IMAGE NAVIGATION FOR EMERGENCY RESPONSE (INFER)

Loey Knapp, Senior Systems Analyst, IBM, U.S.A.
John Turek, Research Staff Member, IBM, U.S.A.
Patricia Andrews, Project Leader, U.S. Forest Service, U.S.A.
Christopher Elvidge, NOAA

Commission II, Working Group 3

KEY WORDS: Developing_Countries, Environment, Forestry, Archiving, Navigation, Technology

ABSTRACT

The technology for collection of data through satellite imagery has made rapid strides over the past decade and projections for the size of image data archives in the next decade are overwhelming. The technology to translate these images into information that is usable, easy to access, and timely has lagged that of collection, resulting in a data processing and distribution bottleneck. This paper addresses developments in the provision of image data over the Internet and associated compression, image search, and user interface technology, all aspects of the process to translate data into useful information. Application to emergency management in general, and wildland fire in specific, is discussed.

INTRODUCTION

Efficient prediction, detection, and response to emergency situations such as wildland fires requires a substantial database from which to draw essential information, for example vegetation coverage and fuel type, adjacent structures and urban areas, and evacuation or attack crew transportation corridors. Assessment of fire potential is enhanced by accurate information on seasonal dryness and change in forest composition. For effective planning and communication this information must be up to date and accessible to a wide range of users. To be useful on a global scale, additional issues arise such as access time, cost, and cross-country applicability.

Spacecraft imagery has long held the promise of global data acquisition; by providing up to date global coverage this data should provide an effective means for augmenting the information available through other databases. The large volumes of data and the physical volume of the data which is currently being collected from a multitude of platforms is already enormous and projected to grow ten- fold in the next five years. For example, the instruments on the first two Earth Observing System (EOS) platforms expected to be launched in 1998 and 2000 will generate data at a rate of over 281 GB/day with most of this data taking a form that can be represented as digital imagery. To date the image analysis process has proved to be a bottleneck with the result that the amount and quality of information extracted from this data has not fully met expectations. The challenge is to provide solutions in the areas of data analysis, storage, retrieval and dissemination that will help alleviate this problem.

This paper describes the research taking place under the Image Navigation for Emergency Response (INFER) project to address these issues. INFER, a subproject under the G7 Global Emergency Management Information Network Initiative (GEMINI) is a collaborative research effort comprising IBM, the USDA Forest Service Research, and the Canadian Forest Service Research. The purpose of the INFER project is to apply new techniques in signal processing, visualization and data management to construct an environment for interactive navigation and exploration of large spacecraft image data sets over a network and within the specific domain of wildland fire mapping and response.

The INFER research is addressing the need to build an infrastructure that can effectively support search, storage, retrieval, and transmission of non-traditional data, e.g. images. Unlike conventional text-based digital libraries and databases, search of image libraries cannot be realized simply through the search of text annotations, or metadata. This is because image data is extremely rich in detail and it would be difficult to provide for automatic annotation of each image without human intervention. The INFER team is exploring mechanisms for extracting meaningful information from images through content-based search. Important components of this work include: 1) the implementation of a variety of content-based search algorithms, e.g. texture and shape; 2) application of these algorithms to compressed data, and 3) data access and dissemination over the Internet. While this technology may eventually have value in many problem domains, the INFER project will focus on the development of the content-based search techniques towards the identification and mapping of wildland fires.

TECHNOLOGY

Content-based Search

The typical search approach to a large image archive is through the metadata. The flexibility of the search depends on the extent of the metadata but common search keys are latitude/longitude, date, and path/row. Metadata information can be extensive but it must be created a priori to the search. Thus searches based on metadata keys work well if the ways in which users want to explore the data have been anticipated. However, there is little room for ambiguity and the desire to limit storage of metadata may radically limit search alternatives.

Our content-based search, in contrast to straight metadata search, is based on a set of algorithms which are implemented at the time of the query and which provide the user with a set of tools to search the data in a more flexible manner. The algorithms are based on an examination of the pixel (or pixels) characteristics rather than on pre-collected information. Thus we can perform content search based on texture, shape, pattern recognition, or, even more simply, on classification algorithms which are run at the time of the query.

For instance, the user may be interested in searching an image archive for evidence of increasing urbanization as indicated by new man made structures. It is possible, though unlikely, to indicate in the metadata of each image that man made structures have or have not been identified in that image. It is even more unlikely that the metadata would include information on which of these structures represented a change from some base period. Such questions are hard to anticipate.

An alternative, content-based, search approach could utilize a combination of pattern-matching and change detection algorithms to determine areas exhibiting new square or rectangular shaped features. Changing the query to select areas of 'extensive' change (a major change to the metadata list), an additional aggregation algorithm might be employed.

As opposed to conventional databases, content search on image libraries cannot be realized only through simple search of text annotations. The problem is that image data are rich in detail and it is difficult to provide automatic annotation of each image without having either some form of human intervention or a set of well defined models describing the domain in which queries are to be run. Thus, content-based query of an image database will require the following two steps:

- Extraction of relevant features from each image through the use of appropriate models.
- Determining whether the combination of features extracted from the image represent the content for which the user is searching.

The first step, namely extraction, can be done either at data ingest or dynamically at run time. In the former case, the extracted features are compiled into feature vectors for each of the images and stored as metadata in a database and the content search proceeds by searching through the stored vectors.

Extraction at ingest is frequently much more efficient than extraction at run time as one can make use of multi-attribute indexing techniques to search on the feature vectors defining the images. However, this approach is not a panacea. First of all, the feature extraction process is by its very nature lossy as the feature vector cannot represent all of the content contained within the image. Furthermore, the processing required to derive the feature vector can be quite expensive and at times impossible; for example, the template matching scheme described below can only be computed at run time unless the template is known a priori (an unlikely event). Finally, the indexing techniques used to store the feature vectors tend to be application-specific and do not typically scale well with a large number of pre-extracted features. Thus, the feature-based databases that are being built today tend to be tailored to a specific domain.

It is our thesis that, although useful and necessary, the use of a predefined feature vector cannot adequately support content-based search. It is necessary to provide the functionality that will allow the user to visualize, define and extract features dynamically thereby performing content-based search directly on the image data.

Content Search Operations

Our system implements a set of image operators that can be used as building blocks to synthesize higher level semantics specified by the user. In this section we briefly describe the three operations that currently provide the bulk of our "general purpose" content search mechanism.

One of the fundamental methods for detecting objects within an image is template matching whereby a template of size n x n is compared pixel by pixel with each n x n subimage. The objective is to find those regions having minimal difference from the template. Typical applications of this kind of mechanism are cross registration of two images for visualization and analysis purposes and detection of a given scene from unregistered images. Template matching is rarely exact as a result of image noise, quantization effects and differences in the images themselves. Seasonal changes alone introduce effects that make the matching process difficult. Thus, additional mechanisms are required to ensure adequate search capabilities.

Texture is frequently used to describe two dimensional variations of an image with a characteristic repetitiveness and is a good candidate for classification and feature recognition in a subimage devoid of sharp edges. By using a taxonomy of texture features or by providing examples, the user can define the information of interest in the image.

Classification of a multispectral image is the process of labeling individual pixels or larger areas of the image according to classes defined by a specified taxonomy. This kind of classification is typically used to generate land cover classifications. We extend this approach by providing two additional extensions:

- We allow the user to dynamically define training classes and perform the classification in real time. Thus, the user can define classes not typically covered by standard classification techniques.
- We allow the user to assign information other than the spectral bands. For example, by allowing the incorporation of texture information into the training process the user can define content that cannot otherwise be extracted from the image.

We are in the process of incorporating other capabilities such as shape (from segmented regions) analysis and specification of spatial relationships into our system. These will be described in a later paper.

Compression

While the price of storage devices continues to drop at a dramatic rate there is no doubt that the major cost of providing a digital library will continue to be in the storage devices. Thus, a reduction of even 30%, by the use of compression, in the storage required by the system results in a significant reduction in the overall cost of the system. At first blush it might seem that the cost associated with processing is higher when images are stored in a compressed format; however, this need not be the case. If the data are organized and stored in the appropriate fashion, the search process can be made more efficient and we do not need to analyze all of the information stored in the image in order to eliminate it from consideration. One of the primary ideas of our project is that it is possible to increase the speed of searching through images stored in a digital library while simultaneously reducing the storage requirements. The remainder of this section goes into more details on this topic.

Compression techniques can be either "lossless" or "lossy". A lossless compression scheme is one which guarantees perfect reconstruction of all of the bits in the original dataset. A lossy compression scheme, on the other hand, does not reconstruct most images exactly but rather allows the loss of some information in order to achieve higher compression ratios. The end-user requirements determine whether lossy or lossless compression schemes can be used.

As discussed below, our analysis scheme takes advantage of the properties introduced by lossy compression. However, remote sense scientists are reluctant to lose any data and, therefore, demand lossless compression. The compression scheme we use for image storage offers a hybrid solution. In order to address this issue, we employ a wavelet-based scheme that allows us to progressively extract image content by specifying both spatial and spectral constraints. Typically, as we relax the spectral constraints we tighten the spatial constraints. However, if one fully relaxes both constraints the output is a lossless representation of the original image. The overall compression ratio achieved by this scheme is competitive with the best lossless compression schemes we have analyzed to date.

In general, applying query and retrieval operations directly on lossily compressed data leads to improved computational efficiencies along two fronts:

- One needs to process fewer bits;
- The features and properties of the data are emphasized by the transformed-based compression.

Query operations including retrieval, evaluation, transmission and visualization of the image (or video) data can be staged progressively, by selectively and adaptively processing limited amounts of information, to minimize the total execution time. The difficulties that get introduced are twofold:

- As the number of coefficients is reduced, dimensionality of the search space is reduced, resulting in many-to-one mapping. Thus, the number of false hits increases.
- The reduction of coefficients implied by the compression results in alignment errors. The net result is that the output produced by the operation may not match the exact value produced by operating directly on the original image.
- We have developed techniques that, for several operations, allow us to quantify the latter effect without having to convert back to the spatial domain. This allows us to guarantee that the results of our filters are identical to those produced by operating directly on the full image.

Access via Internet

Access to and dissemination of the data in the INFER project is provided via the Internet so data management techniques which improve access, search, and retrieval time, interactive visualization techniques, and a Netscape-based user interface are all integral components of the system. While each of these components could be discussed at length, this paper will focus on only one - that of the need for a more complex, or 'smart' user interface in an Internet context.

Through development of the INFER prototype it has become clear that an extremely critical component of an internet application is the user interface. In a local software system, the user interface must be intuitive enough for a user to navigate easily through the system functionality but, in general, a user working with in-house applications has an understanding of the underlying data.

As we move to Internet-based applications it is important to consider the fact that the users are, in effect, logging on to a black box, with little or no knowledge of either system functionality or server databases. In the case of providing a data archive search capability it is essential that the user interface assist the user in understanding both the search functions and the extent of the data available for search. In effect, the result must be a 'smart' user interface.

In a simple browse of a large data archive this means supplying the user with information on the spatial and temporal extent of the data. In an application scenario the problem becomes much more complex. For instance, in the context of a content-based search the user must be able to understand and specify the following individual components of the search and their interrelationships:

- Possible search features or search tools - this can be equated to the 'fields' in a database.
- The data sets available for search, e.g. AVHRR, DMSP.
- The temporal extent of the data sets.
- The spatial extent of the data sets.

Providing this information to the user can present a user interface problem as user selections will progressively change the nature of the available data. For instance, specific search features may only be supported in a subset of the data; by selecting to search on that feature the user must be notified of the potential spatial and temporal extent of the search. On the other hand, beginning a search by selecting a spatial area may constrain the user to a subset of the search features. Essentially, wherever the user starts the search the information relative to the other three components may change and notification is required.

The user would also be interested in knowing what ancillary data sets are available for background information or for complex searches. For instance, the user might want to view the results of a search for fires relative to topography or may want to actually constrain the search by an additional data set such as 'fires in grassland areas'.

Unlike SQL, the result of the image query is a spatial display and, depending on the type of query, the user would also require information on how the data could be displayed. The user interface must provide the user with the relevant visualization options. For example, a temporal display could result in a series of sequential images, a movie format with forward and reverse buttons, or a single display of multiple small images.

Provision of this complex array of information which changes depending on the user selections is a user interface challenge. One option is to base the user interface on a large set of hard-coded, but regularly updated, tables which prescribe the interrelationships across the data and between the data and the display. This option presents a system administration problem. The alternative is to make the user interface 'smart' through a set of database calls, thus providing up to date information for every query. For instance, if a user selects to query data set A, the system will search the database for the temporal extent, spatial extent, and potential search features in data set A and inform the user of the result. This option is compelling but may have significant performance problems.

APPLICATION OF TECHNOLOGY

INFER Problem Scenario: Wildland Fire Application

IBM and U.S. Forest Service research personnel are jointly examining the application of these technologies to the area of wildland fire through the INFER project. The premise of the research is that effective response to fire is dependent on up to date, accessible data for prediction, planning and resource allocation. Typically, however, relevant data such as maps of existing regional fires, fuel type or vegetative greenness, and maps of populated areas that might be threatened, are not easily obtainable nor up to date. The challenge in the area of emergency response is to provide the technology to extract crucial information from images quickly enough to influence the decision-making process.

Wildland fire management in the United States is the responsibility of various agencies ranging from federal and state to rural and private. Fire, however, doesn't recognize administrative boundaries. This has led to interagency and international cooperation in detection and response to wildfires. In some cases, dispatchers and coordinators from various agencies are co-located to facilitate cooperation and information sharing. Nevertheless there is room for improvement in tracking fire activity on a national or regional level across all land ownerships as well as in archiving historical fire data. National or regional mapping of fire locations from satellite data through INFER is useful information in that it will be up to date and easily accessible across multiple agencies; that it can be combined with other ancillary data such as topography, urban areas, and administrative boundaries; and that it can be used in conjunction with fire potential models. It is anticipated that this work might also be extended to early fire detection applications, especially in countries where a fire detection infrastructure is not in place.

INFER System

The initial queries supported in the INFER system relate to fire location at a regional and national level. The data set used in the prototype is the Defense Mapping Satellite Program (DMSP) Operational Linescan System (OLS) imagery which provides daily global coverage with a spatial resolution of eithe r.5 or 2.5 kilometers. The National Oceanographic and Atmospheric Administration (NOAA) National Geophysical Data Center (NGDC) has developed forest fire classification algorithms for the OLS data which rely on detecting areas of active visible and near infrared emission on the planet surface at night, when solar illumination is absent. Intensification of the light emissions results in a data set in which city lights, gas flares, lightning, and fires can be observed.

The purpose of the INFER project is not to develop new classification algorithms but to implement and extend existing algorithms such that they can be used as image search tools. The fact that NGDC had a process for preprocessing the OLS data and extracting fire locations made this data set particularly attractive for the INFER prototype.

The process for locating fire in the INFER system is illustrated in Figure 1.

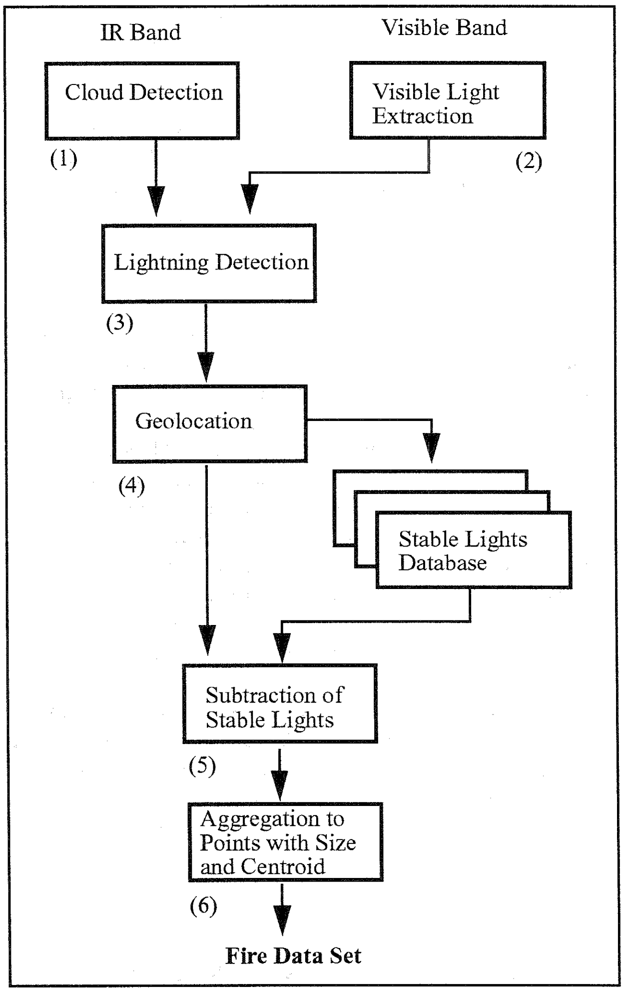

Figure 1: Process for extracting the location of current fires from the DMSP OLS data.

- Step 1: Clouds are automatically detected within the image and a mask is generated. Clouds provide three potential sources of interference. (1) Clouds can cover regions containing active fires and reduce the likelihood of location. This will provide useful ancillary information post processing. (2) The cloud mask is used as input for the lightning detection process. (3) Moonlight reflecting off cloud edges causes glint which may be mistaken as active VNIR emissions.

- Step 2: The visible band of the OLS data is filtered to generate a raster representation of the light sources in the image.

- Step 3: The mask from Step 1 and the raster image from Step 2 are used to identify the lightning sources that are then eliminated from the raster source.

- Step 4: The light sources resulting from Step 3 are geolocated. The result of this step is added to enhance the stable lights image for future dates.

- Step 5: The stable lights database generated from previous dates is used to mask out inhabited areas from the current pass (results of step 4) and the remaining lights are labeled as intermittent lights, or fire.

- Step 6: Aggregation of the remaining light pixels results in the reporting of a centroid and a size (i.e., number of affected pixels).

While portions of this process were put in place by NGDC, full automation has been achieved through the INFER project. Both the cloud and lightning detection steps were previously processed manually; automation of these steps has substantially cut processing time.

With the fire algorithm in place, the user can query the data in a variety of ways. First, the user can request information on current fires at either a national or regional level. These can be viewed relative to a variety of ancillary data such as topography, vegetative greenness, or land cover type. Second, a historical query could be made through which a user could request an animated series of fire images over a specified time period, two dates could be compared to determine relative fire levels, or fire images could be compared to historical data bases. Third, the subproduct of the fire location process, the stable lights filter, could be queried. For instance, an emergency response manager might be interested in a comparison of city lights before and after a disaster to determine area and extent of damage.

Additional search algorithms in the INFER system currently include land cover and burn scar. Scenario two will focus on the local needs pertaining to a specific fire event. Pattern and texture searches on higher resolution data will allow users to query on values at risk, e.g. man made structures, access roads, or land cover type.

Usage scenario

We describe here a scenario of how INFER technology might have been used during a period of high fire activity in the Northern Rockies in August of 1994 (Bradshaw and Andrews in press). A Multi-Agency Coordinating Group (MAC Group) was activated due to extreme burning conditions that posed serious threat to life and property throughout much of western Montana and northern Idaho. The role of a MAC group is to evaluate the fire situation and set priorities for use of limited fire suppression resources (e.g. crews, equipment, and aircraft); they are not involved in the management of individual fires. We expect that INFER products would have supplemented, not replaced, information available from other sources.

During each day's MAC group meeting, members would have used INFER to view the current fire activity for the entire U.S. (satellite imagery in raw pixel form with state boundaries for reference). Verifying that most of the activity is in their area of responsibility and that there won't be additional competition for resources, they would concentrate on (zoomed into) the area they had predefined as their area of concern. Administrative boundaries (e.g. Forest Service, National Park Service, State) would be added for reference and a query made as to the number of fires under each administrative jurisdiction. These numbers can aid in the assessment of overall fire activity even though the actual values may not agree with those from other sources due to limitations of imagery and delays in reporting and compiling. The number of fires as identified by satellite would be readily available for comparison to previous days.

Ongoing fires could be examined with respect to fire potential through models available through the Wildland Fire Assessment System (WFAS) (Andrews and others in press). In setting priorities, the MAC group could query for those fires in areas identified as extreme fire danger.

Using INFER, the MAC group would be able to zoom in to look at higher resolution data for areas of special interest. There was concern about Little Wolf Fire because of its location with respect to a new development. High resolution data on fuel, topography, roads, and structures would be used in conjunction with weather forecasts and fire growth simulation models to project fire growth, leading to a conclusion that the fire would pose significant threats to life and property only in the event of strong west or northwest winds.

The MAC group would choose some images to show at a briefing later in the day, where the audience included members of the press as well as agency personnel. In addition to showing the area of high activity in this part of the country, additional images would show how the season had progressed to date and how it compared to other severe fire seasons in recent memory.

EXTENSION OF TECHNOLOGY

Although this paper has discussed the access, search, and distribution of image data via the Internet in the context of wildland fire, the applicability of content-based search is much broader. As the use of the Internet for commercial purposes increases and the resolution of satellite imagery becomes finer, e.g. one to two meter, it is possible to imagine a wide array of application areas. Emergency managers might search on oil spills or flood areas, transportation managers on road or bridge damage, and resource managers on vegetation health or weed infestations. The key to these scenarios is ready access to data in usable form and within the individual time constraints. In other words the technology for data use must keep up with or exceed that of data collection.

ACKNOWLEDGEMENTS

This work is partially supported by NASA contract NCC5-101.

REFERENCES

Andrews, Patricia L., Larry S. Bradshaw, Robert E. Burgan, Carolyn H. Chase, and Roberta A. Hartford, 1995 (in press). WFAS: Wildland Fire Assessment System-Status 1995. 1995 Interior West Fire Council Meeting, St. George, Utah.

Bradshaw, Larry S., and Patricia L. Andrews, 1994 (in press). Fire Potential Assessment during a Period of High Fire Activity in the Northern Rockies: August, 1994. Interior West Fire Council, Coeur d'Alene, ID.

Castelli, V., C.S. Li, J. Turek and I. Kontoyiannis, 1996. Progressive Classification in the Compressed Domain for Large EOS Satellite Databases. Proceedings of 1996 IEEE International Conference on Acoustics, Speech and Signal Processing.

Elvidge, Christopher D., Herbert W. Kroehl, Eric A. Kihn, Kimberley E. Baugh, Ethan R. Davis, and Wei Min Hao, 1995 (in press). Algorithm for the Retrieval of Fire Pixels from DMSP Operational Linescan System Data. Global Biomass Burning, edited by Joel S. Levine.

Elvidge, Christopher D., Kimberly E. Baugh, Eric A. Kihn, and Ethan R. Davis, 1995. Mapping City Lights with Nighttime Data from the DMSP Operational Linescan System. Submitted to Photogrammetric Engineering and Remote Sensing.

Knapp, Loey, Patricia L. Andrews, and John Turek, 1996. Image Navigation for Wildland Fire Location Mapping. Proceedings of the GIS and Environmental Modeling conference, Santa Fe, New Mexico.

Li, C.S., V. Castelli and C. Elvidge, 1996. Automatic Tracking of Land Cover Transformation through LANDSAT Multispectral Scanner (MSS) Images Using Neural Network Classifiers. 1996 SPIE International Conference.

Pennebaker, W., and J. L. Mitchell, 1993. JPEG Still Image Data Compression Standard. Van Nostrand Reinhold. New York.

Rioul, O., and M. Vetterli, 1991. Wavelets and Signal Processing. IEEE Signal Processing Magazine, vol. 8, n. 4, pp. 14-38.