Evaluating the Usability of a Tool for Visualizing the Uncertainty of the Future Global Water Balance

Terry A. Slocum, Daniel C. Cliburn, Johannes J. Feddema, James R. Miller

ABSTRACT: We describe the development of software that is intended to enable decision makers (and their scientific advisors) to visualize uncertainties associated with the future global water balance. This is an important task because the future water balance is a function of numerous factors that are not precisely known, including the historical climatology, the model of potential evapotranspiration, the soil water holding capacity, and the global circulation models (GCMs) used to predict the effect of increased CO_2 in the atmosphere. In developing the software, we utilized the principles of usability engineering. In our case, we utilized six steps: prototype development, evaluation by domain experts, software revision, evaluation by usability experts, software revision, and evaluation by decision makers. Although this approach led to an improved piece of software, decision makers should have been involved earlier in the software design process, possibly at step two (instead of the domain experts). Decision makers found the notion of uncertainty discomforting, but their positive comments regarding the software suggest that it could prove beneficial, especially with improvements in spatial and temporal resolution. One interesting characteristic of our approach was the utilization of a wall-size display measuring 25 x 6 feet. The wall-size display engendered great interest, but determining whether it is truly effective will require a study that directly compares it with more traditional approaches.

KEYWORDS: Usability, uncertainty, decision-making, visualization, wall-size display

Introduction

Imagine that you are a U.S. senator trying to decide whether you wish to support the President's position on the Kyoto Protocol, which dealt with imposing a limit on CO_2 emissions in order to slow global warming. An important factor in making your decision is recognizing that global warming will, in general, lead to less water available at the surface, which could be detrimental to agriculture. Unfortunately, this simple finding will not apply in all situations, as some geographic regions are expected to warm, while others may actually cool. Moreover, the various global circulation models (GCMs) used to simulate the effect of increased CO_2 can produce dramatically different estimates of temperature and precipitation changes, which in turn can lead to substantial differences in estimated water availability—hence there is considerable uncertainty in determining the water availability at any particular geographic location (IPCC 2001).

The purpose of this paper is to describe our initial efforts to develop and test a visualization software tool that will enable decision makers (and their scientific advisors) to appreciate the uncertainty involved in water resources issues such as these. Because decision-making often requires input from multiple parties, our long-term goal is to develop a collaborative spatial decision-making (CSDM) environment in which individual decision makers can manipulate the display. As a first step in this direction, we developed the present software for a wrap-around wall-size display that enables a group of approximately 10 users to collaborate with one another when visualizing water resources issues. To evaluate the effectiveness of our software, we utilized the principles of usability engineering (Nielsen 1993; Mayhew 1999). The basic notion behind usability engineering is that software should not only be user friendly, but that it should respond satisfactorily to the tasks users expect of it. Thus, our goal was to

Terry Slocum and **Johannes Feddema** are professors in the Department of Geography, University of Kansas, Lindley Hall, 1475 Jayhawk Blvd., Room 213, Lawrence, KS 66045-7613. E-mail: <t-slocum@ku.edu>; <feddema@ku.edu>. **Daniel Cliburn** is a professor in the Department of Mathematics and Computer Science, Hanover College, PO Box 890, Hanover, IN 47243-0890. E-mail: <cliburn@hanover.edu>. **James Miller** is a professor in the Department of Electrical Engineering and Computer Science, University of Kansas, 2001 Eaton Hall, 1520 W. 15 St., Lawrence, KS 66045-7621. E-mail: <jrmiller@ku.edu>.

Cartography and Geographic Information Science, Vol. 30, No. 4, 2003, pp. 299-317

develop software that decision makers could actually utilize.

An edited volume by Medyckyj-Scott and Hearnshaw (1993) was the first to promote the use of usability engineering in geography. Although the technique suggested great promise, few studies in geography have actually employed usability engineering. A notable exception is Buttenfield's (1999) Alexandria Digital Library Project in which usability engineering played a major role. More recently, in a special issue of *Cartography and Geographic Information Science* focusing on "Research Challenges in Geovisualization," Slocum et al. (2001) have argued that developing "…effective geovisualization methods requires a two-pronged effort: theory-driven cognitive research and evaluation of methods via usability engineering principles." One purpose of the present paper is to respond to Slocum et al.'s call by beginning to experiment with usability engineering principles in the context of geovisualization in decision-making.

The remainder of our paper is split into four sections. First, we consider the notion of utilizing wall-size displays and the key characteristics of our display. As we will see, the use of a wall-size display is relatively unique within the geographic community. Second, we consider the nature of water balance models that are used to evaluate water availability on the Earth's surface. A basic understanding of the parameters involved in such models is essential to understanding the software tool that we will describe. Third, we summarize the methods that we utilized to visualize water availability and its associated uncertainty. We will see that our methods focus on 3D displays,[1] taking advantage of the processing power of the Silicon Graphics workstations used in association with the wall-size display. Fourth, we discuss our usability engineering approach, a variant of an approach suggested by Gabbard et al. (1999). Finally, we discuss the results of our usability engineering work and make suggestions for future research.

Utilizing a Wall-Size Display

Florence et al. (1997) were the first to suggest the potential of wall-size displays for geographic applications. Although they did not actually develop any software, they noted several novel capabilities that sophisticated wall-size systems might provide, such as the ability to query through gesture and voice. More recently, several groups of researchers have started to experiment with wall-size displays (the July/August 2000 issue of *IEEE Computer Graphics and Applications* was dedicated to the topic; also see Guimbretière et al. 2001), although most of this work has been outside the field of geography. One common characteristic of such efforts is the potential of wall-size displays for allowing users to interact more naturally with the display, as Florence et al. (1997) originally suggested. We have not yet utilized such novel capabilities with our display (at present, only keyboard and mouse input is supported), but we are developing methods that will enable individual users to manipulate the display via personal input devices (PIDs) and thus permit fuller collaboration.

The wall-size display that we employed measured 25 x 6 feet, covered 120 degrees of the visual field, and provided a 5760 x 1200 pixel resolution. No specialized apparatus was necessary to view the display, as is necessary with the CAVE, a room-like structure for creating a virtual environment (e.g., Wheless et al. 1996). Nevertheless, the display gave one a sense of being immersed in the visualization.

The wall-size display was driven by three SGI InfiniteReality2 graphics subsystems. One advantage of such processing capability is that it permitted us to make heavy use of 3D visualization in portraying water availability surfaces and associated uncertainty. A related characteristic is that the system allowed us to rotate three-dimensional images (consisting of several hundred thousand triangles) in real-time, a capability that was not available on standard desktop PCs when we were developing the software.

The Water Balance Model and Uncertainty

To determine water availability at each geographic location on the Earth's surface, we used a water balance model based on the Thornthwaite-Mather approach (Feddema 1998). Inputs to the model include monthly average temperature and precipitation estimates (we will refer to these as historical climatologies) and a soil water-holding capacity value. An additional important parameter of the model is the sub-model for calculating potential evapotranspiration: Thornthwaite (1948) and Hamon (1966). Uncertainty occurs when one of the inputs (historical climatologies or soil water-holding capacity) or the sub-model for potential evapotranspiration is varied. For example, in

---

[1] Technically, many of the visualizations we describe would be termed 2½D in the parlance of Slocum (1999, p. 19). We use the term 3D to stress the three-dimensional aspect of the visualizations and the fact that such displays are fully 3D from the perspective of the graphics hardware and software.

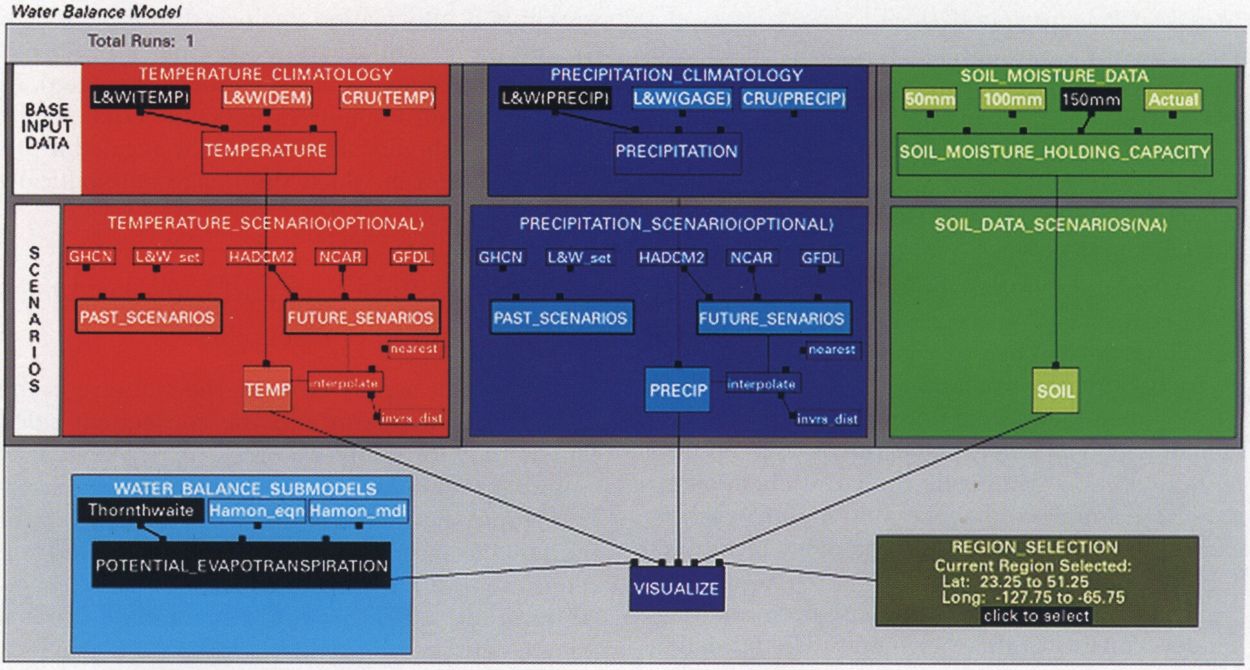

Current Region Selected:
Lat: 23.25 to 51.25
Long: -127.75 to -65.75

Figure 1. The visual programming interface for the software.

the case of temperature, a dataset developed by Legates and Willmott (1990) is based on 24,941 stations in which the number of years of observation varies from 10 to over 100, and Shepard's (1968) algorithm is used to interpolate weather station data to a one-half by one-half degree grid. In contrast, a dataset developed by the Climate Research Unit (CRU) at the University of East Anglia (New et al. 1999) is based on 12,092 stations in which the *same* 30 years of data (1961-1990) are used as input for each station, and a thinplate spline method is used for interpolation. If we run the water balance model with both of these temperature datasets, we obtain two different answers at each geographic location. How do we know which of these is the better answer, or if one is indeed better than the other? This is one of the reasons we are exploring modeling and visualization of uncertainty in our work.

To simulate future water availability, we utilize the methodology of Feddema (1999), which adds GCM climatologies (predicted changes in temperature and precipitation) to the historical climatologies before running the water balance model. For this experimental system, only three GCM climatologies were implemented: Hadley Center (HADCM2), the U.S. National Center for Atmospheric Research (NCAR), and the Geophysical Fluid Dynamics Laboratory (GFDL) (IPCC 2001). To determine a measure of uncertainty, we repeatedly run the water balance model using the different GCM climatologies. For example, it is interesting to contrast the model results for Kansas using the NCAR and the Hadley climatologies (IPCC 2001). The NCAR model suggests that there will be extreme desiccation in Kansas, while the Hadley model suggests wetter conditions during the month of April (which is critical for wheat production).

Visualizing Water Availability and Associated Uncertainty

In this section, we consider the range of approaches that were used to visualize water availability and its associated uncertainty. The images that are shown result from the software as it existed after step five of the usability engineering process (see next section).

Major Windows Used

Conceptually, we split the wall-size display into two parts: the left-hand one-third was utilized as a visual programming window in which users specified input for the water balance model, while the right-hand two-thirds were used to visualize the results of running the model. The one-third, two-thirds split was governed by our desire to have as much space as possible for visualizations, but also to avoid color and lighting differences within a particular display (three separate overhead projectors were used to generate the display).

The Visual Programming Window

The visual programming window is shown in Figure 1. We designed this in a fashion similar to other visual programming environments such as IBM's Data Explorer. The top row of boxes is where the historical climatologies and soil moisture information are specified. The second row of boxes is where the information on GCM climatologies is entered. Note that the GCM climatologies required users to select an interpolation option (the "interpolate" box in Figure 1). This was necessary because the resolution of the historical and GCM climatologies were not the same. Historical climatology data are available on a 0.5° longitude by 0.5° latitude grid, while the GCM climatology data have a much coarser resolution; for instance, each grid cell of the GFDL climatology measured 7.5° longitude by 4.5° latitude.[2]

The bottom row of the visual programming window consists of three elements. The left-hand element allows the user to specify the sub-model for potential evapotranspiration: Thornthwaite or two versions of the Hamon approach. The right-hand element is a region-selection tool. When this tool is clicked, a map of the world appears within which users can delineate any rectangular region. The middle element is a button used to trigger the evaluation of the model and visualization of its results for the currently selected input parameters and region of interest.

The Visualization Area

The right-hand two-thirds of the wall-size display were used to present visualizations of results, as shown in Figure 2. Note that it contains pull-down menus along the top that specify the general form of visualization (e.g., depicting the uncertainty of historical climatology data sets as opposed to the uncertainty of GCM results) and a set of option buttons on the left that specify how the visualization would appear (e.g., orthogonal versus perspective views or a view from the North versus a view from the South). In addition to manipulating the image with the option buttons, users can rotate and drag the image with a three-button mouse. We (and users) found that the ability to rotate three-dimensional images is critical in order to interpret them. Obviously, rotation is essential if information is blocked in a 3D image, but even when information is not blocked, images seemed to "come to life" when they are rotated. In fact, one of our usability experts noted that ideally our system should include what he termed a precession option in which the image automatically jiggled slightly.[3]

Notion of a Base Run and Calculating Uncertainty

An important aspect of using the software is the notion of a base run and the consequent calculation of uncertainty. A base run is created by running the water balance model once, with only one input selected for each of the basic input parameters—precipitation, temperature, and soil moisture—and the sub-model for potential evapotranspiration. The net result of the base run is a single water surplus/deficit value at each grid location associated with the land surface of the Earth. Once a base run is selected, multiple inputs can be selected for each of the basic input parameters, and the model can be run again, once for each additional input selected, with other parameters held constant based on the base run inputs.

To calculate a measure of uncertainty, we compute the *range* of surplus or deficit values modeled at each location. For example, if three temperature input datasets produces values of –20 mm (a deficit), 10 mm (a surplus), and 20 mm (a surplus), the uncertainty is 20 - (-20) or 40 mm. This simplistic approach works well given the limited number of data sets that can produce uncertainty. Our usability engineering confirmed the effectiveness of the approach: all groups tested emphasized the need to make things clearly understandable for decision makers.

In determining uncertainty associated with GCMs, we chose to examine all combinations of GCMs selected by the user. Thus, if all three GCMs are selected for temperature and precipitation, the result is nine possible water surplus/deficit values for each geographic location that can be related to the base run value at each location. The average of these values provides an indication of likely future surpluses/deficits, while the range of values again would indicate uncertainty.

---

[2] In examining Figure 1, note that GCM simulations of climate change apply to climate data only. In the future, we also intend to provide alternative human change simulations of soil degradation and land-cover change, as suggested in the blank area for SOIL-DATA_SCENARIOS.

[3] This finding concurs with that of Hibbard and Santek (1989, p. 56) who noted: "To enhance the depth information of our videotape animations, we usually combine time animation with a slight rocking of the three-dimensional scene."

Visualizing the Base Run

By default, the base run was visualized as a redundant symbology, with water surpluses represented by an elevated surface that is dark blue and water deficits by a depressed surface that is dark red (Figure 2). Although a 3D display is not essential for depicting this concept, we wanted our users to begin working with 3D images early in the visualization process because many of the methods for depicting uncertainty take advantage of 3D space. Blue and red are utilized to depict surplus and deficit moisture values because of many people's association of blue with surpluses and red with deficits (all of our users appeared comfortable with this association) and because blue/red is considered an effective diverging color scheme (Brewer 1996).

In addition to the smooth surface shown in Figure 2, we also permitted users to utilize either a prism-based or mesh (fishnet) surface (Figure 3). Prisms are appropriate in the sense that data input and calculations are done on a grid cell basis. We will see that the mesh surface is desirable when we wish to portray uncertainty in addition to the base run.

Visualizing Average GCM Scenarios

To display the average of GCM scenarios, we colored the surface for the base run with an orange-purple diverging scheme (Figure 4A). We chose this scheme because it looks fundamentally different from the red-blue scheme used to portray the base run,[4] and because it is an effective diverging color scheme (Brewer 1996). We chose the orange end of the scheme to represent less water available (compared to the base run) because we thought users would be apt to associate orange with a drier condition. A user could see both the nature of the base run (high and low points indicating surpluses and deficits) and the impact of the GCM (orange and purple colors indicating drier or wetter future conditions, respectively) by manipulating the resulting image in 3D space.

Visualizing Uncertainty

Previous work in visualizing uncertainty has focused largely on 2D displays, making use of Bertin's visual variables and modifications thereof (e.g., MacEachren 1992; McGranaghan 1993; Howard and MacEachren 1996). We instead wanted to experiment with 3D displays, partly because we had a Silicon Graphics computer environment that could readily create them, and partly because we thought that they might be useful when viewed in the environment of the wall-size display. We recognized, however, that many have questioned the usefulness of 3D graphics, when 2D graphics may be sufficient (e.g., DiBiase et al. 1994, pp. 292-297).

Others who have attempted to visualize uncertainty using 3D displays include Mitas et al. (1997), Pang et al. (1997), and Clarke et al. (1999). Given our utilization of glyphs to display elements of uncertainty, our approach parallels the efforts of Pang et al. (1997). What is unusual about our software is the ability to visualize multiple sources of uncertainty (e.g., to see that separate uncertainties are possible when evaluating water availability for the present and future, and to implicitly identify the source of those uncertainties).

Visualizing Uncertainty of Basic Input Parameters

Gershon (1998) argues that information pertaining to objects in a visual scene (in our case, uncertainty information) can be presented via two basic approaches: intrinsic and extrinsic. Intrinsic approaches vary an object's appearance, while extrinsic approaches rely on additional geometry to portray information about objects. We utilized both intrinsic and extrinsic approaches to depict uncertainty associated with the basic input parameters of the water balance model. An intrinsic approach was the *RGB method*, in which the base surface was shaded by utilizing mixtures of red, green, and blue, respectively, to represent the uncertainty associated with three of the input parameters: temperature, precipitation, and soil moisture.[5] The notion is that uncertainty in just one of these three parameters would be indicated by the respective color, while various combinations of uncertainty could be seen through mixtures of red, green, and blue. While color indicates the source of uncertainty, brightness is used to represent the relative magnitude of uncertainty, where black areas indicate no observed uncertainty. For example, in Figure 5 we see uncertainties due to soil moisture (the greens in the northwest), temperature (the reds in the north-central region), and precipitation (the blues just to the east of the area of temperature uncertainty).

One issue with which we struggled was how to represent the RGB method in a legend. Since each of our uncertainty values fell in the 0 to 1 range

---

[4] Here we are using Monmonier's (1992, pp. 249) notion of "signature hues."

[5] Byron (1994) also developed an RGB approach, but his method was intended for three variables that add to 100 percent.

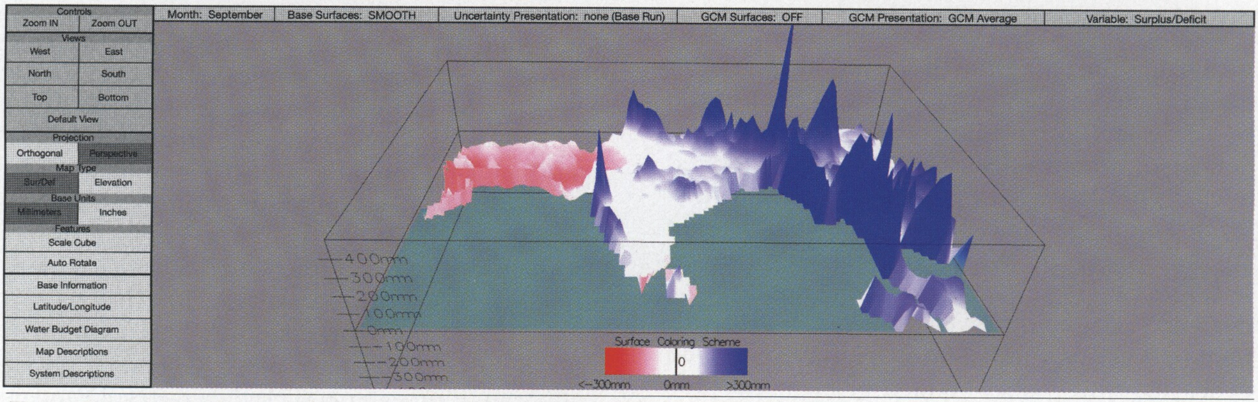

Month: September | Base Surfaces: SMOOTH | Uncertainty Presentation: none (Base Run) | GCM Surfaces: OFF | GCM Presentation: GCM Average | Variable: Surplus/Deficit

Figure 2. The visualization window (the right-hand two-thirds of the wall-size display). In this case a water surplus/deficit surface is shown for the month of September for a region centered on India.

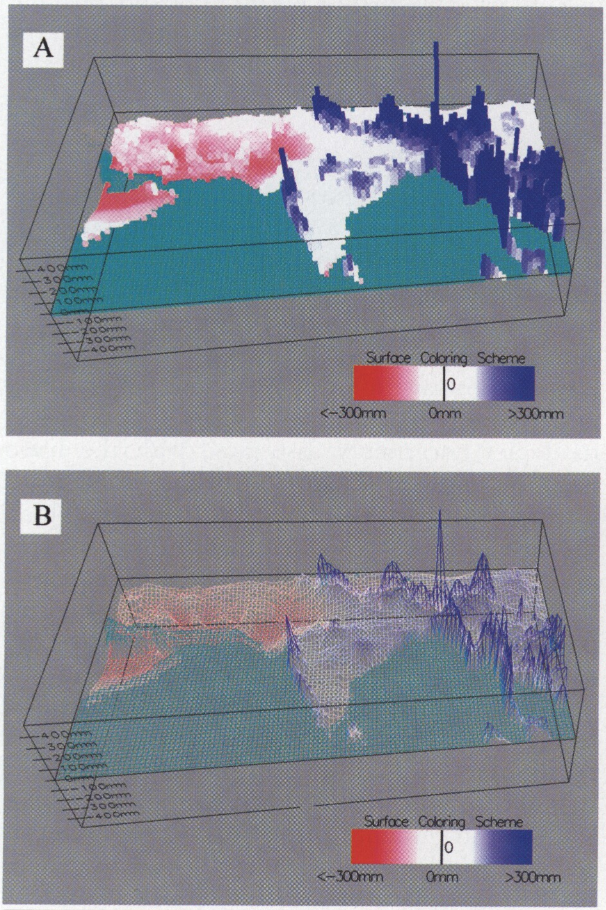

Figure 3. Prism-based (A) and mesh (B) surfaces for a region centered on India in the month of September (compare with Figure 2, which shows a smooth surface).

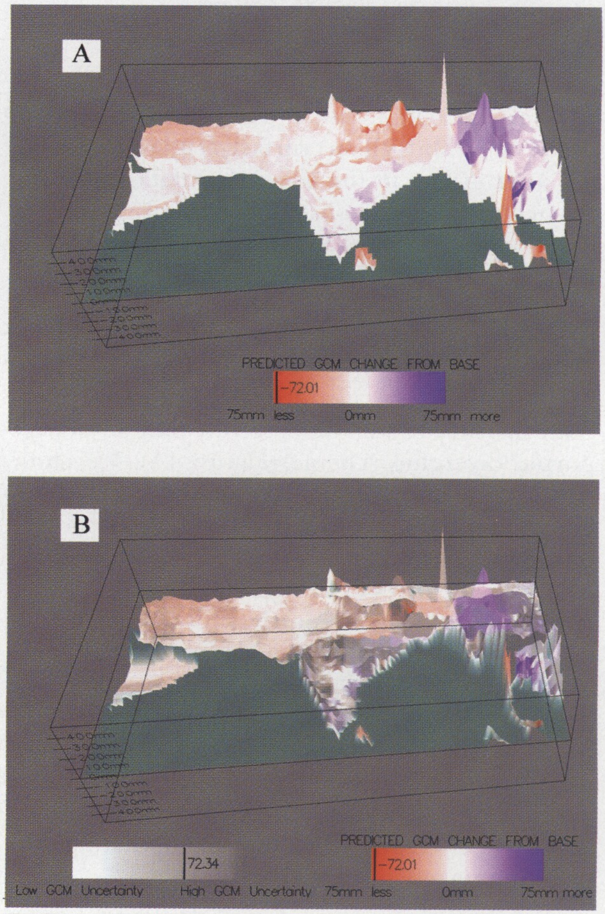

Figure 4. (A) An orange–purple diverging scheme used to depict the average of the GCM scenarios; (B) the uncertainty of the GCM scenarios is indicated by making the image more transparent in areas of greater uncertainty.

(a proportional value was calculated by dividing a grid cell uncertainty by the maximum uncertainty for any grid cell), we could have used a cube with corners (0,0,0) and (1,1,1) to present the range of RGB values. One obvious problem is that one cannot see the colors throughout the interior of such a cube.[6] Instead, we chose to display a triangular slice through the cube using (1,0,0), (0,1,0), and (0,0,1)

---

[6] Brewer (1994) also noted this problem, arguing that a trivariate color scheme should only be used for three variables that sum to 100 percent.

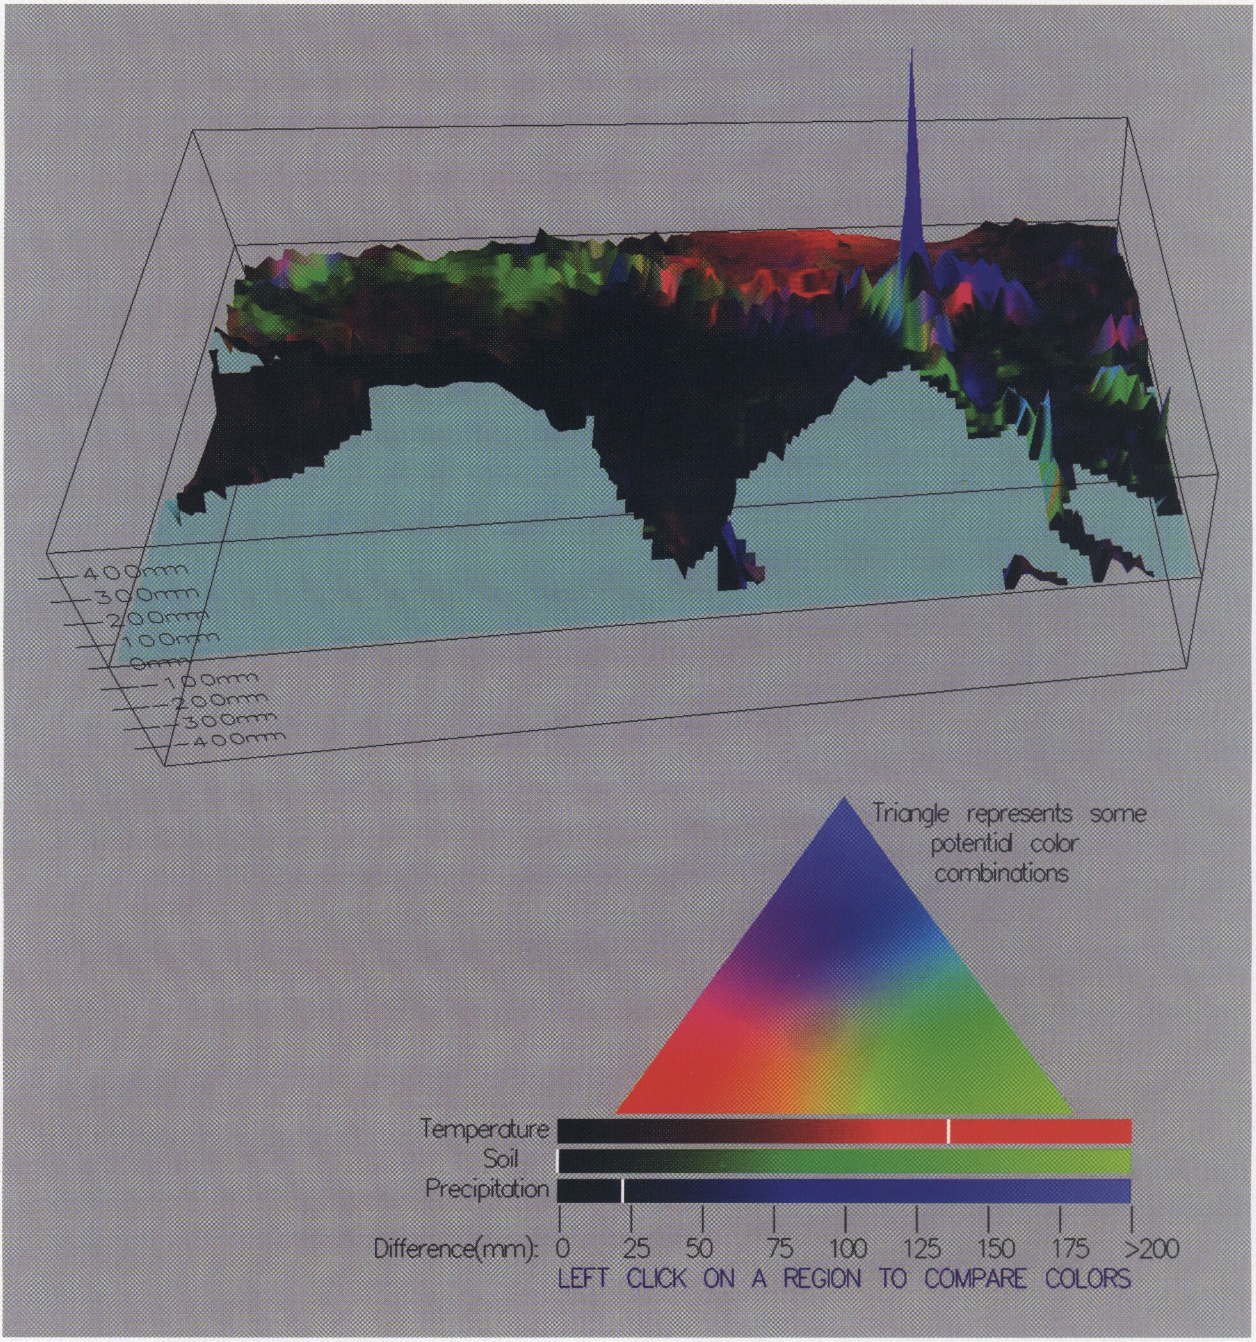

Figure 5. The RGB method for depicting uncertainty associated with three basic input parameters to the water balance model; the three-dimensional aspect of the surface represents surpluses and deficits in the base model.

as the vertices (Figure 5). Note that all points in this triangle have colors whose RGB components sum to 1. Below the triangle are three bars, indicating the extent of uncertainty in the results due to temperature, soil moisture, and precipitation at any particular location selected by the user. Other potential problems with the RGB method are that the uncertainties due to the fourth input parameter (the sub-model for potential evapotranspiration) can not be visualized and that the method is not suitable for users with red-green color deficiencies (Olson and Brewer 1997).

An extrinsic approach we used to depict the uncertainty of basic input parameters employed a set of vertical bars or glyphs (Figure 6). The color of the bar corresponded to one of the four basic input parameters, while the height of the bar corresponded to the magnitude of the uncertainty created by the variations in an input parameter. (We will refer to this as the Basic Input glyph method to distinguish it from a subsequent GCM glyph method.) With this method, the base surface is shown as a mesh so that it does not hide any of the bars.

Visualizing Uncertainty of the GCM Scenarios

We also utilized intrinsic and extrinsic approaches to visualize the uncertainty of GCM scenarios. For the intrinsic approach, we modified the average GCM scenario by making the image more transparent in areas that are more uncertain (Pang et al. 1994). Note that the resulting image (Figure 4B) actually shows three variables: 1) the base surface—the 3D surface, 2) the average change based on the GCMs—the orange-purple scheme, and 3) the uncertainty between the GCMs—the transparency.

For the extrinsic approach, we utilized the GCM glyph method shown in Figure 7. In 7A, we see that, as in the Basic Input glyphs, the size of vertical bars indicates the magnitude of uncertainty. The bars, however, are color coded in purple or orange to indicate whether more or less water will be available under a particular scenario. A bar can have both purple and orange components if one GCM simulation indicates more water will be available, while another GCM simulation indicates less water available.

In Figure 7B, we see that small "pyramids" can be drawn at the end of each bar indicating the GCMs associated with the most extreme predictions. Colors for these pyramids were based on a set that Brewer et al. (1997, p. 413) indicated would be easily named.

Usability Engineering

Overview of our Usability Engineering Approach

In selecting a general approach for usability engineering, we first considered the work of Gabbard et al. (1999), which Slocum et al. (2001) recommended as potentially appropriate for a broad range of geovisualization applications. Gabbard et al.'s approach involves four major steps: an analysis of potential user tasks prior to software development, an evaluation of the software by usability experts, having actual users work with a broad range of software functions, and a comparative evaluation of selected user tasks. Software development or refinement follows each of these steps, producing the eight-step process shown in Figure 8.

We used a variant of Gabbard et al.'s (1999) approach, as shown in Figure 9. In step 1, rather than attempt to analyze potential user tasks, we chose to develop a software prototype, largely based on the third author's domain expertise in water balance models and climatology. We took this approach for several reasons. First, since this type of software had not been developed before, we felt that the intended user group—decision makers—might have difficulty visualizing the nature of water balance models and associated issues of uncertainty. Second, we were unsure ourselves what sort of displays might result when we attempted to visualize the uncertainty of water balance models. Although the third author could hypothesize the nature of such uncertainties, he was curious to see the actual uncertainties displayed. Third, we had done little work with the wall-size display, and so were anxious to develop a geographic application that might serve as an "advertisement" for the display. (We felt that university administrators would be more apt to support our efforts to attain funding if we had something to show them.) Also along these lines, a PhD student with strong computer programming ability (the second author) was interested in undertaking the programming effort as part of his dissertation research. Finally, our intention was to develop a *data exploration* tool—we felt that decision makers could better understand the potential for such a tool if they could see some of the capability demonstrated.

Since we felt the prototype software might be biased toward the third author's interests, we decided to have domain experts evaluate the software in step 2. Our thinking was that the domain experts might be able to suggest a number of tasks that we had not thought of implementing. Based on the domain expert's evaluation, we revised the software (step 3) and had usability experts evaluate it (step 4) in a fashion analogous to Gabbard et al. In step 5, we refined the software based on usability experts' concerns. Finally, in step 6 we had actual decision makers work with the software. We did not implement Gabbard et al.'s comparative evaluations (their step 7), as we felt that such evaluations would only be appropriate after we had developed a fine-tuned application. At this point, we were trying to develop a feel for some of the issues involved in working with decision makers, uncertainty, and the wall-size display. Since this was the first time this sort of software had been developed and we had no previous experience with usability engineering, we expected some pitfalls, but hoped that any lessons learned could be transferred to those developing similar software.

It is important to recognize that this was our first attempt at developing this type of software, and so we were looking for comments from participants that could lead to substantial improvements in the software. As such, we conducted each of our test sessions in an interview format and provided scripts that largely controlled the order in which displays were shown to participants. Thus, it did not make sense to record detailed user interactions with the software; rather, the basic data we collected consisted of tape-recorded comments made by the participants.[7] We made no quantitative analysis of these qualitative data as we were looking for "key" comments that could lead to substantial improvement in the software. In this context, comment "A" made by a single individual could be more important than comment "B" made by several individuals.

Step 1: Develop Prototype Software

Although our prototype software was based largely on the third author's expertise in water balance models and climatology, we all contributed to its development. To accomplish this, we met on a weekly or bi-weekly basis over a period of three months. The bulk of these meetings took place in front of the wall-size display. In fact, this was one advantage of the large display—it enabled us to collaborate easily in developing the software. To be sure, we could have accomplished this in front of a large CRT display, but we think the process was more effective with the wall-size display.

Since we knew that domain experts would be working with the software in the second step, we made no attempt to develop a polished product at this stage. This software lacked the option buttons shown on the left in Figure 2 and no legends were shown (with the exception of the RGB legend). At this stage, the notion of varying transparency to represent uncertainty was applied to a white average GCM-predicted surface as opposed to the orange-purple average GCM scenario. In addition, no base information (such as country boundaries) were available, and when the user pushed a mouse button to rotate the map, the map continued to rotate until the user pushed the button again. (In the completed software, the image only rotates when the user pushes the button and moves the mouse, although continuous rotation is an option.)

Steps 2 and 3: Domain Expert Evaluation and Associated Software Refinement

For the domain expert evaluation, we interviewed those with expertise in water balance models: two civil engineers, two atmospheric scientists, and two geographers. We asked these experts to work through a script consisting of three major sections: an Introduction to the notion of water balance models and associated measures of uncertainty, a Tutorial illustrating major features of the software, and a set of Summary Questions (Appendix A—http: /www.ku.edu/~cagis/SlocApp.doc). We presented the script orally to experts (rather than having them read it) because we felt this would encourage them to respond orally and thus generate discussion about the software. At this stage, however, participants actually operated the controls of the system. Note that we posed numerous questions in the Tutorial section in the hopes that experts would not only evaluate current features of the software, but that they would suggest other features or options that should be included.

Because the second author programmed the system, he attended all of the interview sessions so that he could respond to experts' questions about the software. Additionally, at least one other author attended each session to assist in answering questions about water balance models, uncertainty, or the visualizations. Interview sessions lasted from 1 to 2 hours. This may seem a long time, but none of the experts seemed disturbed by this length; rather, they found the system intriguing.

As it turned out, domain experts tended to focus on usability aspects of the system, as opposed to issues related to their domain expertise. As a result of their comments, some key features added in step 3 were:

- Orthogonal (or bird's eye) views of displays;
- Options to include geographic location information such as country/state boundaries and latitude and longitude;
- Interactive legends for all displays (clicking on the map indicates a value in the legend);
- Banners describing each visualization;
- Applied transparency to the orange–purple average GCM scenario and used the term "visibility" rather than "transparency," as experts were confused by the transparency term; and
- Put rotation under direct control of the user, as opposed to having the image rotate automatically.

When experts did utilize their domain expertise, they generally spoke of how the system might be modified to suit their own interests. For instance, one said "I'm interested in runoff, stream flow, water supply…change in runoff is more important to decision makers than this ET business." Experts

---

[7] The sessions were also videotaped, but due to background noise from the air conditioning system, the audio portion of the videotape was not particularly useful. The visual portion, however, was sometimes used to interpret what was said on the tape recorder.

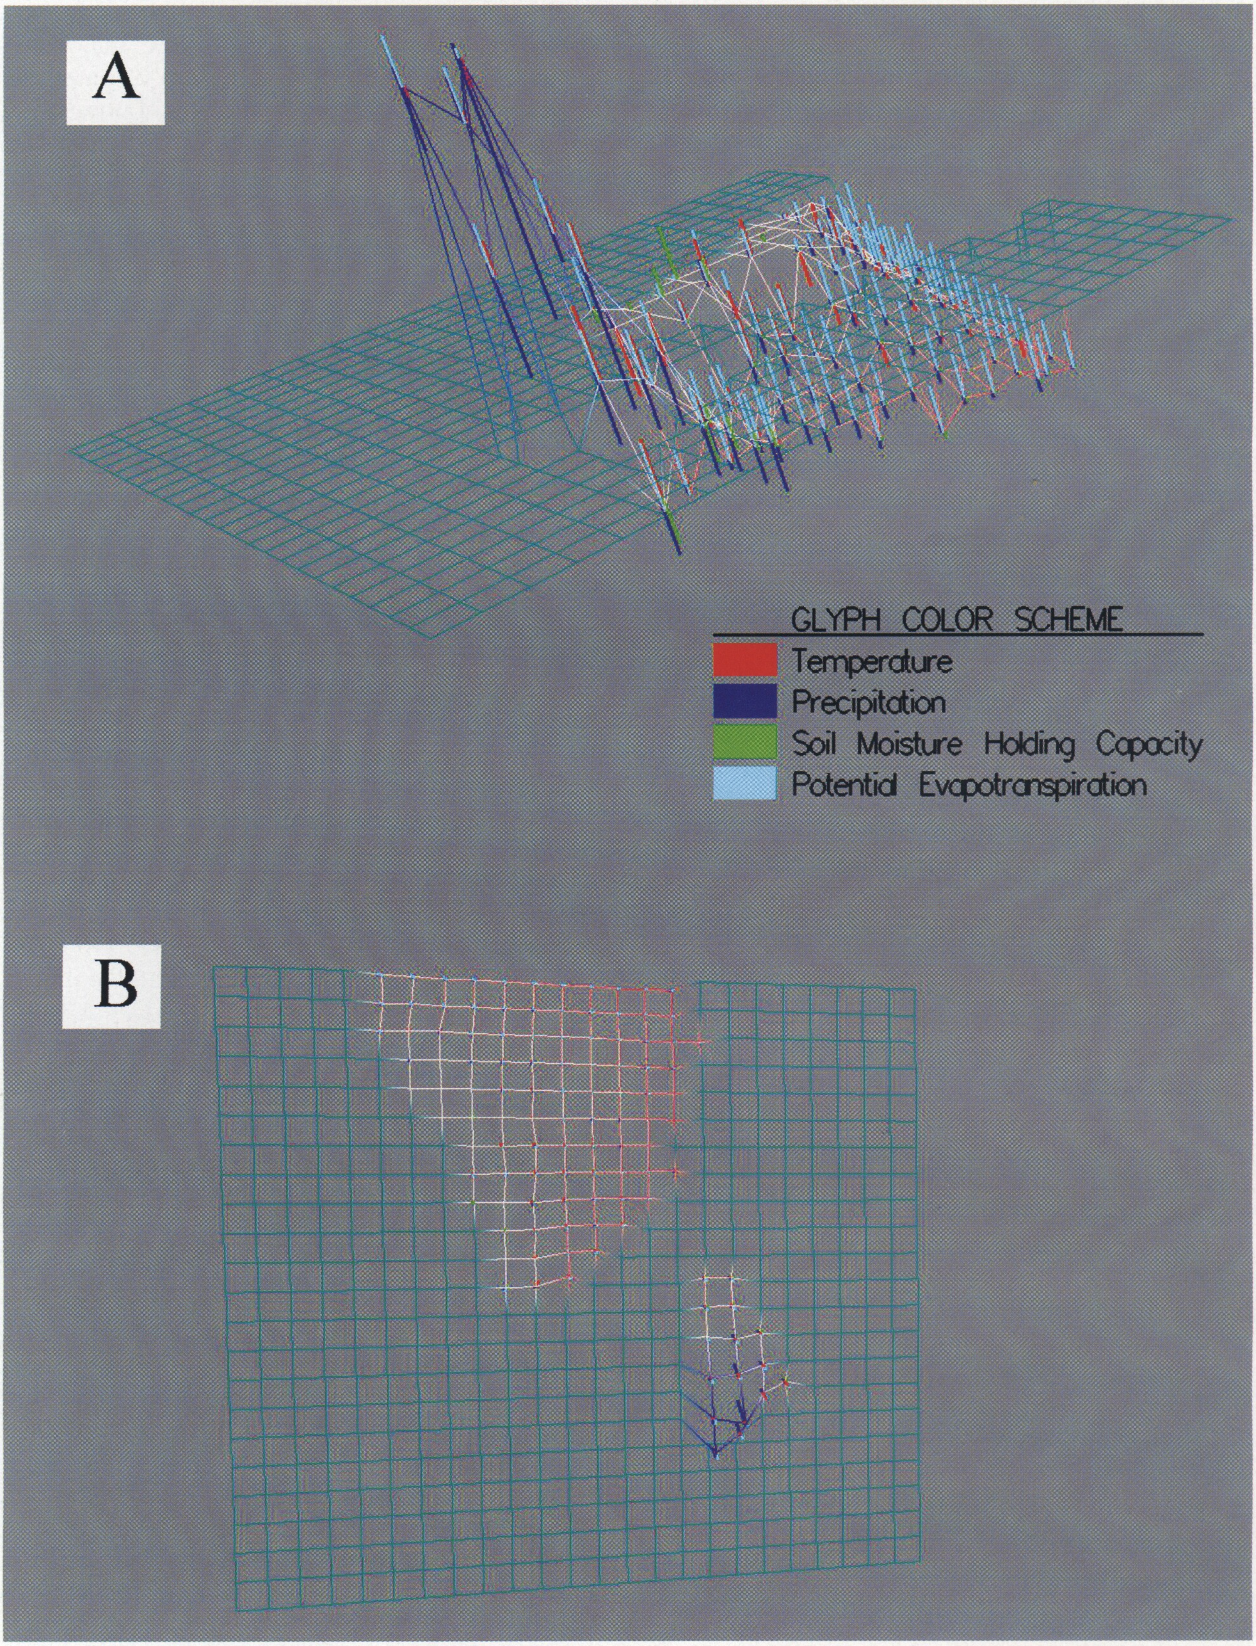

Figure 6. A) The Basic Input glyph method for depicting uncertainty associated with the four possible input parameters to the water balance model; B) a top-down view of the image seen in A, indicating that the area focused on is the southern tip of India and the island of Ceylon. The view in A is from the southwest and about 30 degrees above the horizon. (Note that the base surface is shown as a mesh in each case).

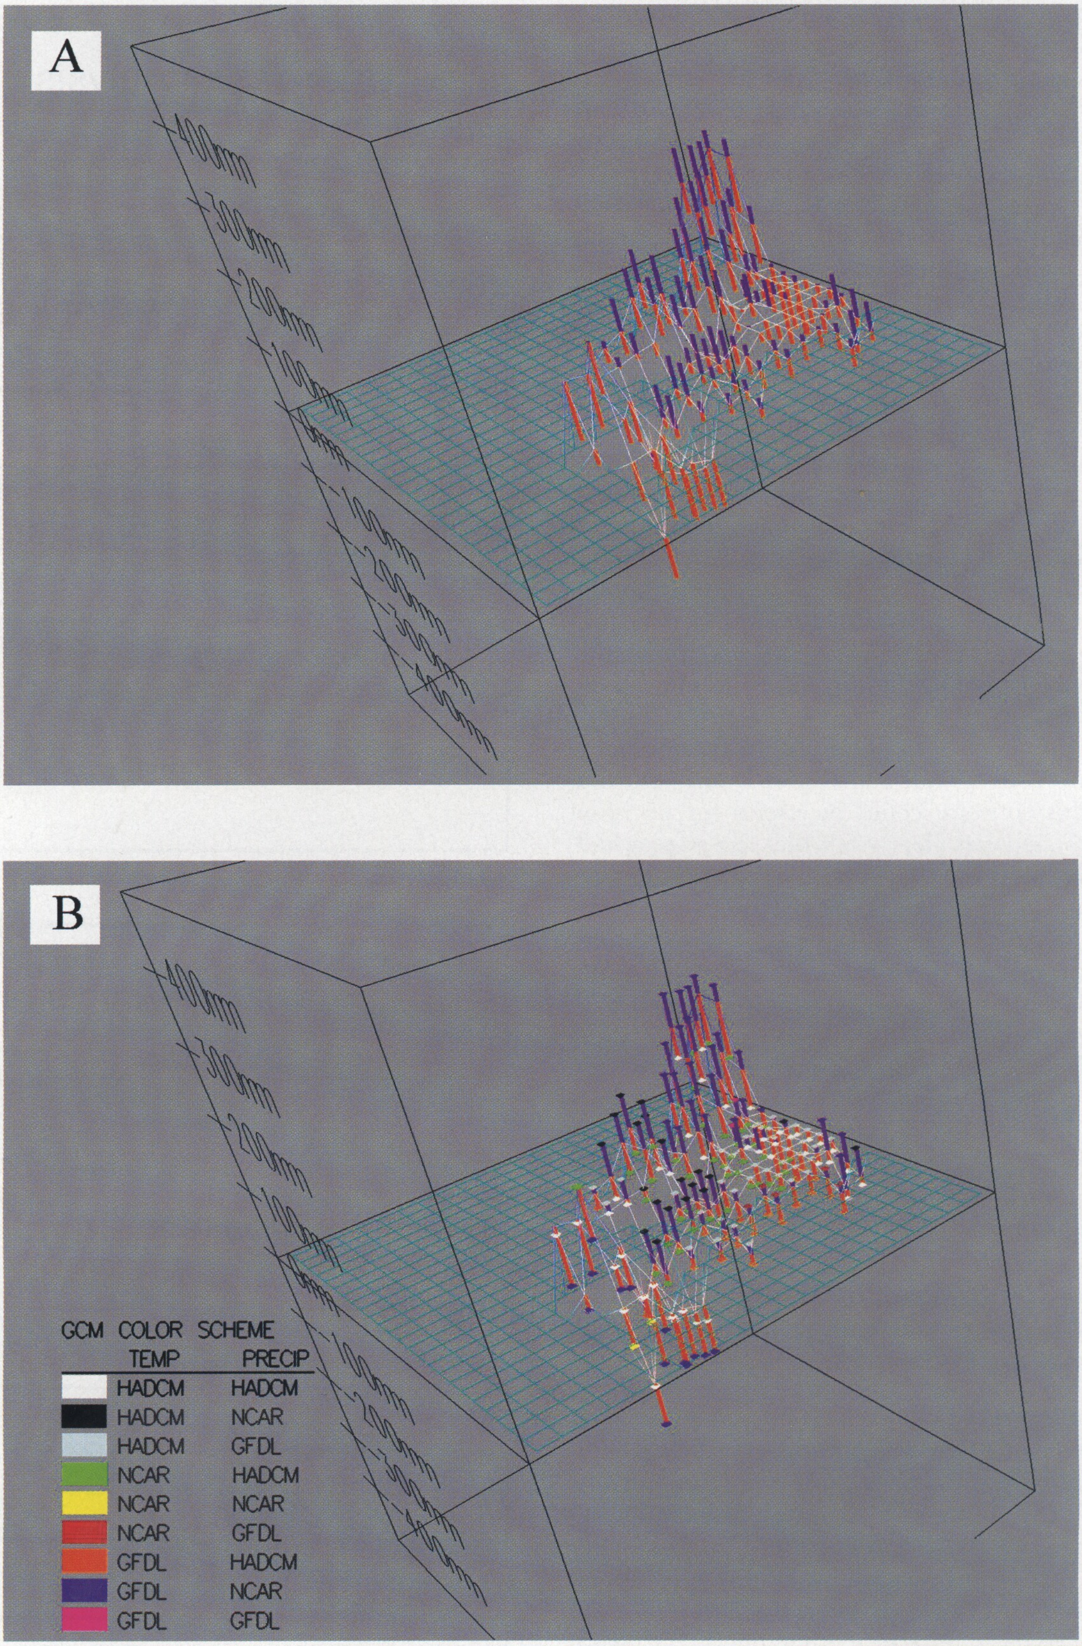

Figure 7. (A) the GCM glyph method for depicting uncertainty associated with GCM scenarios; (B) the GCM glyph method with pyramids included representing each possible combination of GCM. (Note that the base surface is shown as a mesh in each case. Views are of the same region depicted in Figure 6.)

did, however, raise several important issues. One was that the change in water availability due to the GCM scenarios could be less than the variation in water availability from year-to-year; this could lead water managers to say "so what?" Another concern was that our system did not consider the role of irrigation. One expert said, "just because the potential ET exceeds the rainfall doesn't mean they will be short on water…for most of the U.S. west of the Mississippi that's true." A third concern was that we should consider relative deviation rather than raw deviation because "a large deviation may not be important if you have a large precipitation." A fourth concern was the need to show weather station locations for each input dataset so that a user could evaluate which dataset might be best to apply to a particular region. Each of these comments reflects on the intended use and capabilities of our present software. The current prototype is not intended to address all these questions; however, we do propose to address several of them as we improve the water balance model and upgrade the software in the future.

In responding to two of our summary questions, "What are your overall impressions?" and "Do you think this software could be useful for decision makers?", experts clearly felt that the system had great potential, but that it was too sophisticated for a decision maker to work with alone. One expert said, "It is extraordinarily rare that you would find a decision maker with the patience and skill to work through something like this…this would be useful for the staff support for the decision maker…", while another said, "Some of this stuff is probably way over their head. It depends on what level of decision maker you are talking about." Such comments supported what we had already anticipated—that working with the system would require a scientific advisor to support the decision maker.

Steps 4 and 5: Usability Expert Evaluation and Associated Software Refinement

Ideally, we wanted people for our usability evaluation who had extensive training in usability engineering. Because such people were not readily available at the University of Kansas (where testing was done), we chose experts whom we felt would at least be sensitive to usability issues—a psychologist, a computer scientist, and two geographers. All but one of these were faculty members. We asked

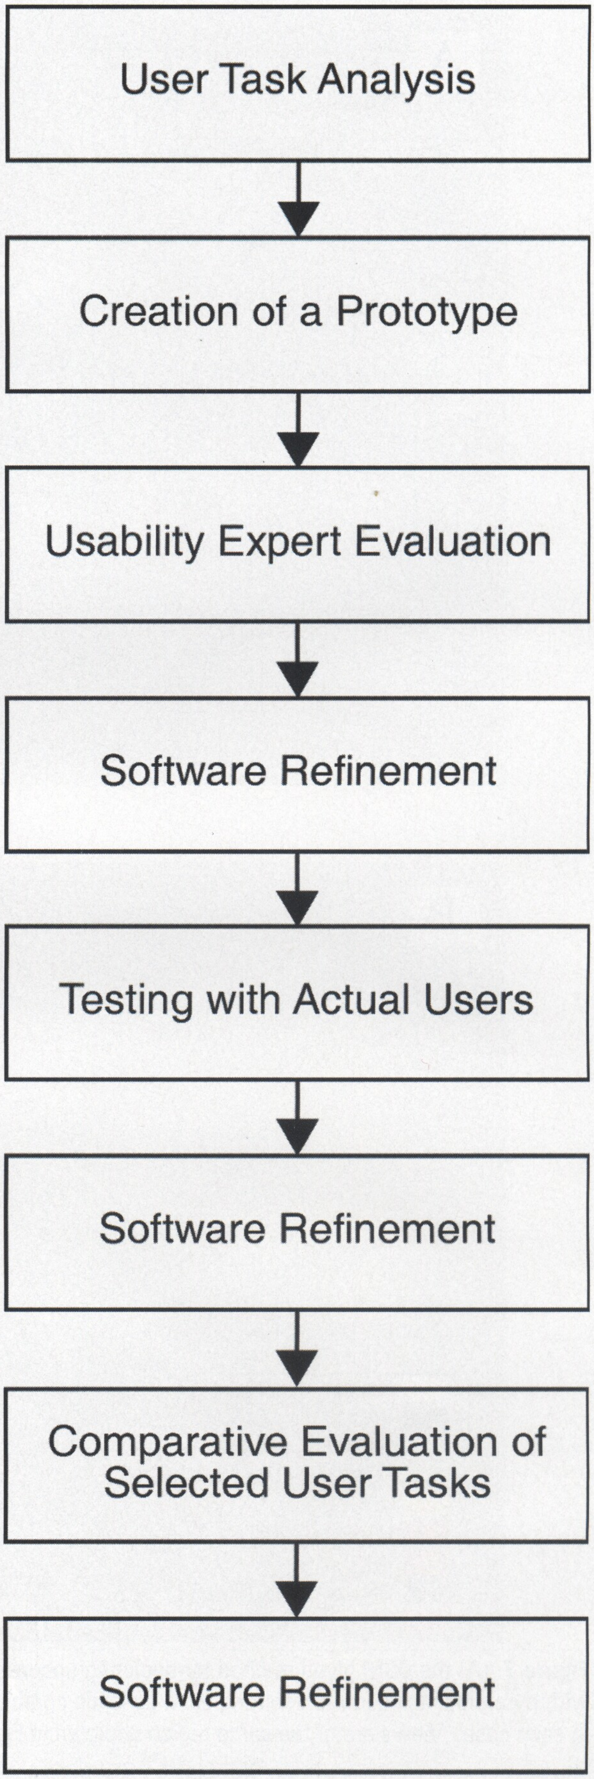

Figure 8. Gabbard et al.'s (1999) approach for evaluating usability viewed as an eight-step process.

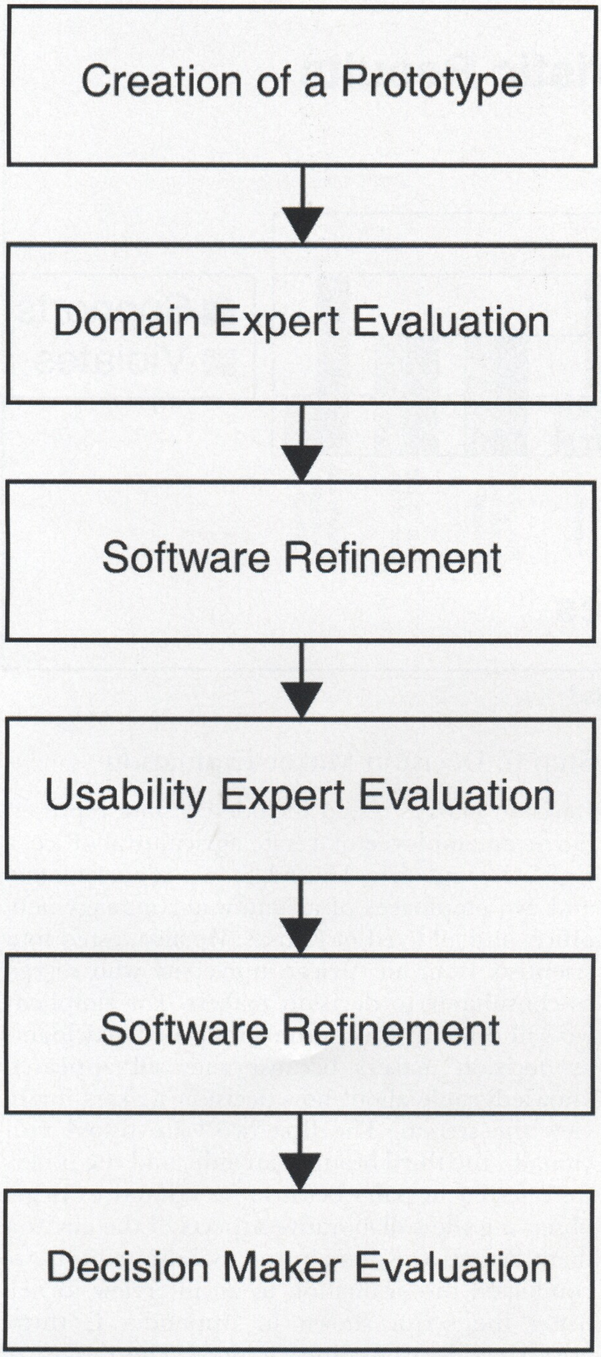

Figure 9. Our approach for evaluating usability of the software.

these usability experts to work through a script consisting of four major sections. The first two sections of the script (Introduction and Tutorial) were similar to the script shown to domain experts, except that the Tutorial reflected the changes we made in step 3 of software development, and it was shorter. In the third and fourth sections of the script, we asked usability experts to complete a task with the software and to evaluate the usability of the software using a set of heuristics previously utilized by Brewer (2001) (Appendix B—http:/www.ku.edu/~cagis/SlocApp.doc).[8] As with our domain experts, we presented the script orally and had participants actually operate the controls of the system. These sessions lasted even longer than those for the domain experts (one-and-one-half to three-and-one-half hours), but again, this did not seem to bother the usability experts as the software also intrigued them.

Usability experts made comments throughout the interview on how we could improve the system. Some of their concerns were: the distance between legends and the map was too long (partly a result of the large screen available); the inability to understand which menu options were currently selected (there was not enough contrast between the text and background of the selected and unselected menu options); no indication on the map of which point location was selected to determine a specific value; unclear map descriptions; and poor color contrast in the visual programming window. Although only one or two experts expressed each of these concerns, it was clear that users would benefit considerably if the concerns were dealt with. Therefore, we implemented the following changes to the software in step 5:

- Moved map legends closer to the surface;
- Made current selections in menus bolder and thus more obvious;
- Added indicators to the display (e.g., a yellow dot) depicting the point location focused on for specific information;
- Made map description banners clearer; and
- Made the color hues in the visual programming window less intense.

Other features desired by usability experts included the ability to:

- Display relative as opposed to absolute change (also suggested by some domain experts);
- Apply the full color scheme to the region currently selected as opposed to the entire world (this would enable a user to contrast areas on the map more easily);
- Animate the display to examine changes from month-to-month;
- Show multiple maps (e.g., the base display adjacent to an uncertainty display); and
- Utilize a "help" system.

Due to time limitations, we did not implement these other options, although ultimately they certainly would be desirable. One expert who was familiar with research on 3D versus 2D mapping also spoke

---

[8] We modified Brewer's (2001) wording slightly for our particular application.

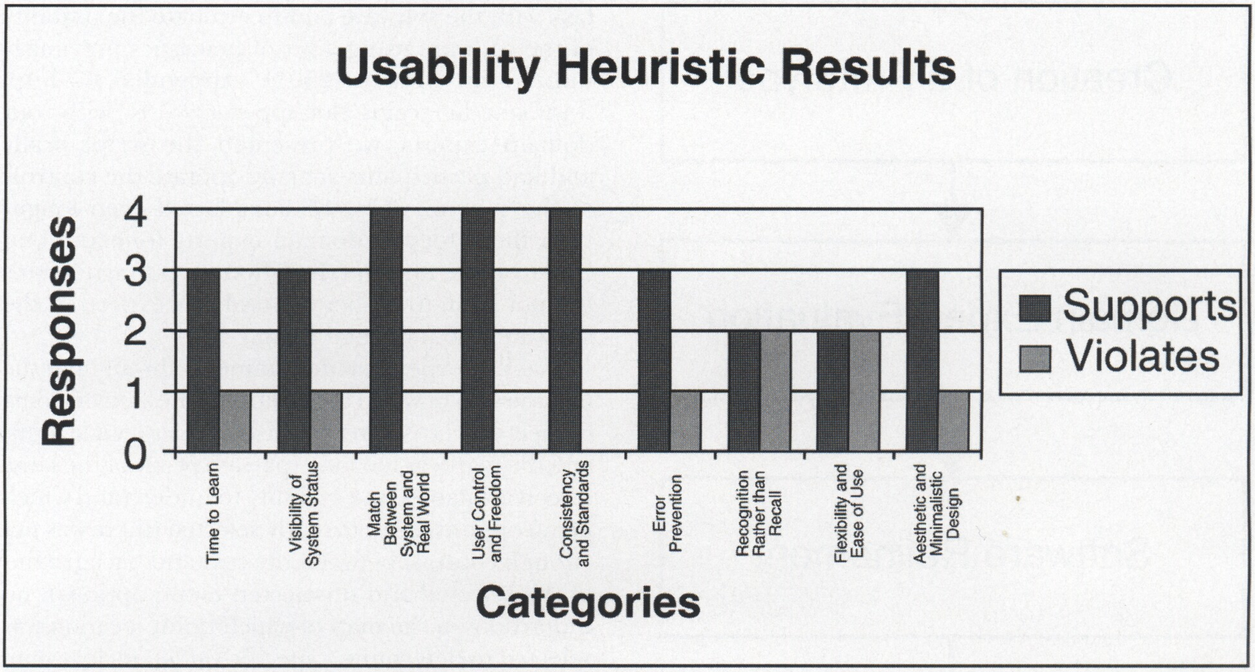

Figure 10. Results of the heuristics evaluated by usability experts.[9]

of the need to proceed with caution when using 3D visualization. He noted that 3D "...can help you develop in the long term a better mental model if 3D information is important...", but he stressed that 3D may be inappropriate if visualization can be accomplished using 2D. The expert made these comments early on in the interview and did not return to them once he saw the complete system. Several experts also noted the difficulty of displaying the uncertainty of water balance models to both novices (decision makers) and experts (scientific advisors to decision makers). As one expert said, "So your users are going to be scientists and decision makers, but those two groups have vastly different knowledge bases. Scientists will look at this and intuitively understand...but if I were [a] decision maker I might want that [referring to an element in the display] there all the time..." Such comments reaffirmed the notion that decision makers would have to work closely with a knowledgeable expert in order to use the system.

Although usability experts noted numerous ways in which the system could be improved, their evaluation using the heuristics in the fourth section of the interview was quite positive. As shown in Figure 10, for each heuristic, at least two out of the four experts felt that the software supported that heuristic, and this was before the above changes were implemented in step 5.

Step 6: Decision Maker Evaluation

Decision makers tested included a state representative, an employee of a state agricultural office, a legislative aid to a United States representative, and two employees of a state water management office, and all lived in Kansas. We also tested four scientists from an African university who served as consultants to decision makers. For simplicity, we will refer to this entire group of participants as decision makers because they all appeared knowledgeable about how decision makers might view the system. The first two were tested individually, the third brought an aide, and the others were tested in pairs because we wanted to begin observing the collaborative aspects of the environment. As with the domain and usability experts, we conducted the evaluation in an interview format using the script shown in Appendix C (http: /www.ku.edu/~cagis/SlocApp.doc). The design of this script, however, was quite different from the former scripts.

Following a brief introduction, we asked participants a series of questions pertaining to the decision making process in their organization (see section II of Appendix C). We then had participants work through a hypothetical scenario involving wheat production under different climate projections (i.e., looking at the impact of climate change on a region's competitive advantage in the wheat industry) (see

---

[9] Reprinted from *Computers & Graphics*, Vol. 26, Daniel C. Cliburn, Johannes J. Feddema, James R. Miller, and Terry A. Slocum, Design and evaluation of a decision support system in a water balance application, pp. 931-949, Copyright (2002), with permission from Elsevier.

section III of Appendix C). To make the evaluation relevant to participants, we chose a region that would be meaningful to them—thus, those from Kansas worked with Kansas as a region, while those in Africa worked with Zambia (Appendix C assumes Kansas as the region). In section III, participants focused on uncertainty related to GCMs, i.e., the transparency and GCM glyph methods. If time permitted, we also had participants experiment with the uncertainty depicted by the RGB and Basic Input glyph methods (section IV of Appendix C). Finally, all participants responded to the summary questions in section V of the Appendix.

Since the evaluations in steps 2 and 4 had shown that the software would be challenging for a decision maker to use on their own, the third author played the role of a scientist familiar with the software and presented a script orally, which was similar to the one shown in Appendix C. The second author manipulated the mouse and keyboard based on directions given by the third author. At least one other author also attended these sessions so that we would develop a full appreciation of how useful the system might be.

There were several interesting findings from having decision makers work with the software. One was that we developed an understanding of how the decision process works at the state and federal levels in the U.S. and, consequently, how our software might ultimately fit into the process. Test participants indicated that either government agencies or lobbyist organizations influence state or federal representatives (the ultimate decision makers) in an effort to influence legislation. In this process, scientists from government agencies or lobbyist organizations present information to representatives and/or their aides—the decision makers then use this information along with other information (such as written materials and public opinion) to form a decision. Participants stressed that the process is often collaborative in the sense that one or more scientists present information to one or more decision makers. Given this process, it seems that one approach for us would be to convince government agencies or lobbyist organizations of the usefulness of this sort of software.

Comments from the state legislator were especially useful in suggesting how we might actually reach such agencies or organizations. The legislator stressed that one of the keys to decision-making is determining who the important decision makers are: "They are going to lead the questions, they are going to frame the debate, they are going to be the ones that define decision-making that occurs. As scientists or students of government, it's your job to figure out which four, five, or six people really matter and then get to them." Further, he argued that a key is to find an "advocate" or "champion." He stated, "There are a number of good ideas…that die because they don't have that champion pushing." He also indicated that this advocate could be either an individual or a group of individuals, such as a legislator or state agency. These seem to be important ideas that anyone developing software for decision-making could utilize. In our case, this particular legislator was sufficiently enamored with our software that we expect he could serve as our advocate.

A second finding arising from the evaluation was that decision makers are likely to feel uncomfortable with the notion of uncertainty. One participant bluntly said: "Politicians have trouble with uncertainty." Another said, "That's really powerful, the visual of global warming. But when you start bringing in all the models … that would really frustrate my boss, he would want to see one model." The danger is that decision makers may simply ignore problems that have a lot of uncertainty. For instance, the latter participant also said, "When political people hear and see all this uncertainty they would say 'well we don't need to put money into stopping global warming.'" The fact that decision makers are likely to feel uncomfortable with uncertainty suggests that we should not merely present the uncertainty, but also suggest how to deal with the uncertainty. For example, in the case of climatology information, additional weather stations and their placement might be suggested by the system.

A third finding was that participants clearly were excited by the potential provided by maps, which they termed "visuals" or "visualizations." One said, "Well, the traditional way to do [these] kinda' things is with tables and there are some people who deal with [these] things better than others. Numbers are a lot harder for me to read and to assimilate than a graphic is." Another said, "It is a spatial presentation, and I have a preference for spatial presentations. I think for the most part when you are talking about a large area, spatial presentation is the best way to go because it is the best way to get the mind on the whole picture." Such comments may not necessarily indicate the effectiveness of our software, but rather that participants were not aware of what maps might be used for.

Fourth, we found a desire for software that could provide higher spatial and temporal resolution. Regarding the spatial resolution, one participant said: "How small can you break this down…I am working with an interest group…in Sedgwick, Kansas…[They] got two inches of rain the other day.

I didn't...It is not going to be well received by those people who are looking at the impact on [Bob Jones' farm]." In terms of the temporal resolution, the same participant said, "most academics I know are dealing in things that are 20-30 years out...you've got to give me something that I can say I voted for this because it's good for [my constituents]." At the same time, the participant felt the software could be very useful at the regional or national level for long-term planning.

Finally, participants warned us that some decision makers could feel intimidated by the software. One said, "I'm looking at it and saying is that going to be a threat to someone who didn't finish high school or only finished high school..." Another said, "[decision makers]...may not be as scientific as you are, so I think simplification would be important." Such comments support our notion that a scientific advisor would have to work closely with the decision maker who might benefit from utilizing such software.

In spite of the limitations noted, it was apparent that all participants felt considerable knowledge could be gleaned from the software. One participant said: "...a model like this ought to drive K State's researchers to develop crops that need less moisture. I am convinced that ultimately western Kansas has to go to dryland farming or else you've got to shut down all the towns and move people out..." This notion was supported by another participant who said: "Well, I think you have [to] question whether to continue with the current cropping pattern or should we be focusing some of our efforts on developing crops that are less water intensive." The former participant also suggested the notion of combining the present software, which is based on climatic data, with The GreenReport, which utilizes remote sensing to monitor the health of crops and natural vegetation throughout the United States (see http://www.kars.ukans.edu/products/greenreport.shtml). Clearly, he felt that considerable benefit could be derived by linking these two approaches.

Discussion and Future Research

Steps in Usability Testing

Overall, we feel that our usability testing led to an improved piece of software that ultimately could be used by decision makers working in close association with scientific advisors. Those wishing to develop similar software, however, will need to carefully consider the steps utilized in usability engineering and how those steps are implemented. In our case, we chose to create a prototype and then have domain experts evaluate that prototype. Although we still feel that a prototype was useful, we probably could have attained more useful input from the domain experts if we had "driven" the system ourselves rather than let them drive it—as a result, we suspect that they would have focused more on the content (and associated potential tasks) than on its usability. Once domain experts have examined a prototype, it might also be useful to have them interact in focus groups (Morgan 1998).

We also should have considered getting our decision makers involved earlier in the software engineering lifecycle, possibly using them at step 2 rather than the domain experts. We chose not to do this because we felt that we needed to show a more polished product to the decision makers, but, on reflection, getting them involved earlier might have produced a more useful product. If we were to get them involved earlier, we would probably have wanted to get a better understanding of how they might use such software in their own decision making. As it was, our evaluation tended to focus on the particular tool we had developed rather than on finding out how it might enhance their decision-making.

It must be recognized that the usability evaluation of all of the visualization techniques used (including those for uncertainty) was subjective, as we collected no hard data on the ability to work with the techniques. Ultimately, we need to contrast the various visualization methods in controlled experiments where we can collect hard data (this would be along the lines of Gabbard's step 7—comparative evaluation of selected user tasks). We see two general types of experiments that can be done. The first would be to focus on a comparison of particular techniques (e.g., to contrast the Transparency (or Visibility) and GCM glyph approaches) for a range of user tasks (e.g., acquiring specific information and examining general patterns). Such tests ultimately could lead to improved symbology for depicting uncertainty. The second type of experiment would be to have decision makers (and their scientific advisors) actually utilize the software themselves—in this process we would not drive the software, but rather would observe the process. Researchers at Penn State currently are developing a suite of tools that could assist in analyzing data resulting from such experiments (Haug et al. 2001).

Effectiveness of Methods for Presenting Uncertainty

For each of our usability tests, we asked participants to contrast the effectiveness of methods for presenting uncertainty. Domain experts tended to prefer the extrinsic methods (Basic Input and

GCM glyphs), as they provided more detailed information; the experts seemed concerned with examining all information available for a problem. In contrast, decision makers tended to prefer intrinsic methods (RGB and Visibility methods) because they wanted to get the "big picture." The intrinsic and extrinsic methods each had their limitations. The intrinsic methods were awkward for acquiring specific information, while extrinsic methods appeared very complex when a large area was shown.[10] To handle the problem of showing numerous glyphs over a large area, a generalized glyph routine could be developed in which the uncertainty of several grid cells would be averaged.

Although many participants found the RGB method useful, some struggled with it. The difficulty of understanding the mixtures of colors in the RGB method might be handled by portraying each form of uncertainty on a separate map. In fact such an approach could enable all four basic inputs to be depicted (presumably cyan would be used for the sub-model for potential evapotranspiration). Such an approach, however, would require that users visually correlate maps, which could prove problematic in understanding various combinations of uncertainty (Olson 1981).

Effectiveness of the Wall-size Display and Collaboration

Virtually all of the participants in our evaluations were enamored with the wall-size display. For the future, we plan to directly compare the wall-size display with more traditional displays such as large CRT screens and LCD projection systems. Our suspicion is that the traditional approaches will be less effective for several reasons. One is that CRTs and LCDs will not permit one to readily examine both the visual programming and visualization windows in detail simultaneously, as users did with the wall-size display. Secondly, users will not get the impression of being immersed in the data as they would with the wall-size display. Although we did not explicitly pose this question to users, *we* certainly felt more immersed in the data than when using a traditional CRT. Finally, the size of the display seems appropriate for initiating discussion among collaborators.

To achieve a truly collaborative system, we must allow individuals to control the display using their own personal input devices rather than forcing individuals to play "musical chairs" to get access to the keyboard and mouse. We are developing an approach based on the use of a Remote Application Controller (RAC), which is a Java program capable of running on a small hand-held Personal Input Device (PID) (Miller 2003). Each participant will have a PID running a RAC. The RAC is able to locate applications running in a shared environment and establish connections to them to perform dynamic viewing transformations, select objects, make menu choices, and so forth.

Improvement to the Water Balance Model

In addition to making improvements and modifications in the technology for disseminating our software, we plan to make improvements in the water balance model. These improvements concern some of the questions raised during the evaluation, such as the partitioning of water to model ground water recharge and surface water flows. In addition, we will develop a number of new methods for simulating human impacts on the water balance not related to atmospheric change (e.g., soil degradation and land cover change).

Summary

Our goal was to develop a visualization software tool that would enable decision makers to understand the uncertainty associated with water resources problems such as those involving climate change. As a first step in achieving this goal, we have developed a tool for a wall-size display that will enable decision makers to visualize both present-day and predicted future uncertainties associated with a basic water balance model. We feel that this tool will serve as the basis for a more elaborate tool that can be utilized for collaborative visualization when collaborators are at the same or different locations.

We developed our tool via the principles of usability engineering. We chose to modify the standard usability engineering approach by having domain experts evaluate a prototype software tool (rather than utilize the actual decision makers). We feel that domain experts might have provided more useful information if they had not been permitted to "drive" the system. On reflection, it might have been better to utilize decision makers to evaluate the prototype rather than domain experts. Ultimately, testing with actual decision makers will require that they (and their scientific advisors) have control over the system,

---

[10] For a detailed discussion of participants' comments on the methods for depicting uncertainty, see Cliburn (2001).

and that we are able to track their interaction with the system.

Although it was clear that participants were intrigued by the wall-size display, it is essential that its effectiveness be compared with more traditional technologies. Only when the wall-size display is proven to be more effective can the cost of such a system be justified.

ACKNOWLEDGMENTS

The development of the wall-sized display was funded by NSF Grant CDA-94-01021. Programming support for the second author was provided by grant 2300502 from The University of Kansas Center for Research Development Fund. We also wish to thank the participants in our usability studies who so graciously contributed their time.

REFERENCES

Brewer, C. A. 1994. Color use guidelines for mapping and visualization. In: A. M. MacEachren and D. R. F. Taylor (eds), *Visualization in modern cartography*. Oxford, U.K.: Pergamon Press, pp. 123-47.

Brewer, C. A. 1996. Guidelines for selecting colors for diverging schemes on maps. *The Cartographic Journal* 33(2): 79-86.

Brewer, C. A., A. M. MacEachren, L. W. Pickle, and D. Herrmann. 1997. Mapping mortality: Evaluating color schemes for choropleth maps. *Annals, Association of American Geographers* 87(3): 411-38.

Brewer, I. 2001. The design and implementation of temporal, spatial, and attribute query tools for geovisualization. Paper presented at the Association of American Geographers Annual Meeting, New York, New York.

Buttenfield, B. 1999. Usability evaluation of digital libraries. *Science & Technology Libraries* 17 (3/4): 39-59.

Byron, J. R. 1994. Spectral encoding of soil texture: A new visualization method. *GIS/LIS Proceedings*, Phoenix, Arizona. pp. 125-32.

Clarke, K. C., P. D. Teague, and H. G. Smith. 1999. Virtual depth-based representation of cartographic uncertainty. In: *Proceedings of the International Symposium on Spatial Data Quality*, Hong Kong. pp. 253-59.

Cliburn, D.C. 2001. Representing multiple uncertainties and evaluating usability in a spatial decision support system. PhD Dissertation. Department of Electrical Engineering and Computer Science, Lawrence, Kansas.

DiBiase, D., C. Reeves, A. M. MacEachren, M. Von Wyss, J. B. Krygier, J. L. Sloan, and M. C. Detweiler. 1994. Multivariate display of geographic data: Applications in Earth system science. In: A. M. MacEachren and D. R. F. Taylor (eds), *Visualization in modern cartography*. Oxford, U.K.: Pergamon, pp. 287-312.

Feddema, J. 1998. Estimated impacts of soil degradation on the African water balance and climate. *Climate Research* 10: 127-41.

Feddema, J. 1999. Future African water resources: Interactions between soil degradation and global warming. *Climate Change* 42: 561-69.

Florence, J., K. Hornsby, and M. J. Egenhofer. 1997. The GIS wallboard: Interactions with spatial information on large-scale displays. In: M.-J. Kraak, M. Molenaar, and E. Fendel (eds), *Proceedings of the Seventh International Symposium on Spatial Data Handling* (*Advances in GIS Research II*). London, U.K.: Taylor & Francis. pp. 449-63.

Gabbard, J. L., D. Hix, and J. E. I. Swan. 1999. User-centered design and evaluation of virtual environments. *IEEE Computer Graphics and Applications* 19(6): 51-9.

Gershon, N. 1998. Visualization of an imperfect world. *IEEE Computer Graphics and Applications* 18(4): 43-5.

Guimbretière, F., M. Stone, and T. Winograd. 2001. Fluid interaction with high-resolution wall-size displays. UIST'2001, Orlando, Florida. [http://interactivity.stanford.edu/publication.html].

Hamon, W. 1966. An equation for estimating potential evapotranspiration in terms of a temperature function and day length. In: *Methods for Estimating Evapotranspiration, Irrigation and Drainage Specialty Conference*. American Society of Civil Engineers, New York, New York. pp. 235-6.

Haug, D., A. M. MacEachren, and F. Hardisty. 2001. The challenge of analyzing geovisualization tool use. In: *Proceedings of the 20th International Cartographic Conference*, Beijing, China. CDROM.

Hibbard, W., and D. Santek. 1989. Visualizing large data sets in the Earth sciences. *Computer* 22(8): 53-7.

Howard, D., and A. M. MacEachren. 1996. Interface design for geographic visualization: Tools for representing reliability. *Cartography and Geographic Information Systems* 23(2): 59-77.

IPCC. 2001. Intergovernmental Panel on Climate Change. http://www.ipcc.ch.

Legates, D., and C. J. Willmott. 1990. Mean seasonal and spatial variability in global surface air temperature. *Theoretical and Applied Climatology* 41:11-21.

MacEachren, A. M. 1992. Visualizing uncertain information. *Cartographic Perspectives* (13):10-19.

Mayhew, D. J. 1999. *The usability engineering lifecycle: A practitioner's handbook for user interface design*. San Francisco, California: Morgan Kaufman Publishers, Inc.

McGranaghan, M. 1993. A cartographic view of spatial data quality. *Cartographica* 30(2&3): 8-19.

Medyckyj-Scott, D., and H. M. Hearnshaw (eds). 1993. *Human factors in geographical information systems*. London, U.K.: Belhaven Press.

Miller, J. R. 2003. The remote application controller. *Computers and Graphics* 27(4):605-615.

Mitas, L., W. M. Brown, and H. Mitasova. 1997. Role of dynamic cartography in simulations of landscape processes based on multivariate fields. *Computers & Geosciences* 23(4): 437-46.

Monmonier, M. 1992. Summary graphics for integrated visualization in dynamic cartography. *Cartography and Geographic Information Systems* 19(1): 23-36.

Morgan, D. 1998. *Planning focus groups*. Thousand Oaks, California: SAGE Publications.

New, M., M. Hulme, and P. Jones. 1999. Representing twentieth-century space-time climate variability. Part I: Development of a 1961-90 mean monthly terrestrial climatology. *Journal of Climate* 12(3): 829-56.

Nielsen, J. 1993. *Usability engineering*. Boston, Massachusetts: Academic Press, Inc.

Olson, J. M. 1981. Spectrally encoded two-variable maps. *Annals, Association of American Geographers* 71(2): 259-76.

Olson, J. M., and C. A. Brewer. 1997. An evaluation of color selections to accommodate map users with color-vision impairments. *Annals, Association of American Geographers* 87(1): 103-34.

Pang, A., J. Furman, and W. Nuss. 1994. Data quality issues in visualization. In: *SPIE Proceedings on Visual Data Exploration and Analysis*. vol. 2178, pp. 12-23.

Pang, A., C. Wittenbrink, and S. Lodha. 1997. Approaches to uncertainty visualization. *The Visual Computer* 13(8): 370-80.

Shepard, D. 1968. A two-dimensional interpolation function for irregularly spaced data. In: Proceedings of the 1968 ACM National Conference. pp. 517-23.

Slocum, T. A. 1999. *Thematic cartography and visualization*. Upper Saddle River, New Jersey: Prentice Hall.

Slocum, T. A., C. Blok, B. Jiang, A. Koussoulakou, D. R. Montello, S. Fuhrmann, and N. R. Hedley. 2001. Cognitive and usability issues in geovisualization. *Cartography and Geographic Information Science* 28(1): 61-75.

Thornthwaite, C. 1948. An approach to a rational classification of climate. *Geography Revue* 38(1): 55-94.

Wheless, G. H., C. M. Lascara, A. Valle-Levinson, D. P. Brutzman, W. Sherman, W. L. Hibbard, and B. E. Paul. 1996. Virtual Chesapeake Bay: Interacting with a coupled physical/biological model. *IEEE Computer Graphics and Applications* 16(4): 52-7. ■