Intelligent Interactive Dynamic Maps

Suguru Ishizaki and Ishantha Lokuge

Visible Language Workshop, Media Laboratory
Massachusetts Institute of Technology
20 Ames St, Cambridge, MA 02139

ABSTRACT

This paper presents an experimental intelligent map system—GeoSpace—which allows information seekers to explore complex spaces of geographic information using dialogue-like interaction. GeoSpace progressively and selectively provides information as an information seeker enters queries while visually maintaining the larger context. Domain knowledge is represented in a form of information presentation plan modules, and an activation spreading network technique is used to determine the relevance of information based on information seeking queries. The reactive nature of the activation spreading network, combined with visual design techniques, such as typography, color, and transparency, enables the system to fluidly respond to the continuous changes in the information seeker's goals and intentions.

INTRODUCTION

The exploration of complex geographic data spaces in an age where both technology and information are growing at exponential rates is a challenging task. Recent developments in interactive computers with high-quality visual displays have provided the GIS designer an opportunity to create more comprehensive environments for presenting complex information. However, most of existing systems fail to support an information seeker's continuous exploration of information and gradual construction of understanding. In other words, although they provide highly sophisticated functionality and displays, they do not relate one presentation to another in response to an information seeker's goals and intentions.

We have applied an activation spreading network technique as a representation scheme of domain knowledge, along with an abstraction of information seeking goals and presentation plans, in order to provide interactive maps which embody the following characteristics:

Continuity: We assume that an information seeker's goals are achieved through a series of input queries, or a *dialogue*. The system should be able to consider previous queries as well as to anticipate the forthcoming queries in order to respond to a current input query. We adopted dialogue as a fundamental model of interaction in our system.

Fluid response: When interaction is taking place in a form of dialogue, the system should be able to generate a map, or visual response, in a fluid manner—*as if it was a continuous conversation*. We have used an activation spreading network to create a responsive display which allows the map display to progressively respond to a series of information queries over time.

Visual clarity: Dialogue-based interaction allows the system to limit the range of information that need to be displayed simultaneously, so that the display becomes visually clarified and highly comprehensive, as opposed to a dense display which contains many information elements. GeoSpace uses an activation spreading network to determine the levels of importance for each information element so that an appropriate set of information elements can be chosen to be more visually dominant than the others.

Context preservation: In addition to the visual clarity requirement, we also must visually preserve the larger context in a map, so that an information seeker does not get lost during interaction. GeoSpace achieves this by using various degrees of transparency, type size, and color according to visual design rules.

Two main areas of research have influenced the work presented in this paper. The first area of research involves visual techniques and direct manipulation as a means of exploring complex information space. One such approach is the use of overlapping multiple layers of information in which individual layers are accessible (e.g., Belge 1993, Colby 1991). Most multi-layer approaches provide users with an interface that controls the display based on layers and regions in order to visually simplify the map display. However, this type of interaction becomes cumbersome when the volume of information is large, or when the information seeker does not have prior knowledge about a particular geographic database.

While the above approaches emphasize direct manipulation and visual techniques, other interface displays have been proposed that incorporate domain and presentation knowledge (e.g., Feiner 1993, Maybury 1993, Roth 1993). Maybury introduces an interactive visual presentation method that considers visual presentation as communicative acts (e.g., graphical, auditory, or gestural) based on the linguistic study of speech acts (Maybury 1993). A multimedia explanation is achieved by using rhetorical acts, which is a sequence of linguistic or graphical acts that achieve a certain communicative goal such as identifying an information entity. Rhetorical acts are represented in a form of a plan, which is similar to our representation. Although the system introduced by Maybury enables sophisticated presentation based on a user's single query, it does not have a mechanism to maintain a model of the user's information seeking goals from one query to another.

In this paper, we propose a software architecture for creating intelligent and responsive geographic display systems that allows an information seeker to incrementally asks questions in order to gradually achieve his/her information seeking goals. In the following sections, we first outline the basic functionality of GeoSpace using a simple interaction example. Then, we present a technical framework for implementing the software architecture of GeoSpace. Finally, we discuss potential directions in which our research can be extended.

A TYPICAL SCENARIO

Most GIS users often find it difficult to formulate their information seeking goals in one request. Hence, we believe that an information display that gradually augments this process would greatly enhance the user's comprehension.

As a consequence, we have used the following simple scenario of conversation between an information seeker (IS) and information provider (IP) as an interaction model for GeoSpace. The first query by the IS makes the IP guess what is important to show. After the IP provides information based on the first query, the IS may ask the second query based on what is provided. The IP then determines what is important to show next considering both the first and the second queries. The information seeking dialogue may continue until the IS is satisfied.

An IS's information seeking process can be top-down, bottom-up, or a combination of both. For example, imagine a situation where an IS is trying to locate a new apartment. The IS may start a dialogue by stating that s/he is looking for an apartment. This can be considered top-down since the IS provided the ultimate goal of the dialogue. In this case, the IP is not certain about what kind of detailed information the IS is aware of. On the other hand, the IS may ask for a particular location (e.g., "Where is Cambridge?"). This can be considered bottom-up since, it targets a specific item of data. In this case, the IP is not certain about what the IS's ultimate goal is.

Figure 1. Map of Boston area showing the dense nature of the display.

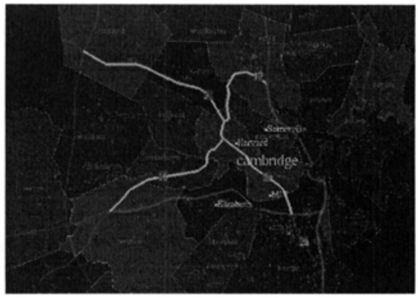

Figure 2."Show me Cambridge."

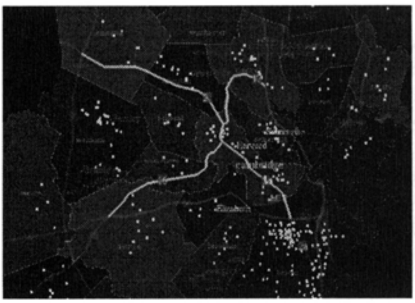

Figure 3. "Show me crime distribution."

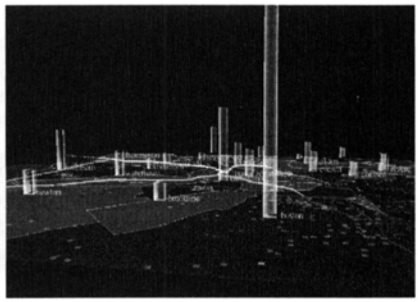

Figure 4. "Show me crime statistics."

In GeoSpace, we consider the IP to be an expert in both domain information and visual presentation and the IP's knowledge is canonicalized in a form of reactive patterns. Instead of deliberately reasoning about what to present every time the user asks a question, the IP simply reacts to it by using canonical presentation techniques.

Based on this scenario, we have developed: (1) a knowledge representation scheme for representing domain knowledge together with visual design knowledge, (2) a computational mechanism whereby the system reacts to a series of user requests (i.e., information seeking goals) while maintaining overall context.

There have been interface approaches to interactive map systems which use queries coupled with graphical displays both for narrowing down information to be presented (e.g., Ahlberg and Shneiderman 1994, Goldstein and Roth 1994) and for supporting users' exploration of the data space (Egenhofer 1990). The information seeking scenario used in GeoSpace emphasizes the latter in its purpose.

The rest of this section presents a simple interaction example of GeoSpace in order to introduce its basic functionality which was developed based on the scenario described above. Figure 1 shows a snapshot of the initial state of the display. The visual complexity of this map display makes it hard for users to discern specific information while interacting without getting lost. The following interaction examples illustrate how GeoSpace dynamically generate a map display according to a series of user queries. Imagine a person new to the Boston area tries to explore the information around the area so as to look for a place to live. Having heard of the perilous life styles of people in Boston, suppose that the person is interested in crime distribution statistics and accessibility to hospitals in the neighborhood.

First, the IS asks the system "Show me Cambridge." Then, the type size of the text label Cambridge and its opacity value gradually increases resulting in a sharper focus of Cambridge (Figure 2). Notice also that related information, such as hospitals, highways around Cambridge became visually prominent, but to a lesser degree compared to the label for Cambridge. This first query exemplifies the reactive nature of discerning visually complex map display into a relatively simple and comprehensible design.

The power of using an activation spreading network to control the visual dynamics is exemplified in Figure 3, where the user requests to see crime distribution following the previous query. This shows a spatial distribution of crime data for the greater Boston area, while maintaining the Cambridge context from the previous query. Now, crime date represented by a collection of small read dots is most prominent in the display, while information related to Cambridge has become secondary but it can be still distinguished from the rest of the data. If the IS asks the relational statistics of the crime data instead of its distribution, the user can obtain a three dimensional view of crime data in the form of a bar graph as shown in Figure 4.

GeoSpace is implemented in C++ and GL graphics language on a Silicon Graphics Onyx workstation with Reality Engine graphics.

DOMAIN KNOWLEDGE

Information seeking goals and presentation plans are the basic components of this approach. A plan consists of a list of sub-plans, a list of conflicting plans, and a list of effects. The effect-list contains a set of goals that are achieved by executing the presentation plan. The sub-plan list contains a set of goals that must be achieved in order to accomplish goals in an effect-list. The conflict list contains a set of goals that are either semantically irrelevant or visually conflicting with the plan. Knowledge about semantic conflicts helps the system to identify a shift of interest. When large amount of data exist in a database, it is often the case that same visual features (such as color, typeface, orientation, or motion) are used by more than one visual element. Knowledge about visual conflicts helps the system to identify visually confusing situations.

Figure 5 shows a typical presentation plan. The plan (a) says, in order for a user to know about transportation, a user must know about bus routes, subways, and place names. The plan also indicate that hospitals and bookstores are not relevant when a user wants to know about transportation. In the current knowledge representation, semantic and visual conflicts are not distinguished. Plan (b) is much simpler; it has neither sub-plans nor conflicts. The activation level specifies the threshold energy required to realize the plan.

Plan:	**{Show_Transportation}**	
Sub-Plans:	**{Know_Place_names, Know_Bus_routes, Know_Subways}**	
Conflicts:	**{Know_Hospitals, Know_Bookstores}**	
Effects:	**{Know_Transportation}**	
Realization:	**ø**	
Activation:	**0.8**	(a)
Plan:	**{Show_Bus_map}**	
Sub-Plans:	**ø**	
Conflicts:	**ø**	
Effects:	**{Know_Transportation}**	
Realization:	**#<bus_map-object>**	
Activation:	**0.3**	(b)

Figure 5. Typical presentation plan.

The domain knowledge-base is independent of the geographic information databases. Figure 6 shows the relationship between a database and the domain knowledge-base, which can be created either manually by a cartographic designer or automatically by the system. In the current implementation of GeoSpace, the domain knowledge is encoded manually by the designer. In the future implementation, we expect to partially automate this process by developing rules that can infer relationships among information in a database. Spatial proximity relationships, information types and visual design principles will provide a criteria for creating the rules.

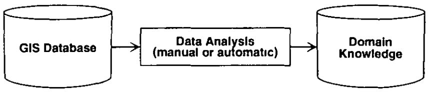

Figure 6. A process of generating the domain knowledge.

ACTIVATION SPREADING NETWORK

The system uses an activation spreading network (Anderson 1983, Maes 1990) to determine priorities of plans based on the user's request. The activation spreading network can be viewed as a graph in which plans are viewed as the vertices, and sub-goals, conflicts, and effects are viewed as the edges. A plan module's activation level is changed by the user's immediate goals, and when their activation level exceed the threshold, positive and negative activation energy is sent to other plan modules connected by hierarchical links and conflicting links respectively. The current system iteratively injects a constant amount of energy to fluidly change the overall activation state. In every iteration, activation levels of all the plan modules are normalized to the most active plan. This also results in the gradual decay of plans whose links are not explicitly specified. In addition, a presentation plan can also be activated by the dynamic changes of information it is representing. For instance, a presentation plan for a highway section can be activated by the dynamic changes of traffic information (assuming it is available).

When the user specifies a query such as "Show me transportation", *know_transportation* becomes the current information seeking goal. The system then injects activation energy to the plans that contain *know_transportation* in the effect-list. When a plan module's activation level reaches a certain threshold, it spreads energy to the plans which contain the sub-goals in their effect list. A plan also spreads activation energy upwards to the higher level plans whose effect-list contains *know_transportation* as sub-goals. This upwards activation results in activating indirectly related information. Figure 6 shows a simple example of an activation spreading process. Every iterative activation spreading

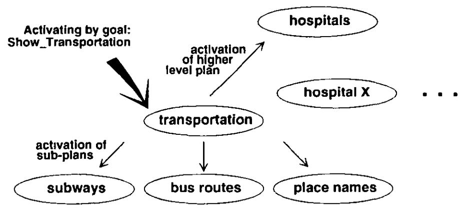

Figure 6. Schematic diagram of typical activation spreading.

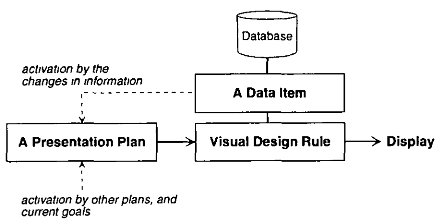

Figure 7. Shows the method for selecting graphical style of data item.

change of the display. Figure 7 shows a schematic diagram of how an activation level of a presentation plan are used to determine the graphical style of a data item.

An activation spreading network not only maintains the immediately relevant information, but it can also preserve the history of a user's exploration process. When a user requests new information, the system seamlessly transforms the previous state into the new state. The network can also prepare for the user's future request by activating plan modules that are potentially relevant in the following interactions. This could greatly assist users to formulate subsequent queries towards satisfying a particular goal.

VISUAL DESIGN

The map display involves many layers of information each of which corresponds to a different set of data. The system is intended to incorporate various visual techniques, such as translucency and focus which helps clarify visual information without loosing overall context. We have incorporated these new techniques along with traditional graphical design techniques in the design of the map display. Most important information is displayed with a higher level of opacity, and related information is displayed with medium translucency. Irrelevant information is displayed almost transparent. Since, the display can show related information using relative transparency, the user has a chance of realizing a new question to ask next. Also, previously displayed information can be shown with medium to high transparency so that the user can maintain a continuos dialogue.

Plans may or may not have a graphical presentation. For example, a plan to show highways does not have a graphical representation, but each highway has a graphical representation. Those plans that have a graphical representation change their graphical style according to their activation levels. Currently, the energy levels are scaled and mapped to transparency values and/or typographic sizes on the cartographic display. The mapping from the activation levels to graphical styles is achieved by simple procedures that are implemented according to design principles. In other words, visual design knowledge is embedded in those procedures and presentation plans. Thus, the quality of visual presentation, such as legibility, readability, and clarity are significantly enhanced.

Intelligent dynamic maps guide user' attention to regions in the display that are important in a fluid manner. The mechanism described above can implicitly chain presentation plans by hierarchically spreading activation energy, and can respond to an immediate shift of interest by spreading negative energy to conflicting plans. This spreading of energy can be driven by temporal information such as weather and traffic data which makes the display truly dynamic. In such cases it is critical that users are focusing their attention at the relevant regions of the map display to comprehend the data being presented. This is accomplished by the visual techniques described previously. The result is an intelligent and highly reactive cartographic display.

CONCLUSION AND FUTURE DIRECTIONS

We have presented an intelligent and responsive map display for interactively exploring complex geographic data space. We have shown that the knowledge representation scheme which uses the activation spreading network, along with an abstraction of information seeking goals and presentation plans, provides the map display with a reactive capability. The mechanism can chain presentation plans by hierarchically spreading activation energy, and can respond to an immediate shift of interest by spreading negative energy to conflicting plans. The system can also maintain the context of a continuous dialogue in a fluid manner by gradually changing the states of activation. Dynamic use of various visual techniques, such as translucency, type size and color, are associated with activation levels of plans in order to visually maintain overall context during a course of information seeking dialogue.

Having completed the first generation of GeoSpace presented in this paper, we have begun to develop the next generation in order to further enhance its functionality. The following issues are currently being investigated: First, GeoSpace currently uses relatively small amount of data in order for us to carefully examine the behavior of the activation spreading network. We are in a process of increasing the size of database so as to further examine the potential of this technique. Second, in the first generation, the domain knowledge is built manually by a designer. The next generation will include a graphical interface for building domain knowledge, and a mechanism that automatically constructs the initial domain knowledge base for certain types of geographic information. Third, the activation network is relatively sensitive to the amount of activation energy spread. We are continuing to experiment with varying energy levels to find the optimal network configuration. Fourth, we are experimenting with the use of weighted links (Maes 1992) to support varying degrees of relationship among presentation plans. Finally, the second generation includes a learning mechanism which allow a particular IS to customize the domain knowledge base. Both an explicit learning and implicit learning mechanisms are being experimented.

ACKNOWLEDGEMENTS

We would like to thank Prof. Muriel Cooper, Ron MacNeil, and Dave Small for their continued support and advice at all times during the course of this project. Special thanks to Prof. Whitman Richards, Louis Weitzman, Hasitha Wickramasinghe, and Yin Yin Wong for commenting on the earlier drafts of this manuscript. We would also like to thank the other members of the MIT Media Laboratory's Visible Language Workshop for providing valuable suggestions as we developed our ideas.

This work was in part sponsored by ARPA, JNIDS, NYNEX and Alenia.However, the views and conclusions expressed here are those of the authors and do not necessarily represent that of the sponsors.

REFERENCES

Anderson, John. A Spreading Activation Theory of Memory. Journal of Verbal Learning and Verbal Behavior 22, pp.261-295, 1983.

Ahlberg, Christopher and Shneiderman, Ben. Visual Information Seeking: Tight Coupling of Dynamic Query Filters with Starfield Displays. Proceedings of SIGCHI Human Factors in Computing Systems, 1994.

Belge, Matt, Lokuge, Ishantha and Rivers, Dave. Back to the Future: A Graphical Layering System Inspired by Transparent Paper. INTERCHI'94 Conference Companion, 1993.

Carberry, Sandra. Plan Recognition in Natural Language Dialogue, A Bradford Book, 1990.

Colby, Grace and Scholl, Laura. Transparency and Blur as Selective Cues for Complex Visual Information. International Society for Optical Engineering Proceedings, Vol 1460, 1991.

Egenhofer, Max J. Manipulating the Graphical Representation of Query Results in Geographic Information Systems, Proceedings of the IEEE Workshop on Visual Languages, 1990.

Feiner, Steven and McKeown, Kathleen. Automating the Generation of Coordinated Multimedia Explanations. In: Intelligent Multimedia Interfaces, ed. Mark T. Maybury, AAAI Press/The MIT Press, 1993.

Goldstein, Jade and Roth, Steven. Using Aggregation and Dynamic Queries for Exploring Large Data Sets. Proceedings of SIGCHI Human Factors in Computing Systems, 1994.

Maes, Pattie. Situated Agents Can Have Goals, Designing Autonomous Agents: Theory and Practice from Biology to Enginering and Back, ed. P. Maes, MIT Press/Bradford Books, 1990.

Maes, Pattie. Learning Behavior Networks from Experience, Towards a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life, ed. F.J. Varela & P. Bourgine, MIT Press/Bradford Books, 1992.

Maybury, Mark. Planning Multimedia Explanations Using Communicative Acts. In: Intelligent Multimedia Interfaces, ed. Mark T. Maybury, AAAI Press/The MIT Press, 1993.

Roth, Steven and Hefley, William. Intelligent Multimedia Presentation Systems: Research and Principles. In: Intelligent Multimedia Interfaces, edited by Mark T. Maybury, AAAI Press/The MIT Press, 1993.

Small, David, Ishizaki, Suguru., and Cooper, Muriel. Typographic Space. SIGCHI Conference Companion, 1994.