge Level. AI Magazine, 2(2), 1981.

[Nilsson 1980] N.J. Nilsson. Principles of Artificial Intelligence. Tioga Publishing, Palo Alto (CA), 1980.

[Robinson 1987] V.B. Robinson and R.N. Wong. Acquiring Approximate Representations of Some Spatial Relations. In: N.R. Chrisman, editor, Eighth International Symposium on Computer-Assisted Cartography, Baltimore (MD), 1987.

[Smith 1986] T. Smith et al. KBGIS-II: A Knowledge-Based Geographic Information System. Technical Report TRCS86-13, Dept. of Computer Science, University of California, Santa Barbara (CA), May 1986.

[Smith 1984] T. Smith and M. Pazner. Knowledge-Based Control of Search and Learning in a Large-Scale GIS. In: International Symposium on Spatial Data Handling, Zurich (Switzerland), August 1984.

[Winograd 1975] T. Winograd. Frame Representations and the Declarative/Procedural Controversy. In: D.G.Brobow and A.M. Collins, editors, Representation and Understanding, Academic Press, New York (NY), 1975.