MVA'90 IAPR Workshop on Machine Vision Applications Nov. 28-30,1990, Tokyo

Identifying Regions of Interest in Complex Imagery *

Ronald P. Blanford and Mark R. Thomsen
TRW Inc.
Data Systems Center
One Space Park, Bldg O2/1779
Redondo Beach, California 90278

Abstract

We develop methods for manipulation and display of multidimensional data objects. These methods assist in the identification of precise subregions of interest when the data objects are extremely large and complex. Our results extend previous work in scientific visualization by generalizing the number of data dimensions, and by integrating novel interactive and analytic tools for focusing on regions of interest.

1 Introduction

People in many fields are increasingly using computers to explore data spaces in creative ways. Through interaction with the data they discover new relationships, view new perspectives, and gain insight into cause and effect. These discoveries are often fleeting observations that must be acted on quickly - change the data displayed or processed, modify parameters, move through the data, etc. An interactive interface is usually best to involve the person in an exploratory journey. Matching to the human bandwidth and focusing on "the interesting" more than "the mundane" are desirable features of the coming computer systems. Bandwidth is mostly a hardware issue. Focus is more of a design and software issue.

More precisely, determining the regions of interest and dedicating resources to them is the chief prioritization goal for an exploratory system. The amount of complex imagery data being collected in medical imaging, automated inspection, earth remote sensing, and other fields is making detailed human review of all the data impractical. However in exploratory work, final determinations on what is key should be reserved for human expertise. A promising avenue for research is using the computer to screen out the determinably mundane and cue on the determinably interesting.

We first discuss candidate measures of importance (each application potentially will use different measures). Automatic and manual techniques for directing resource use are then described and illustrated. Future directions are discussed.

*Research funded by TRW Large Scale Information Management Pilot IR&D.

2 Measures of Importance

Measures of inportance take a variety of forms in multidimensional data objects. Examples include edges, extrema, modes, patterns, deviation from expected patterns, change as a function of one or more variables (e.g., change with time), local independence of generally correlated variables, and extreme gradients in natural systems. The results in the sections that follow, on automatic and manual techniques, are essentially independent of specific measures because they assume weights are used. A single measure or a suite or measures (combined) are used to weight regions of the multidimensional data object.

One measure exemplifies this approach: edges. An *edge* is defined as a sharp gradient between regions (a discontinuity). In image and signal processing literature an edge in an array of numerical values is modeled as a zero crossing of the second derivative [7] or as numerical models of singularities. We note that the edge may be spatial, temporal, spectral, or in any other appropriate dimension. In the natural environment the edges may correspond to

- A spatial boundary between materials,
- A shadow from illumination of topographic features,
- Different temperatures in an infrared image,
- Temporal change between video frames,
- Frequency transition due to absorption, reflection, emission in the spectral domain.

The weights for regions of interest would be high for those places where edges exist - where the data is changing values the most. For transmission the edges would be transmitted first. For image sequence processing those areas where something has changed would be highlighted. Edges become less useful when highly textured data is present, leading to more sophisticated measures as alternatives or to supplement edge measures.

The selection of measures should be made with an eye towards to the use. For visualization the measures should be oriented to the types of displays and manipulations supported. The field of visualization is evolving [8]. Given this together with the discovery processes we wish to support, we recommend that a system should provide many measures that the user can select from. Tailoring will be a key feature of the better exploratory systems, based on our experience.

We conclude that the most useful measures are those that can be computed locally, that are computationally efficient, and that best fit the final use. Given measures we now explore how they may be used.

3 Automatic Techniques

Automatic techniques for focussing on regions of interest are useful for two reasons. The first is technical, and the second functional. Often the speed at which data can be presented to the user is restricted because of technical limitations. Communications bandwidth and graphical display speeds are two areas in which limitations are commonly encountered. Progressive transmission (or refinement) has been proposed as a means to ameliorate these limitations [9, 11].

Briefly, progressive transmission encompasses a variety of techniques used to reorganize image data so as to initially present a coarse approximation of the image which is then progressively refined to higher and higher resolutions. The low-resolution image approximations may be generated by data reductions in a number of domains, such as spatial resolution, greyscale quantization, frequency components present in the image, or a combination of these. Each reduction step produces a reduced image approximation and a *residual* image. The transmission sends the most reduced approximation first, then the residual data needed to produce the next approximation. Subsequent residual transmissions generate successively closer approximations to the original data. The residual data may be quantized and compressed to conserve transmission bandwidth, at the possible cost of precise image recovery.

In some implementations, the residual data for one approximation need not be completely transmitted before commencing with the next. Rather, data is sent and displayed where it produces the best effect in the resulting approximation [3, 1]. Most progressive transmission techniques have been implemented for use with image data, though experiments have also been performed with transmission of three-dimensional data [2].

In addition to technical limitations, there are also functional reasons to pursue automatic techniques of focussing on regions of interest. The data can often be so voluminous that the very quantity overwhelms the user, making it impossible to effectively identify and select those portions which are significant. In this case, progressive refinement may be employed as a means of *data abstraction* in which the data is initially presented to the user in summary form and the user permitted to select regions to view in more detail. It is neither desirable nor necessary to have the amount of detail present in the summary data be predetermined at some fixed value. The ideal situation is for the user's display to first present an extremely general picture which gradually increases in detail as time passes. The key factor is to give the user additional data in response to his needs and interactive input.

The user's needs may be roughly determined by the measure of importance which he has selected as being significant. Detail should be added first to those portions of the display which are most important by that measure. The importance measure thus serves as a *priority function* with which to rank regions of the data space for refinement. The refinement becomes equivalent to a hierarchical search for detail guided by the specified priority function.

The user should also be allowed to modify the selected priority interactively, for a number of reasons. The overall importance measure he has selected may not capture all of his criteria, and he may wish to provide additional direction to the search as it is in progress. His criteria may change as he notices additional features in the data. Indeed, he may not have a clear idea of what constitutes an important feature in the data, and should be allowed the freedom to explore the data space interactively. We permit the user to indicate with a pointing device the regions which should be refined to greater detail (*cf.* [5]). This refinement is performed concurrently with that dictated by the priority function, so that other areas of importance are not inadvertantly passed over.

The question of what representation to use for the summary data has been explored in the literature. Any number of filters have been proposed for generation of the reduced resolution images used in progressive transmission. For our purposes, reduction in spatial resolution is most appropriate. Features of interest are generally described using terms of spatial significance such as size, shape, color, or location rather than frequency or phase, and even with special hardware assistance the transformations to and from the frequency domain are computationally expensive. We have adopted Knowlton's *comp/diff* approximation to averages of pixels over a region as the best general-purpose reduction algorithm, because it combines the error-minimizing characteristics of the average with the perceptual enhancement of contrast stretching at minimal computational and storage cost [6, 5]. The one-dimensional encoding described may be extended to multiple dimensions by coding and reducing pairs of elements in each dimension in turn.

Vector-valued data and time sequences pose additional problems to the process of encoding and refinement. Vector component elements may be encoded individually, but to avoid aliasing the refinement process must ensure that all elements are displayed at the same resolution. Similarly, there should be a coherence among the successive frames of a sequence. If the user has manually indicated a point of interest in one frame, that location in the data space should be given priority in subsequent frames as well.

4 Manual Techniques

Much of our work has been on exploring phenomena in remote sensing earth science images through mostly visual presentations. The evolving art of three-dimensional visualization (x, y, z) has established high expectations for computer graphics, and provides the basis for manual techniques for identifying regions of interest [4].

Given that the control is with the human, the real-time loop is the data is presented, the human points to a region (using any interface approach), and the system enhances that region. The enhancement directly relates to resource use - disk access, algorithm computation, memory allocation, transmission bandwidth allocation. Our work uses a projection model, mapping N-dimensions into two display dimensions (x, y) and one frame dimension (z). The x, y dimensions are manipulated with the usual image processing techniques - pan, zoom, rotate, etc. The z dimension is used for time series analysis (movie loops), arithmetic comparisons (subtract, add, average, fade-in, composite), and spectral analysis. The dimensional mapping can be changed at any time (which of the N dimensions map to x, y, z). Coordinated display of multiple images is provided for.

Following selection of the display axes, transformations may be applied to the data to produce enhancement effects. One of the most useful is an interactive color lookup table manipulation system which permits enhancement of specific values or ranges within the displayed data. This capability is particularly useful for locating discrepancies or change when combined with functions which difference the data in successive frames.

5 Examples

Figure 6 illustrates a portion of a four-dimensional data set containing information on ozone concentrations over the south polar region. The four dimensions are latitude, longitude, pressure (altitude), and time. The first image in the sequence shows a south polar topographic image over which the ozone data is placed. Subsequent images show the concentration of ozone at a fixed pressure over a period of months in the fall of 1985. The sequence graphically illustrates the depletion of ozone over time and the widening of the south polar hole.

Figure 2 also shows ozone concentrations, but from two different sensors. The first image was generated from data collected by the TIROS Operational Vertical Sounder (TOVS) sensor, while the second came from the Total Ozone Mapping System (TOMS) sensor. Both depict average concentrations of ozone for the month of August, 1985. The third image in the sequence is the difference between the two sets of measurements, and points up a discrepancy between the two sensors in the south polar region which indicates a need for additional sensor calibration.

6 Future Directions

We continue to extend our prototype system to incorporate more sophisticated visualization and selection tools. We are investigating navigational aids for control over positioning within and traversal of the n-dimensional data space. We envision more flexibility in allowing the user to incorporate customized transformations and analytic filters for the data prior to selection and display.

Future systems will support data types in addition to simple n-dimensional raster arrays. Map, vector or sparse matrix data, for example, can be displayed as overlays on other data to provide reference and boundary information.

A possible extension to the automatic refinement might permit important features to be enhanced at the coarse resolutions using a reduction function which *stylizes* the data rather than simply filtering it. Tanimoto [10] describes desired characteristics of stylizing functions.

References

[1] Ronald P. Blanford. Progressive refinement using local variance estimators. *IEEE Transactions on Communications, to appear.*

[2] Indranil Chakravarty and Bruce G. Nichol. A hierarchical technique for the compression and incremental transmission of three-dimensional digital data. In *Proceedings of the SPIE Conference on Methods of Handling and Processing Imagery (757)*, pages 16–22, Los Angeles, California, January 1987.

[3] Howard M. Dreizen. Content-driven progressive transmission of grey-scale images. *IEEE Transactions on Communications*, COM-35(3):289–296, March 1987.

[4] R. Friedhoff and W. Benzon. *Visualization: The Second Computer Revolution.* H.N Abrams, Inc., New York, 1989.

[5] F. S. Hill, Jr., Sheldon Walker, Jr., and Fuwen Gao. Interactive image query system using progressive transmission. *Computer Graphics*, 17(3):323–330, July 1983.

[6] Ken Knowlton. Progressive transmission of grey-scale and binary pictures by simple, efficient, and loss-less encoding schemes. *Proceedings of the IEEE*, 68(7):885–896, July 1980.

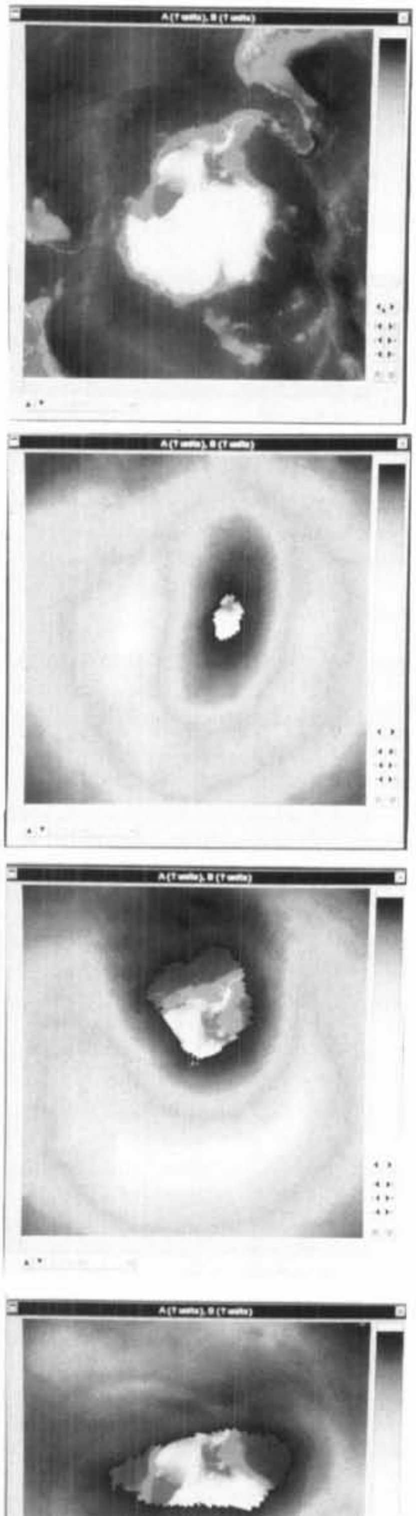

Figure 1: Time series showing ozone concentrations over the south pole.

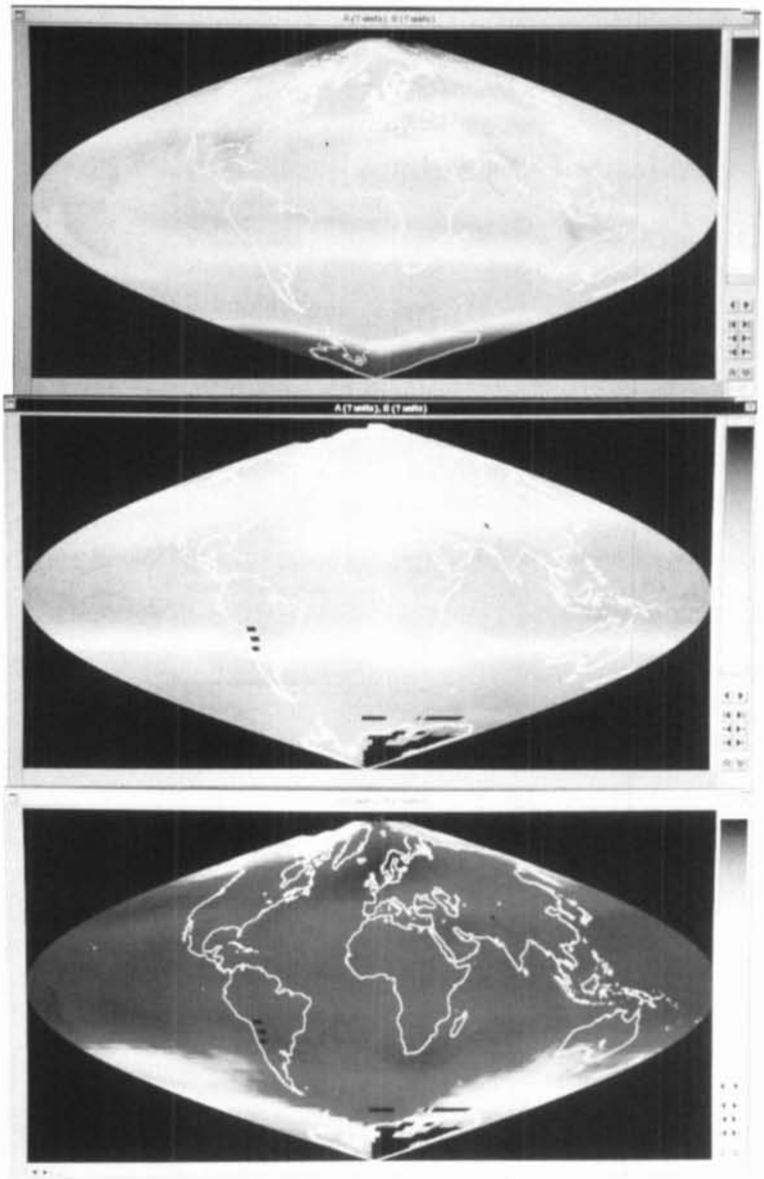

Figure 2: Ozone sensor comparison.

[7] D. Marr. *Vision*. W.H. Freeman and Co., New York, 1982.

[8] B. H. McCormick, T. A. DeFanti, and M. D. Brown, editors. *Visualization in Scientific Computing*, volume 21(6) of *Computer Graphics*. ACM SIGRAPH, November 1987.

[9] Steven L. Tanimoto. Image transmission with gross information first. *Computer Graphics and Image Processing*, 9(1):72–76, January 1979.

[10] Steven L. Tanimoto. Stylization as a means of compacting pictorial data bases. *Policy Analysis and Information Systems*, 3(2):67–89, December 1979.

[11] Kou-Hu Tzou. Progressive image transmission: A review and comparison of techniques. *Optical Engineering*, 26(7):581–589, July 1987.