How can we build more effective weather visualizations?

Lloyd A. Treinish IBM Thomas J. Watson Research Center, Yorktown Heights, NY, USA

lloydt@us.ibm.com http://www.research.ibm.com/people/l/lloydt

1. Introduction

Visualization has matured sufficiently that the design of content can be driven by user needs rather than enabling technology. Improvements in technology have typically led to generalized systems as the preferred mechanism to address a diversity of visualization strategies. Such flexibility is useful for research activities and applications development. However, their inherent lack of focus makes them less suitable in environments with relatively fixed tasks or user goals. This is especially true in operational situations, where there is no need to master generalized interfaces, whose many facilities may be viewed as superfluous.

To overcome this barrier, an understanding of user goals and how they map to visualization tasks is required. For example, Domik and Gutkauf have modelled user needs [3] while Card and Mackinlay [2] have creating a taxonomy for visualization design. Since domain-specific content in the visualization is required Jung among others have matched both interface and composition design to task [5]. These efforts achieved their decomposition after significant iteration with the users. Another approach is by automating the design process. For example, Zhou and Feiner discuss an expert system-based implementation focused on design elements, which also uses a taxonomy of data characteristics [16]. But this approach is only viable with limited visualization techniques and data. For more general scientific problems, the available techniques and the diversity of the data do not lend themselves to a tractable, expert-system solution. The user has considerable domain expertise that defines the tasks and is required for the interpretation of results. Having the user (intelligence) in the visualization process enables that expertise and the human capacity for pattern recognition to be used more effectively. As an alternate approach, a set of visualization tasks coupled with appropriate designs are developed a priori, and then refined through modest iteration. Generalized design elements are employed and tested to more efficiently develop focused visualizations.

2. Task-Based Visualization

To begin, consider three steps to defining visualization tasks:

i. Definition of the application in terms of user needs

ii. Composition of design elements and interface actions to implement that definition

iii. Establishing different techniques for various user goals

Prototypes are used to help focus step i. and to converge on results for steps ii. and iii. During that refinement, the tasks are decomposed hierarchically by recognizing that the user's tasks are not the same as the visualization tasks. For example, a given user may require one or more visualization tasks, and a specific technique may support more than one user task. The goals for the user are driven by the desire for specific results, such as feature or event identification or communication of the results. The visualization tasks consist of graphical or interface actions such as select, interact, animate and interrogate, which are used for specific composition like browse, analyze or present. To test these concepts for task decomposition in visualization design requires their application to an interesting set of problems, namely visualization of meteorological data for operational weather forecasting.

3. Related Visualization Work in Meteorology

Visualization in meteorology predates computing when scientists drew contour maps of weather data by hand. While researchers in atmospheric sciences have been early adopters of modern three-dimensional visualization methods, operational weather forecasting has focused on two-dimensional visualizations. The majority of turnkey visualization systems for meteorology are designed from the perspective of "one size fits all" - one interface and style of visualization independent of task to support a single class of users. Although such systems have been successful, there are user goals or operational efficiencies that are not addressed by their focused design.

For mission-critical tasks, improvements in speed and effectiveness have significant impacts, which is why weather agencies have developed focused visualization tools. One example is the Advanced Weather Interactive Processing System (AWIPS) deployed by the United States National Weather Service (NWS), which provides two-dimensional visualizations [10]. Similar tools are available from a plethora of other organizations worldwide, which are termed Class I. They provide colormapped or contoured two-dimensional scalar fields for analysis tasks by forecasters with minimal direct (graphical) interaction at a specific "layer", either the ground or at a constant atmospheric pressure. Given a flat canvas for visualization design, these tools can only show a few parameters simultaneously (e.g., overlaying wind as barbs or arrows, another scalar variable as line contours).

An example of a Class I visualization is illustrated in Figure 1 generated by the visualization tools discussed below. In this case, color contour bands of forecasted temperature are shown with overlays of wind barbs as well as coastline and river maps and political boundaries.

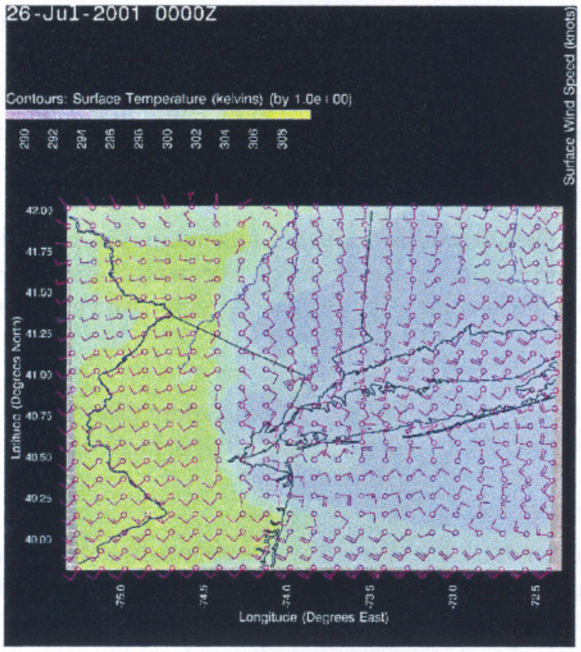

Fig 1. Class I Visualization.

Despite the generation of large, three-dimensional data sets, Class I techniques dominate in operational settings, for which there are few exceptions. Chief among them is Vis-5D developed at the University of Wisconsin. It has a fixed user interface with specific three-dimensional visualization tools. The implementation focuses on manipulation of regularly gridded data, preferably compressed to byte precision to increase computational efficiency. This yields an highly interactive tool that maps well to many meteorological data sets [4], which is in use by several operational weather centers, focused on analysis. However, for other forecasting tasks like model assessment and dissemination, Vis-5D does not have the ideal interface nor design.

An earlier effort by Forecast Systems Laboratory (FSL) provided operational three-dimensional visualization for analysis of weather models [11]. To eliminate duplication of the capabilities of Vis-5D, this work has changed direction to build directly upon Vis-5D. The efforts now concentrate on providing an interface consistent with other facilities (i.e., AWIPS) used primarily for analysis, based on evaluation of user preferences and tasks, especially in operational environments [12].

Fraunhofer Institut für Graphische Datenverarbeitung (FIGD) has taken a different approach. They have implemented independent systems that are focused on specific tasks. The first was Triton, oriented toward generating two-dimensional visualizations for the non-meteorologist [8]. The second, TriVis, is based upon a related goal - providing two- and three-dimensional visualizations for television broadcasts [9]. The third system, RASSIN, is designed to provide analysis facilities directly on the native grids of meteorological data [6].

4. Compositional Guidelines

To enable a set of design elements useful for a variety of tasks, visualization methods have been developed within a natural coordinate system to provide a context for three-dimensional display and interaction. They provide representations of the atmosphere consistent with the data source that are registered with ancillary or reference data (e.g., terrain and political boundary maps). The design of the content and choice of coordinate system has been dictated by the particular task to support both conceptual and physical realizations.

Since color is a critical aspect of design, knowledge of human perception is applied via a rule-based advisory tool that is sensitive to the spatial frequency of data and the visualization task [7]. It is employed in the design of specific elements, which are integrated into the final composition provided to users. For example, noisy data such as wind speed are primarily mapped into luminance, while smoothly varying data such as temperature are primarily mapped into opposing saturation pairs to impart an isomorphic or continuous representation. For moisture-related data (e.g., humidity and precipitation), two colormaps are combined such that dry regions are mapped to brown ranging through yellow to green for modest values. At high levels, the data are mapped into blue, with decreasing luminance. When contouring is used to map the data onto a set of bands, a segmented colormap with perceived ordinality is applied (e.g., Figure 1). For discrete three-dimensional representations (e.g., cloud surfaces), uniform but complementary colors are chosen to minimize the effects of color mixing. For direct volume rendering, the same hue is employed, but coupled with simultaneous mapping into luminance and opacity.

Several techniques are implemented for surface wind velocity, which are pseudo-colored by wind speed draped over a topographic surface. Vector arrows of fixed size are used to eliminate misleading motion cues during animation and to show gross atmospheric movement. In contrast, streamlines with directional arrows are superior at capturing fronts, convergence zones, vortices, etc. On the other hand, waving flags (rigid or furled) that point in the direction of the wind have been effective for the non-meteorologist.

The combination of these approaches provided a base of techniques to present to forecasters allowing greater effort on development rather than progressive refinement. Subsequent iterations in the composition were relatively minor such as improvement of specific colormaps or the choice of visualization task for analysis.

5. Results

The task decomposition leads to four other classes of visualization: Class II (2½-dimensional analysis), Class III (three-dimensional browsing), Class IV (three-dimensional analysis) and Class V (decision support).

5.1 Class II: Two-and 2½ – dimensional analysis

Class II can be viewed as a superset of Class I to include three dimensional enhancements. Its focus is for analysis by forecasters, particularly to support data comparison. Because the appearance of the visualizations may be complex, direct manipulation is provided. Up to five parameters may be visualized simultaneously. These two-dimensional variables may be any combination of surface or upper air layers from the same or different sources. They may be illustrated redundantly by applying multiple techniques (e.g., color and height). The variables and techniques can be independently selected interactively. An example is shown in Figure 2 with results of a mesoscale model, where the primary features are lifted index as the height of a shaded, deformed surface, pseudo-colored by temperature. Color-coded maximum reflectivity contours at 10 dBZ intervals are overlaid along with streamlines of wind. The wind direction is indicated by arrows and the speed by color. The surface is also draped with local coastline and political boundaries (black) and river (blue) maps.

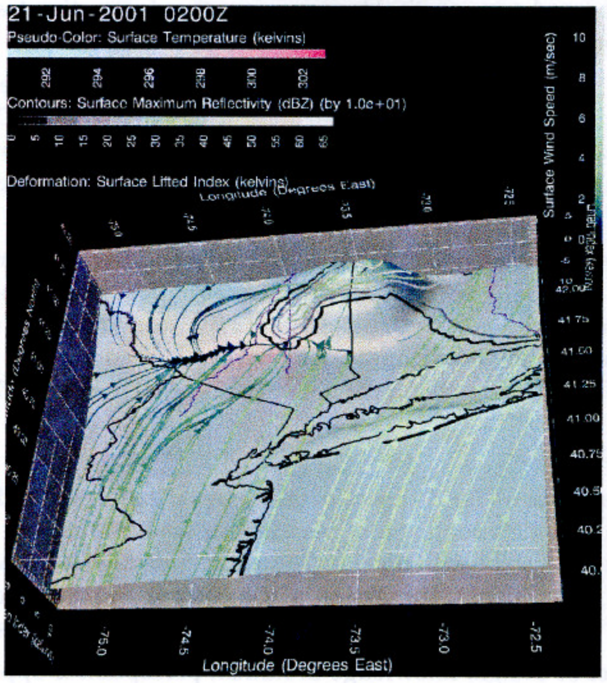

Fig 2. Class II Visualization.

5.2 Class III: Three-dimensional browsing

Class III enables forecasters to create qualitative three-dimensional representations for both interactive investigation and production of animation. The focus is placed on surface conditions and precipitation for general forecasting. The visualizations are composed of a set of simplified techniques for two purposes: gross assessment, and source material suitable for public dissemination (e.g., media, World-Wide-Web). They are presented in a geographic coordinates - cartographically projected horizontally and terrain-following (i.e., true height) vertically. The techniques requires high-resolution data (temporally and spatially) to enable a coherent presentation. An example is shown in Figure 3, which illustrates predicted cloud structure as translucent, white isosurfaces of cloud water density at 5x10^{-5} kg/kg. The cloud surfaces are registered with a terrain map overlaid with coastline (black), county (gray) and state (black) boundary maps, where major cities and airports are marked. This representation can show atmospheric motion and potential distribution of moisture.

The terrain is pseudo-colored by total precipitation to indicate where and how much rainfall is predicted, where heavy rainfall is shown as blue puddles. Translucent, cyan isosurfaces in the interior of the clouds are forecast radar reflectivities at a threshold of 30 dBz, approximating rain shafts. The correspondence between the rain shafts and the regions of relatively heavy precipitation is quite clear. The topography is also overlaid with streamlines of surface wind velocity, color-coded by speed. This visualization shows a correct prediction of thunderstorm activity in the vicinity north of New York City.

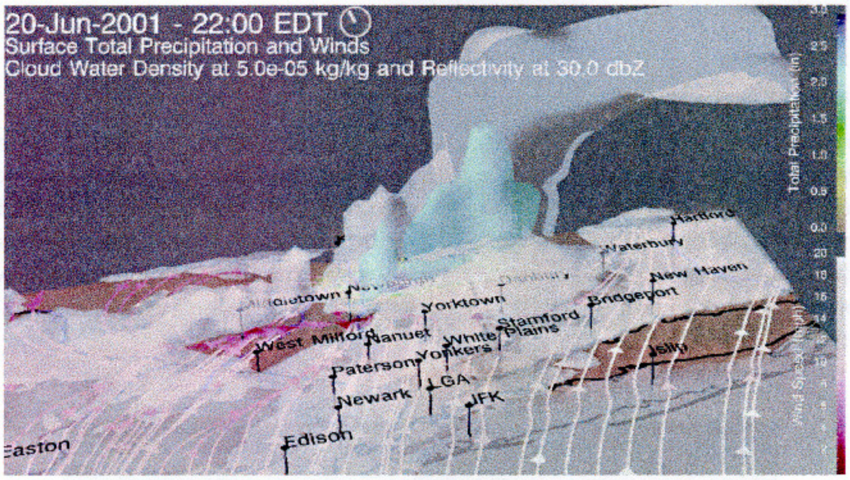

Fig 3. Class III Visualization.

5.3 Class IV: Three-dimensional analysis

Class IV provides viewing and interaction tools presented in geographic coordinates - cartographically projected horizontally but at standard pressure levels vertically. This class is similar to the visualization tasks addressed by Vis-5D and RASSIN, but with greater emphasis on direct manipulation and the introduction of new realization methods. Since these presentations can be visually complex even after application of complementary colormaps, facilities to interrogate and estimate data values are provided. The notion of a virtual metstation is introduced – a realization of direct manipulation as a graphical analogue for a simulated atmosphere to instrumentation that would be used to observe the real atmosphere.

An example is shown in Figure 4, also applied to mesoscale modelling results, but focused on the Hawaiian Islands. A surface variable (precipitable water) is displayed as pseudo-color contour bands, which are overlaid on a topographic map. Coastlines (black) are draped on the surface. An upper air variable (relative humidity) is displayed as an isosurface at 90% in translucent white as a representation of a cloud boundary. Another field (vertical wind speed) is shown as a vertical slice, which is pseudo-color contoured. Any of the three-dimensional fields available from the model can be visualized with either of these methods.

The upper air three-dimensional wind velocity is visualized via interactive marking of geographic locations for virtual soundings. At each location (two in this case), a vertical profile is extruded through the atmosphere. Each profile is real ized as a pseudo-colored tube, which is contoured by the variable selected for isosurface realization (i.e., humidity). The wind velocity along the profile is shown by a set of vector arrows that point in the direction of the wind. Horizontal speed at these points is indicated by the color and length of the arrows. Points along the profile are used as seeds for stream-ribbon integration, which are visible within the volume, and are also colored by horizontal speed. The visualization for the profile marked 2 can help illustrate vortex shedding in the lee of the island of Hawaii.

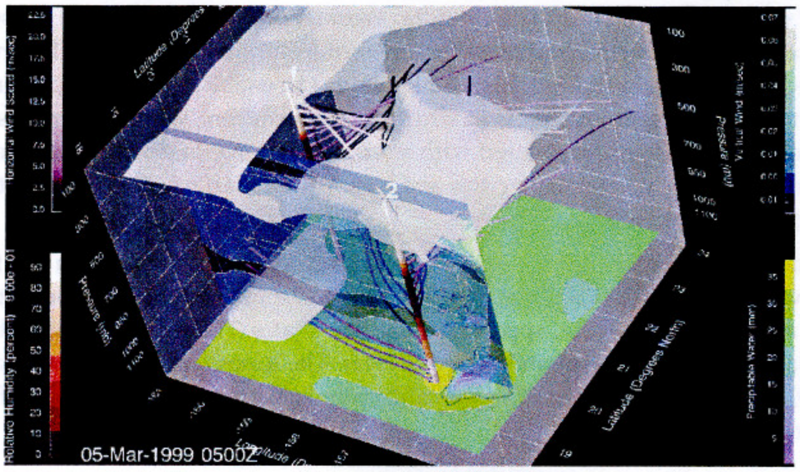

Fig 4. Class IV Visualization.

5.4 *Class V: Decision support*

User-driven design has potential for other applications of numerical weather prediction like aviation, agriculture, energy and emergency planning. But a refinement of the task decomposition is necessary and weather data must be correlated with other information relevant to decision making [15]. This idea is illustrated simply in Figure 5. It contains two frames of a 24-hour animation with 10-minute time steps. Each image shows a topographic map overlaid with coastlines, and state and county boundaries. An additional overlay is present, which is a map of the southeastern portion of the major components of the electricity transmission system in New York State (i.e., lines of capacity greater than or equal to 115 kV).

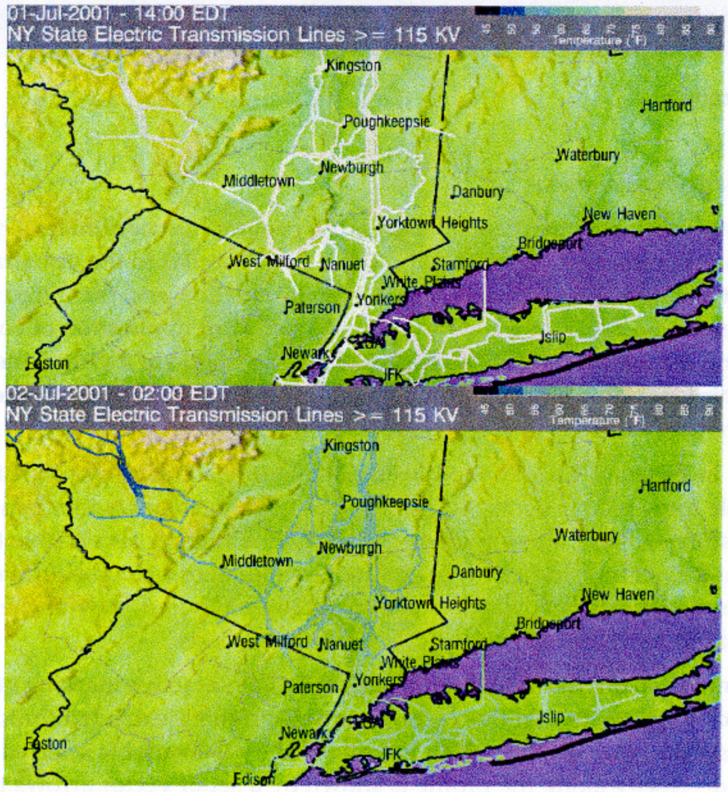

Fig 5. Class V Visualization.

The forecasted surface temperature from the three modelling nests are combined into a multi-resolution structure [13] and then interpolated to the geographic location of the transmission lines. The results are color contoured and fused with the other maps. The model prediction shows considerable variation in temperature along this system over a 12-hour period, thus illustrating the potential for mesoscale modelling in support of electricity transmission operations.

6. Implementation

A suite of tools for Linux, AIX and Windows workstations provides the facilities for all five visualization classes. They present simple user interfaces based upon XWindow/Motif for indirect interaction and OpenGL for direct three-dimensional interaction implemented in cartographic coordinates.

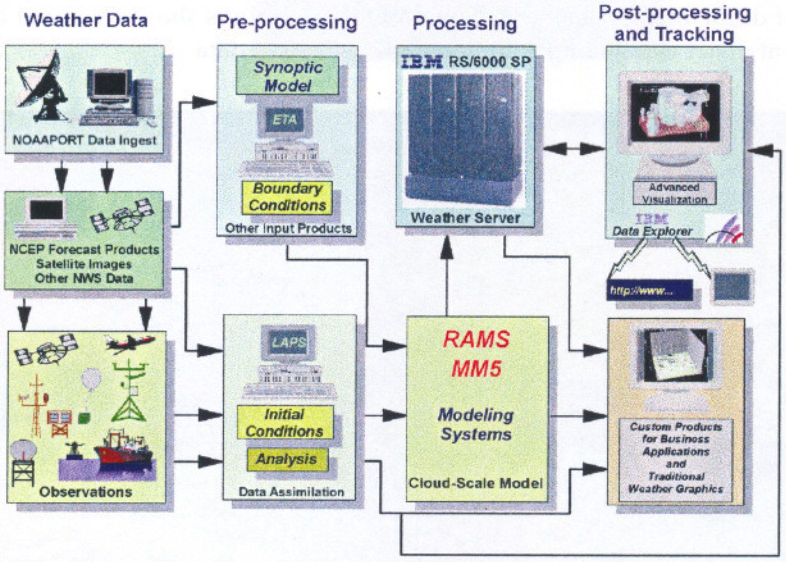

Fig 6. Architecture of Deep Thunder.

The tools provide forecasters interactive capabilities with three-dimensional representations of the state of the atmosphere (e.g., thermodynamics, moisture, clouds) derived from the data. They have been implemented with Data Explorer (DX) [1], a portable, open source, general-purpose software package for visualization and analysis (http:/ / www.research.ibm.com/dx and http://www.opendx.org). It employs a client-server architecture with an extended data-flow execution model.

This suite is part of an integrated mesoscale forecasting system, called "Deep Thunder". Additional information and visualization examples are available on the World-Wide-Web (http://www.research.ibm.com/weather/DT.html). It provides operational mesoscale numerical weather prediction via the architecture shown schematically in Figure 6. The system leverages the available infrastructure of weather data to do localized forecasts - observations for initialization via an assimilation step and synoptic-scale forecasts for boundary conditions. The simulation portion of the system is parallelized for rapid operational execution on an IBM RS/6000 SP distributed memory computer. All the aforementioned visualizations classes are integrated, with Class V being represented at the lower right with the others at the upper right. They have also been extended to included automated generation of images and animations for web browsers on the same parallel server [14].

7. Utilization

The Browser application (Class III) enables model assessment. Typically, an animation with 10-minute resolution over the full model run (24 hours long) is created after the forecaster selects the variables, techniques and geographic view. These would remain invariant throughout the animation. One or more animations are generated for local playback at work-station resolution to support media briefings to complement what is produced for the web. To aid in this selection process the forecaster would interactively move through the geographic scene, experiment with different displays and do limited animation either during or after model execution. To illustrate this process, consider the montage of seven images in Figure 5, which are sequenced from left to right, top to bottom. Start with a three-dimensional representation of the local area - a terrain map overlaid with state, county, coastline and river maps, and marked with major cities. Then, predicted temperatures colored by the scale at the upper right to show their continuous variation are overlaid. Surface winds are displayed as a set of arrows pointing in the wind direction, while the color corresponds to speed. Next, three-dimensional representations of predicted clouds are added, which are illustrated as a white, translucent isosurface as a "boundary" where the density of water is above a certain threshold. From this information, the simulation can be examined in more detail. Rather than temperature, wind chill is shown since the data showed fairly windy conditions and low temperatures. To illustrate regions where the wind chill might be particularly low, the data are shown as a set of bands, each with a distinct color that increase in value from dark to light. To examine the predicted winds in another way, they are shown as colored streamlines with arrows still indicating the direction. This illustrates a front moving through the area (i.e., lines bunched up at the lower right). Now, predicted wind chills for the specific time are shown for major cities. From this interaction, a representation that might be useful to publish in a newspaper or on the web to illustrate a forecast is chosen. To simplify the visualization, the wind data and some of the annotation are removed, and the geographic viewpoint of the map is changed. Now only the terrain is shown, colored by bands of wind chill prediction and values at major cities along with the maps and cloud data.

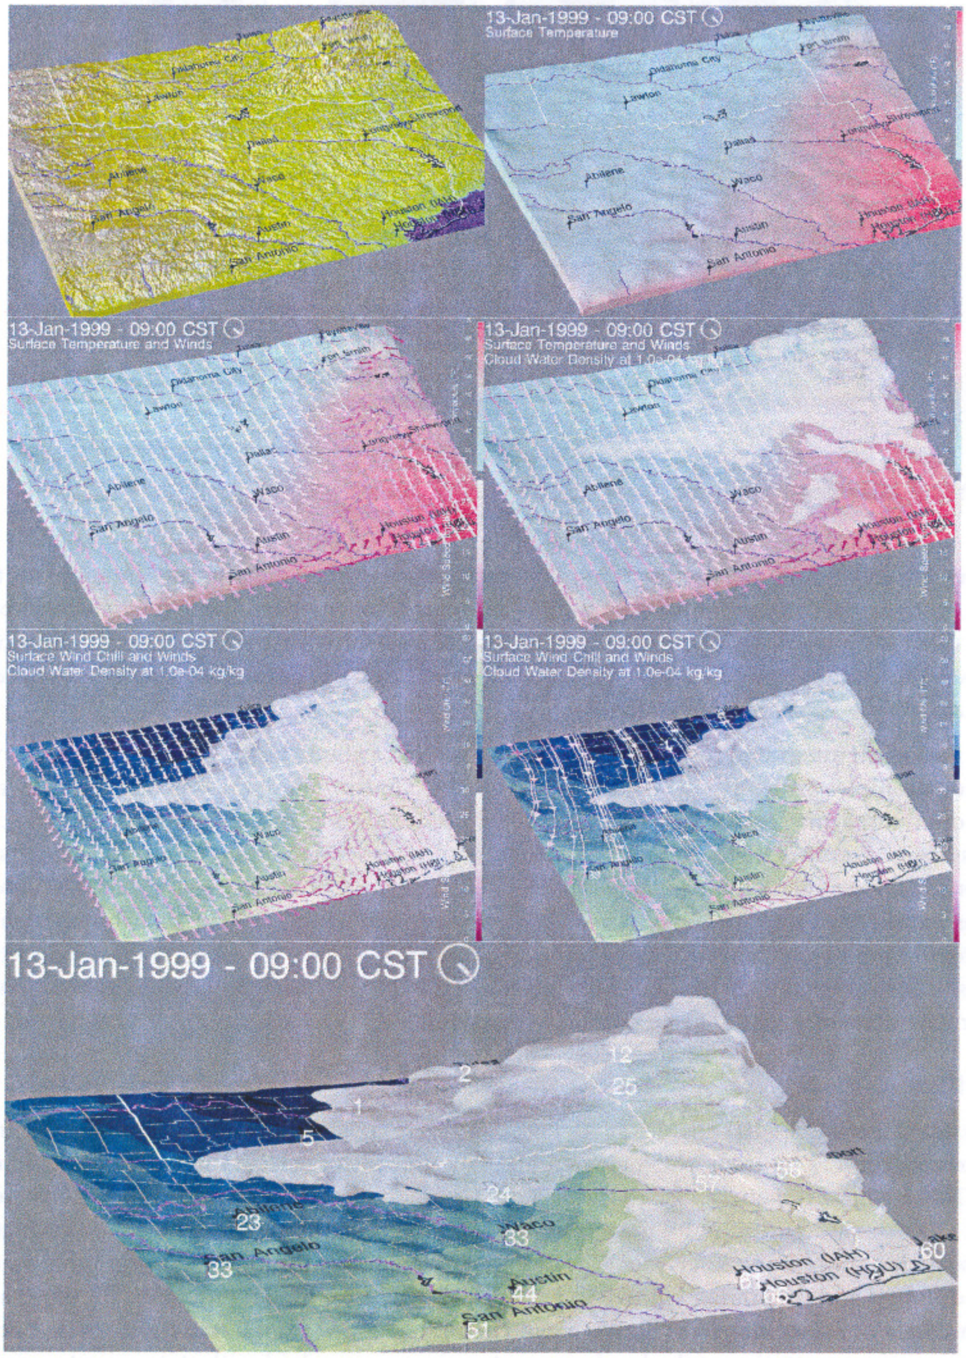

Fig 7. A sequence of images illustrating a typical use of the Browser application.

After each model or analysis execution, all of the results are collected and reorganized at hourly resolution for analysis with the Slicer and Viewer applications (Class I, II and IV). To illustrate their capabilities, consider the sequence of four images from left to right, top to bottom in Figure 8. The first image shows two surface scalar fields, mean sea level pressure as a continuous color field, and line contours of relative humidity. Another display shows the humidity contours as color-filled bands using the same segmented colormap, but now overlaid with 850 mb temperature values at specific locations and 750 mb winds as vector arrows, colored by speed. These fields are combined in the Class II visualization at the upper right, but only showing surface variables. The height of the deformed surface corresponds to lifted index, which reflects instability in the atmosphere. This representation is very effective, especially in animation, of showing the motion of a front. The image at the bottom, now shows contours of surface pressure, but combined with a three-dimensional representation of relative humidity, temperature and winds. The humidity data are shown as an isosurface at 75%, corresponding to a simple cloud boundary and sampled along two vertical profiles. The temperature data are visualized as a single vertical slice that is contoured. The wind data are shown as arrows as part of virtual wind profiler and as streamlines using the profile points as seeds.

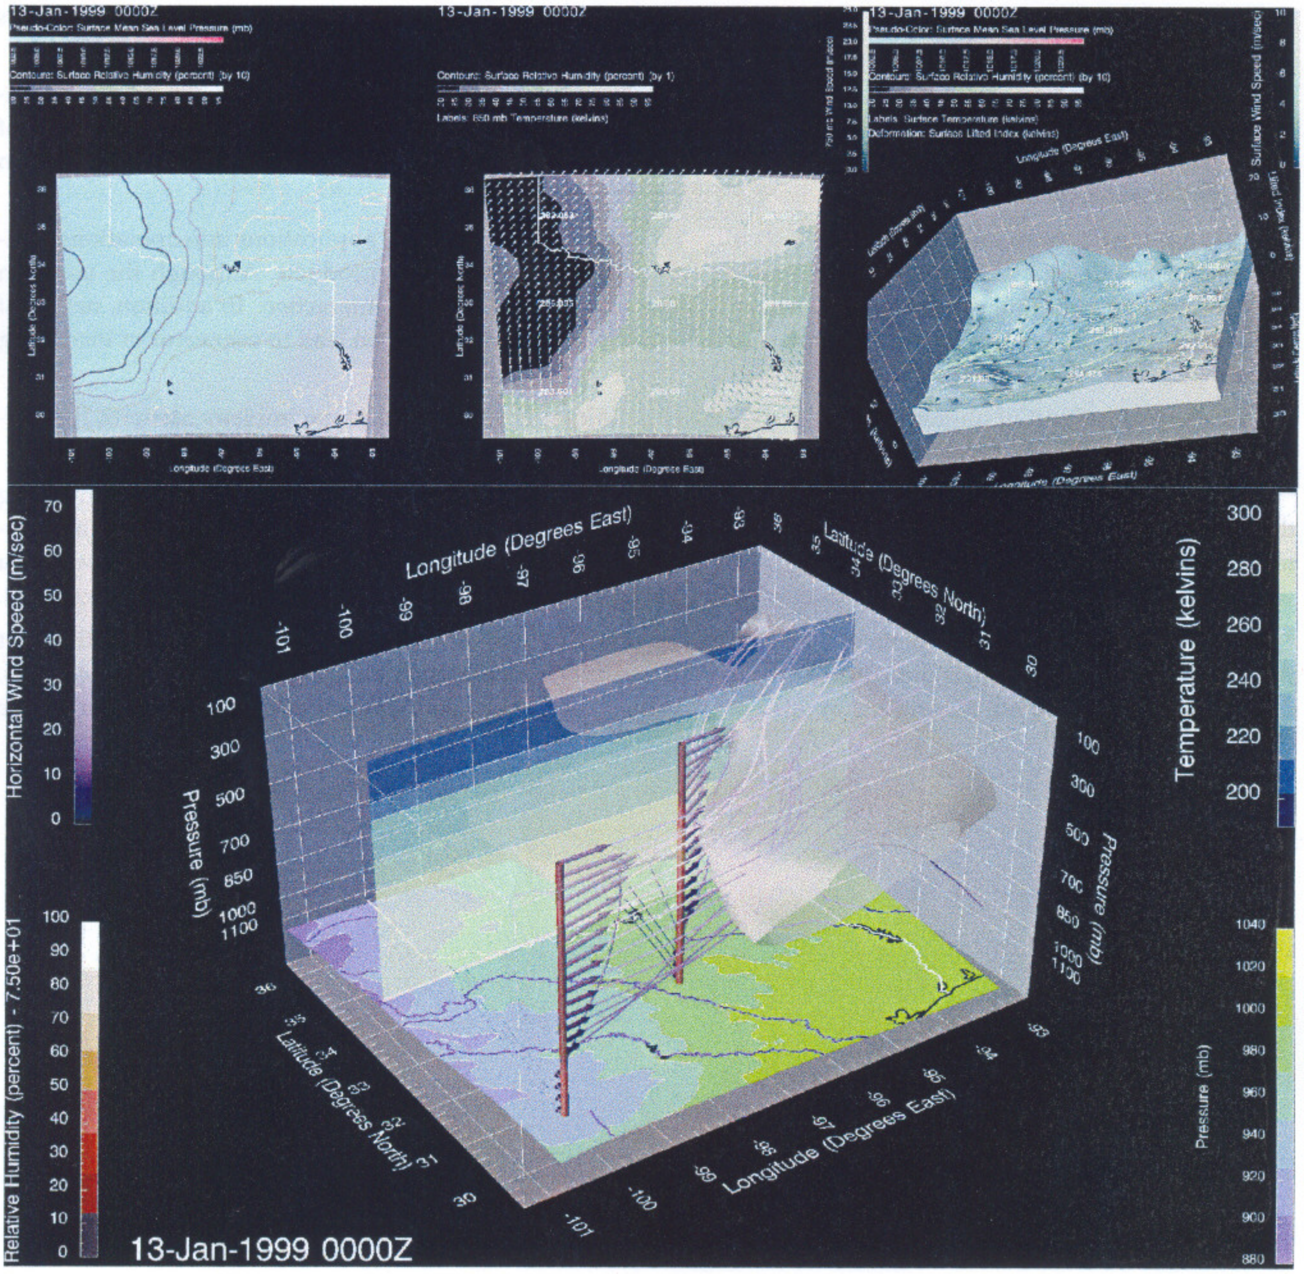

Fig 8. A sequence of images illustrating a typical use of the Slicer and Viewer applications for analysis.

8. Conclusions and Future Work

Specialized interfaces and tools matched to user goals and underlying visualization tasks can provide new facilities for operational activities. They can be characterized as being easy to master, even if the underlying capabilities may be sophisticated. Although a generalized system can be employed to provide similar functionality, the lack of focus in its interface increases learning time. Customized systems can reduce training costs but increases the expense of development. However, a generic tool can be used for both prototyping and efficient implementation by promoting high-level reuse of tools and design elements. Employing a generic toolkit (DX) also eliminated the need to implement a graphics and computational infrastructure. This is in contrast to the low-level reuse (renderer) in efforts by FIGD or code-level modifications to a turn-key tool (Vis-5D). Since DX is built upon an unified data model that enables direct operations without transformation or compression, customization for data types was not required as well as preserving fidelity during visualization.

Unlike FSL, the work herein considers a wider variety of user goals and visualization tasks. FSL has focused primarily on interactive visualization for analysis. Further, they have addressed the problem of training and usability by developing an interface consistent with other tools in operational forecasting environments [12]. Although FIGD has developed task-specific visualization content, they present different user interfaces and design elements requiring additional effort in both development and training.

Class III visualizations proved to be more effective than expected by virtually eliminating the laborious evaluation of numerous Class I images. Conceptual models that would require inference from a significant amount of two-dimensional data are obvious in three-dimensional animations. Further, one could easily infer vertical motion based on a three-dimensional display of clouds forming.

The subsequent introduction of the Slicer (Class I and II) and Viewer (Class IV) applications into operations complements Class III, but uncovered problems inherent in utilizing typical data products. Although the user can easily select a data of interest, the organization of the data is not designed for interaction. In addition, not all of the variables are consistently populated and have incomplete metadata, which can lead to user error or increasing the complexity of the application to compensate.

Class V is the subject of on-going but related work. In addition to several prototype implementations, the concept has been applied to operational decision support for local government agencies and transportation [15].

Visualizations for the web can be generated by the current applications after an intermediate step of migrating the products to a web server, which is advantageous in an operational environment because the forecaster has content control. However, the task decomposition can be further refined by considering direct generation within a web browser, where the user interface and the content must be simplified and automation can be introduced [14].

The notion of task-driven customization of content and interface has been successful in weather forecasting, but the idea can be applied to other domains. Likely candidates include measurements collected from medical scanners, the output of data mining algorithms applied to relational data bases, and utilization of results from terascale physics simulations.

References

[1] **G. Abram and L. Treinish**. An Extended Data-Flow Architecture for Data Analysis and Visualization. Proceedings of the IEEE Visualization 1995 Conference, October 29-November 3, 1995, Atlanta, GA, pp. 263-270.

[2] **S. K. Card and J. Mackinlay**. The Structure of the Information Design Space. Proceedings of the IEEE Information Visualization 1997 Conference, October 20-21, 1997, Phoenix, AZ, pp. 92-99.

[3] **G. O. Domik and B. Gutkauf**. User Modeling for Adaptive Visualization Systems. Proceedings of the IEEE Visualization 1994 Conference, October 17-21, 1994, Washington, DC, pp. 217-223.

[4] **W. L. Hibbard, B. E. Paul, D. A. Santek, C. R. Dyer, A. L. Battaiola, M. F. Voidrot-Martinez**. Interactive Visualization of Earth and Space Science Computations. IEEE Computer, 27, n. 7, July 1994, pp. 65-72.

[5] **V. Jung**. A System for Guiding and Training Users in the Visualization of Geographic Data. Proceedings of the 1st Conference of GeoComputation. School of Geography, University of Leeds, September 17-19, 1996, Leeds, UK, pp. 470- 482.

[6] **M. Lux and T. Fruhauf**. A Visualization System for Operational Meteorological Use. Proceedings of the Sixth International Conference in Central Europe on Computer Graphics and Visualization (WSCG'98), Plzen, Czech Republic, February 9-13, 1998, pp. 525-534.

[7] **B. Rogowitz and L. Treinish**. How Not to Lie with Visualization. Computers in Physics, 10, n.3, May/June 1996, pp. 268-274.

[8] **F. Schroder**. Visualizing Meteorological Data for a Lay Audience. IEEE Computer Graphics and Applications, 13, n. 5, September 1993, pp. 12-14.

[9] **F. Schroder and M. Lux**. TriVis - Professional Television Weather Presentation.
http://www.igd.fhg.de/www/igd-a4/ projects/docs/trivis/trivis_e.html, October 1997.

[10] **W. R. Seguin**. AWIPS - An End-to-End Look. Proceedings of the Eighteenth International Conference on Interactive Information and Processing Systems for Meteorology, Oceanography and Hydrology, American Meteorology Society, January 13-17, 2002, Orlando, FL, pp. J47- J51.

[11] **J. S. Snook, J. M. Cram, J. M. Schmidt**. LAPS/RAMS, A Nonhydrostatic Mesoscale Numerical Modeling System Configured for Operational Use. Tellus A, Dyn. Meteorol. Oceanogr. (Sweden), 47A, n. 5, pt. 1, October 1995, pp. 864- 875.

[12] **E. J. Szoke, U. H. Grote, P. T. McCaslin, P. A. McDonald and . D3D**: Overview, Update and Future Plans. Proceedings of the Eighteenth International Conference on Interactive Information and Processing Systems for Meteorology, Oceanography and Hydrology, American Meteorology Society, January 13-17, 2002, Orlando, FL, pp. J128-J133.

[13] **L. Treinish**. Multi-Resolution Visualization Techniques for Nested Weather Models. Proceedings of the IEEE Visualization 2000 Conference, October 2000, Salt Lake City, pp. 513-516, 602.

[14] **L. Treinish**. Interactive, Web-Based Three-Dimensional Visualizations of Operational Mesoscale Weather Models. Proceedings of the Eighteenth International Conference on Interactive Information and Processing Systems for Meteorology, Oceanography and Hydrology, American Meteorology Society, January 13-17, 2002, Orlando, FL, pp. J159-161.

[15] **L. Treinish**. Coupling of Mesoscale Weather Models to Business Operations Utilizing Visual Data Fusion. Proceedings of the Third Symposium on Environmental Applications: Facilitating the Use of Environmental Information, American Meteorology Society, January 13-17, 2002, Orlando, FL, pp. 94-101.

[16] **M. X. Zhou and S. K. Feiner**. Data Characterization for Automatically Visualizing Heterogeneous Information. Proceedings of the IEEE Information Visualization 1996 Conference, October 28-29, 1996, San Francisco, CA, pp. 13-20.