Development of Thematic Browse Products to Aid in the Analysis of National Satellite Image Data Sets

B. Guindon
Canada Centre for Remote Sensing
588 Booth Street
Ottawa, Ontario, CANADA
K1A 0Y7
Tel. (613) 947-1228
Fax (613) 947-1383
Email bert.guindon@ccrs.nrcan.gc.ca

Abstract

Currently, there a number of remote sensing programs underway that will lead to the creation of comprehensive national and global data sets of processed, high resolution satellite scenes. While these data sets represent a sampling of only a fraction of all scenes held in raw image archives, collectively they still constitute a massive volume of potential information beyond the capabilities of users to digest. This paper proposes the development of a new form of browse product to accompany these full resolution data sets. The purposes of this browse are to support extraction of synoptic information measures and to quickly locate example areas illustrating user-specified thematic occurrences. To accomplish these goals, a browse product consisting of coarse thematic layers is proposed. The structure of a complementary browse analysis system is also described which builds upon content-based image retrieval (CBIR) currently under development for generic image databases.

Background

Rapid improvements in computer storage, processing and graphics as well as the development of the Internet has made on-line access to large image libraries a reality. The content of such digital image databases ranges from photographic records of art and museum holdings to archival satellite imagery of the earth while the databases themselves may contain hundreds of thousands of distinct data items (individual images).

Currently, a key challenge is the effective exploitation of these databases by a broad audience whose main interests and expertise lies in the thematic content of the image files not image processing or computer technology. This area of research, referred to as content-based image retrieval (hereafter abbreviated to CBIR), has led to the development of a number of proto-type retrieval systems including, for example, Photobook (Pentland et al., 1996), QBIC (Flickner et al., 1995; Ogle and Stonebraker, 1995), SOM-AIR (Zhu et al., 2000) and I.Q. (Agouris et al., 1998). With the exception of the last two named cases, most studies have addressed generic content extraction and retrieval issues.

The generic structure of a typical CBIR system is shown in Figure 1. While the systems listed in the previous paragraph differ significantly in detail, all include three principal processing modules that support content extraction, interactive querying and a matching.

Content Extraction

Traditionally, content extraction has been accomplished through visual interpretation with the results stored in the form of textual meta-data. A simple example from the early days of remote sensing is the estimation of percent cloud cover in satellite images. For very large, dynamic databases, manual processing is no longer feasible and automated techniques have been sought in which image processing techniques are employed to extract image primitives from which content-related cues can be inferred. These primitives and cues may be saved as raster overlays or summarized in the form of textual meta-data. Numerous primitives have been proposed and incorporated in systems including;

(a) texture measures (in QBIC, Photobook and SOM-AIR),
(b) hierarchical descriptions of image regions through segmentation (Fuh et al., 2000; in Photobook and QBIC),
(c) wavelet/moment histogram descriptors (Mandal et al., 1998),
(d) Hough and frequency domain transforms to encapsulate object orientation (Celentano and Di Sciascio, 1998; Ben-Arie and Wang, 1998),
(e) Local spectral and spatial 'interest points' within images (Schmid and Mohr, 1997; Ramesh and Sethi, 1995),
(f) 'eigenimage' representations based on Kahruen-Loeve expansions (Pentland et al., 1996).

A number of research studies have been directed at remote sensing content issues including (a) the development of physics-based, multi-spectral invariants to address problems related to atmospheric and seasonal differences between satellite images (Healey and Jain, 1996) and (b) an investigation of Gibbs-Markov random field models to describe SAR image spatial structure (Schroder et al., 1998).

Query Interface

The goal of most CBIR systems is to efficiently locate and extract example data items that meet some query criteria. Three common forms of query are;

(a) textual. These requests involve keyword and parameter searches of meta-data files.

(b) search by example. The operator presents an image to the system and requests that 'similar' images be located. Searching can be done at the meta-data or raster levels.

(c) search by sketch. In this case the operator generates a 'sketch' of an object or configuration of objects. The system then extracts the salient geometric and radiometric cues from the sketch and uses these as selection criteria.

In the latter two cases, the system content extraction tools are applied to the example image or sketch to ensure consistency with the content characterizations of the database members.

Match Engine

This is the process whereby the attributes/cues quantifying the query are compared with those of the database items to come up with a list of candidate matches. There are a number of keys issues related to matching.

(a) Since CBIR is meant to support or replace visual selection, computer metrics of similarity must mimic human judgement of similarity. This does not mean that computation must parallel visual processes but rather that machine and human judgements of similarity be 'correlated' (Pentland et al., 1996).

(b) Complex queries may involve searching a feature space consisting of a diverse set of image attributes. This presents two important challenges, namely, (i) how to relate each attribute to their importance in meeting visual similarity goals and (ii) how to combine diverse attributes within the context of a decision-making framework such as Bayesian reasoning. Most match engines employ simple approaches such as combining attributes into vectors and utilizing vector divergence as a measure of dissimilarity (e.g. Celentano and Di Sciascio, 1998; Androutsos et al., 1999). This approach is tractable in cases where individual attributes can be represented as scalar measures. In more complex cases where spatial context/relationships are relevant, similarity measures must account for these through, for example, template matching (e.g. Smith and Li, 1999) or relational graph matching (e.g. Kitamoto et al., 1993).

(c) An efficient search strategy is needed since a 'brute-force' approach generally is not feasible in terms of response time for large databases. A preferred approach may involve iterative/interactive refinement in which the system returns some initial candidates that are visually assessed by the operator. This assessment, leading to a ranking of successful candidates and a discarding of others is used by the system to further refine match criteria and attribute weights leading to an improved candidate list.

(d) Many performance tests of CBIR systems are based on the number of false matchesreturned (commission error rate) which reflects the query philosophy of the form 'find similar examples of ….'. Unfortunately this would not necessarily be the goal of a CBIR system for a satellite browse inventory where the query requirement couched as 'find all examples of ….' better reflects the need for completeness of examples, i.e. minimization of omission errors.

Application of CBIR to Satellite Image Archives

Satellite image archives encompass large numbers of complex images and would, at first glance, appear to be ideal benefactors of CBIR technologies. Unfortunately, there are a number of reasons to believe the CBIR is insufficiently developed for general remote sensing application.

(a) Satellite data reception and archiving must be highly automated and operational considerations (e.g. data volume and timely reception and dissemination) preclude expensive content extraction at the time of reception.

(b) CBIR is at present an immature technology. Current proto-type systems have been designed for and tested on databases of simple images, i.e. images containing one or at best a few objects. Satellite scenes, on the other hand, cover extensive geographic areas and tend to contain a complex mix of natural and man-made features.

(c) Ground stations archive images in a downlinked format. Significant processing is required to produce useable, calibrated scenes from which content might be extracted. Since such processing is undertaken only on a user-request basis, the population of readily useable scenes only constitutes a small fraction of current archives.

In conclusion, the application of CBIR technologies to satellite image archiving and dissemination activities, while attractive in principle, remains impractical at present. On the other hand, there is a restrictive problem that may benefit from CBIR. In recent years a number of initiatives have been undertaken to identify and process sets of scenes that collectively provide definitive terrestrial temporal and/or geographic coverage. Example programs include NALC (Lunetta et al., 1998; Landsat MSS coverage of the U.S. and Mexico at 3 epochs), MRLC (Loveland and Shaw, 1996; Landsat TM coverage of the U.S.) and GEOCOVER (Landsat MSS and TM global coverage). While these data sets are limited in scope relative to overall Landsat archives, they still constitute immense data volumes in their own right and present real challenges to novice users in terms of manipulation and information extraction.

A tractable application of CBIR in this arena might be in the creation of a ‘content-based’ browse products that would support synoptic thematic-based queries. Traditionally, satellite browse data have taken the form of sub-sampled imagery, automatically generated at the time of reception and made available to users either in hard copy form (e.g. microfiche) or in compressed digital form via the internet. The rudimentary form of conventional browse limits its primary use to that of a visual inspection tool to aid users in assessing synoptic scene quality especially in terms of gauging atmospheric degradation. Some high level goals/characteristics of a content-based browse are summarized below.

(a) Each browse product would constitute an interpretation of a parent scene in the form of a suite of thematic overlays and textual meta-data. For national data sets, the stringent time constraints of current browse generation (i.e. at data reception time) can be relaxed since the creation of a national image data set is a long-term endeavour (typically 2-4 years) with sufficient time for parallel thematic information extraction.

(b) The proposed browse should attempt to capture thematic structure at the spatial resolution of the parent sensor. This differs significantly from traditional browse images that are generated either by a sub-sampling process, leading to aliasing, or by spatial averaging, leading to a loss of spatial detail. This implies that primary content extraction must be carried out on the full resolution parent image and that an alternate encoding scheme is required to produce a compressed thematic interpretation.

(c) Since a primary driver of many of the data set initiatives is the study of landscape change, integration of thematic browse derived from different national/regional image data sets is desireable.

(d) The primary purposes of the new browse would be to allow users to (i) conduct synoptic assessments of large data sets to delineate and quantify large-scale (spanning multiple scenes) thematic trends (ii) to locate specific examples of sub-images where these thematic phenomena are occurring. Thus the proposed CBIR system and browse would form only an initial module of a larger decision-making system that fully exploits the parent data sets.

(e) The match engine component of the associated CBIR system should seek to minimize omission errors (true example completeness) even at the expense of increased commission error levels since false alarms can be detected and discarded during subsequent, detailed analyses. To meet this goal, the interpreted content of the browse must be of high reliability. Given this requirement and the fact that automated content extraction is desirable, the thematic level of interpretation may be of a rudimentary nature. On the other hand, the information content of the browse must still be detailed enough to support synoptic queries related to first order impacts of issues such as deforestation, urbanization, wetland loss, etc,

An Example Scenario

In this section we develop a set of browse specifications that would be applicable for satellite image data sets, generated by moderate resolution (10 to 100 m) sensors, and spanning national coverage. Each parent image is assumed to geocoded with a pixel size comparable to the sensor IFOV. As a specific benchmark, we take the case of Landsat TM imagery. Although the imaging swath width of the sensor is approximately 180-185 km, a typical full scene geocoded product will cover an area of 250 km x 250 km. If we assume a pixel spacing of 25 m (i.e. the value of current operational geocoded TM products), this leads to an image size of 10000 x 10000 pixels per band and a total of approximately 600 scenes for national coverage. While it is true that

Landsat 7 also includes a higher resolution (15m) panchromatic band, thematic information extraction necessary to generate thematic browse is envisioned to require the spectral dimensionality of the conventional 30 m bands.

A reasonable browse product size would be 500 x 500 pixels corresponding to a scale reduction of 20. Unlike current browse imagery, this reduction would be achieved by capturing thematic information in 20 x 20 pixel blocks from the parent scenes and encoding the content of each block as a set of values or ‘bands’. For our example TM scenario each block (i.e. each browse ‘pixel’) corresponds to a 0.5 km x 0.5 km footprint. Since a major application of the browse would be to study large-scale (i.e., multi-scene) trends, it is likely that users would want to have access to the complete set of browse products. If these were to be distributed on a conventional CD medium, one Mbyte of storage would be available per scene or the equivalent of 4 8-bit ‘bands’ of thematic characterization.

Current CBIR systems have been designed to deal with simple images (i.e containing few dominant objects of interest) and to treat constituent images as independent entities, (i.e. no content inter-dependencies). These design assumptions are not applicable to a national satellite image data set. First, since the scenes are geocoded to a common cartographic grid, there is a spatial ordering or linkage between scenes. Second, since the imaging swaths of adjacent satellite tracks overlap in coverage, there is a significant measure of information redundancy especially in the case of high latitude image frames. One approach is to view overlap as undesirable and to eliminate it by merging scenes into a seamless mosaic. We take an opposing view, namely, that overlapping coverage constitutes a powerful information supplement. It can provide supporting evidence or validation for a thematic inference drawn from a scene as well as enhance temporal sampling of phenomena since the parent scenes of an overlap region may have been acquired months or even years apart.

This inter-scene dependence suggests that the image data set structure should be viewed at two levels. At the first level, the overall data set consists of distinct scenes. At the second level, each complex scene can be viewed as a spatially ordered set of image blocks (in our example 0.5 km x 0.5 km in size). This block or ‘browse’ pixel should then be considered to be the fundamental granule of the data set, with each granule having one or more content descriptions depending on the number of images contributing to it. An added advantage of this view is that it naturally supports integrated interpretation of multiple national data sets from different ‘epochs’ or sensors as long as they conform to a common block structure, and, consequently, supports content queries related to long-term temporal issues. It should be pointed out that reduction of images to an ordered set of image sub-areas or blocks has been alluded to in earlier papers (e.g. Celentano and Di Sciasscio, 1998), however, in these instances it was employed to determine spatial context of a dominant object within an image.

Up to now we discussed structural aspects of the browse but not the thematic content to be portrayed in each browse pixel. CBIR systems have notable limitations in their application to satellite data. First, most match engines are designed to find images that are similar to a reference image or sketch. In most cases this can be carried out without explicit reference to the thematic content since the matching involves comparison of image primitives. We contend that in the satellite case, content extraction should lead to a higher level of thematic interpretation such as some form of image classification. Since it is desireable that these explicit content measures be reliable and consistent across a national scale, a rudimentary thematic stratification is called for. Simple measures such as those based on, for example, vegetation indices and/or brightness-greenness since they can be derived by an unsupervised process, however, their subsequent thematic interpretation requires a level of remote sensing expertise beyond that of many users especially those at the policy and decision-making level.

A better target level would be a thematic classification comparable to, for example, Anderson Level 1 (Anderson et al., 1976). Numerous studies (see Lillesand and Kiefer, 1987) have shown that this level can be reliably inferred from Thematic Mapper imagery and that it is sufficiently detailed to directly support synoptic queries on issues such as deforestation, urbanization. There are a number of significant implications to employment of an explicit classification such as Anderson. First, it would require a national network of 'ground truth' sites to support classifier training and validation. On the other hand, such a network would probably needed to support more detailed interpretation of the full resolution data set and this browse requirement would be a subset of that requirement. Second, an encoding scheme would be needed to summarize the class content of each 0.5 km x 0.5 km block. This could take the form of simple measures such as the percentages of pixels in each class or landscape metrics, such as contagion, compaction (e.g. Dillworth et al., 1994) that encapsulate spatial distributions of class content within a block.

Finally, Figure 2 illustrates a proposed structure for a satellite CBIR system. Since content extraction issues were discussed in the previous paragraph, we have omitted this module for the sake of clarity. The system structure can best be understood by considering a typical analysis case. Initially, the user describes a rudimentary thematic scenario of interest through the front-end query interface. Query specification could be accomplished with the aid of a keyword dictionary. This request is converted to a set of attribute specifications that drive the match engine. Matching is carried out on the browse data sets. The match process returns lists of the geographic coordinates of candidate blocks, synoptic summaries of the distributions of these blocks (e.g. in the form of a customized map) and quantitative estimates of the overall phenomenon of interest (e.g. number of hectares of agricultural land lost to urbanization). If the browse is directly linked to the parent image data set(s), full resolution image 'chips' of the candidate blocks could also be presented to the user. Since these chips comprise small data volumes, real-time interactivity over the Internet should be feasible. The advantage of this feature is that it would allow one to exploit a major strength of current CBIR systems, namely, interactive feedback. For example, a rudimentary query may return a large number of candidate blocks including numerous commission errors. If the user then selects example sub-sets that are the most and least representative of the specific case of interest, these can be used to refine the query and improve the synoptic results.

References

Agouris, P., Stefanidis, A. and J.D. Carsewell, 1998, 'Intelligent Retrieval of Digital Images from Large Geospatial Databases', **Proceedings of the 1998 ISPRS Commission III Symposium**, Vol. 32, pp. 515-522.

Anderson, J.R., Hardy, E.E., Roach, J.T. and R.E. Witmer, 1976, 'A Land Use and Land Cover Classification System for Use with Remote Sensor Data', **USGS Professional Paper # 964.**

Androutsos, D., Plataniotis, K.N. and A.N. Venetsanopoulos, 1999, 'A Novel Vector-Based Approach to Color Image Retrieval Using a Vector Angular-Based Distance Measure', **Computer Vision and Image Understanding**, Vol. 75, pp. 46-58.

Ben-Arie, J. and Z. Wang, 1998, 'Pictorial Recognition of Objects Emplying Affine Invariance in the Frequency Domain', **IEEE**

Transactions on Pattern Analysis and Machine Intelligence, Vol. 20, pp. 604-618.

Celentano, A., and E. Di Sciascio, 1998, 'Feature Integration and Relevance Feedback Analysis in Image Similarity Evaluation', **Journal of Electronic Imaging**, Vol. 7, pp. 308-317.

Dillworth, M.E., Whistler, J.L. and J.W. Merchant, 1994, 'Measuring Landscape Structure Using Geographic and Geometric Windows', **Photogrammetric Engineering and Remote Sensing**, Vol. 60, pp. 1215-1224.

Flickner, M. et al., 1995, 'Query by Image and Video Content: the QBIC System', **IEEE Transactions on Computers**, Vol. 28, pp. 23-32.

Fuh, C.-S., Cho, S.-W. and K. Essig, 2000, 'Hierarchical Color Image Region Segmentation for Content-Based Image Retrieval Systems', **IEEE Transactions on Image Processing**, Vol. 9, pp. 156-162.

Healey, G. and A. Jain, 1996, 'Retrieving Multi-Spectral Satellite Images Using Physics-Based Invariant Representations', **IEEE Transactions on Pattern Analysis and Machine Intelligence**, Vol. 18, pp. 842-848.

Kitamoto, A., Zhou, C. and M. Takagi, 1993, 'Similarity Retrieval of NOAA Satellite Imagery by Graph Matching', **Proceedings of the SPIE**, Vol. 1908, pp. 60-73.

Lillesand, T.M and R.W. Kiefer, 1987, **Remote Sensing Image Interpretation**, John Wiley and Sons, New York, 721p.

Loveland, T.R. and D.M. Shaw, 1996, 'Multi-Resolution Land Characterization: Building Collaborative Partnerships', **Proceedings of the ASPRS/GAP Symposium**, Charlotte, North Carolina, pp. 83-89.

Lunetta, R.S., Lyon, J.G., Guindon, B. and C.D. Elvidge, 1998, 'North American Landscape Characterization: Dataset Development and Data Fusion Issues', **Photogrammetric Engineering and Remote Sensing**, Vol. 64, pp. 821-829.

Mandal, M.K., Aboulnasr, T. and S. Panchanathan, 1998, 'Illumination Invariant Image Indexing Using Moments and Wavelets', **Journal of Electronic Imaging**, Vol. 7, pp. 282-293.

Niblack, W., Barber, R., Equitz, W., Flickner, M., Glasman, E., Petkovic, D., Yanker, P., Faloutsos, C. and G. Taubin, 1993, 'The QBIC Project: Querying Images by Content Using Color, Texture and Shape', **Proceedings of the SPIE**, Vol. 1908, pp. 173-187.

Ogle, V.E. and M. Stonebraker, 1995, 'CHABOT: Retrieval from a Relational Database of Images', **IEEE Transactions on Computers**, Vol. 28, pp. 40-49.

Pentland, A., Picard, R.W. and S. Sclaroff, 1996, 'Photobook: Content-Based Manipulation of Image Databases', **International Journal of Computer Vision**, Vol. 18, pp. 233-254.

Ramesh, N. and I.K. Sethi, 1995, 'Feature Identification as an Aid to Content-Based Image Retrieval', **Proceedings of the SPIE**, Vol. 2420, pp. 2-11.

Schmid, C. and R. Mohr, 1997, 'Local Grayvalue Invariants for Image Retrieval', **IEEE Transactions on Pattern Analysis and Machine Intelligence**, Vol. 19, pp. 530-536.

Schroder, M., Rehrauer, H., Seidel, K. and M. Datcu, 1998, 'Spatial Information Retrieval from Remote-Sensing Images Part II: Gibbs-Markov Random Fields', **IEEE Transactions on Geosciences and Remote Sensing**, Vol. 36, pp. 1446-1445.

Smith, J.R. and C.-S. Li, 1999, 'Image Classification and Querying Using Composite Region Templates', **Computer Vision and Image Understanding**, Vol. 75, pp. 165-174.

Zhu, B., Ramsey, M. and H. Chen, 2000, 'Creating a Large-Scale Content-Based Airphoto Image Digital Library', **IEEE Transactions on Image Processing**, Vol. 9, pp. 163-167.

Figure 1. Schematic representation of a generic CBIR system.

Figure 2. Schematic representation of a browse-based CBIR system for national satellite image data sets.