Visual Database Systems
T.L. Kunii (Editor)
Elsevier Science Publishers B.V. (North-Holland)

Model-Driven Image Analysis to Augment Databases

Gio Wiederhold, Jim Brinkley*, Ramin Samadani, and C. Robert Clauer

Stanford University
Departments of Computer Science, Medicine, and Electrical Engineering
Stanford CA, USA

Abstract In this paper we consider how information may be obtained from images. To search large image collections we need to search on secondary parameters. We may look for images containing certain types of objects, for images where the objects are of a certain size or shape, or for images having certain features. Since we now have techniques to rapidly acquire and store many images, we need techniques for automatic image analysis to generate such parameters. This paper describes a promising category of image analysis, namely model-driven methods. Two examples, operating in very different domains, are presented.

1. Introduction

To be able to retrieve images stored in databases we must associate identifying parameters with each of the images. It is through these parameters that we can select stored images for display, comparison, and further analysis. Primary parameters are produced when the images are obtained, and describe the imaging event and its process. For satellite images of earth we would have the time of observation, the setting of the scanner and the satellite, and, immediately derived from those data, we can determine the satellite's path and the portion of the earth seen. For medical images the primary parameters are the name of the patient, the date, and the operator of the scanner, positioning information, and the purpose of the scan.

Any each of those technologies can generate thousands of images. For systematic analysis we need to retrieve images based upon their contents. The pixels which make the images themselves are rarely suitable for direct search. If we can provide secondary parameters indicating the type of objects seen on the earth or the type of abnormality seen in a body, then database searches could be enabled that directly serve the end-users' objectives. Large scale on-line image storage without such a retrieval capability will not be justifiable [Arvidson 1986]. This need for image-oriented search is well recognized, for instance [Chang 1988] presents a suitable iconic query language and shows examples, but does not explain how the icons are being generated from the images.

Today, the majority of secondary parameters, when available at all, are entered by humans after visually scanning an image. There are some exceptions, generally based on a bottom-up analysis of image contents. By bottom-up we mean pixel-based inductive processing, in order to find pixel agglomerations as lines, intersections of lines, or bounded

* now at the Department of Biomedical Structures, Univ. of Washington, Seattle WA.

areas. If the pictures are well-behaved, like the maps presented to PSQL queries, then direct manipulation may be feasible [Roussopoulos 1988]. Analysis of structured images is also feasible [Kasturi 1988].

1.1 Model-based Image Analysis

The alternative, top-down image analysis, starts with a general expectation of what is to be found. The expectation is represented by a model, and if the model cannot be matched, then the object or feature is assumed not to be present. For instance, in high-energy physics, programs will scan automatically for *interesting* events, which generate images that are easily described, as the track angles after an impact. Quality control imaging applications compare images versus a model image of a correct part.

Often the objective of image analysis is to look for variations in objects. The two basic approaches for representing variation are:

1 placing the variation in the matching function

2 placing the variation in the model

The first approach, which is more common, is to devise prototypical instances of a particular shape, then develop a metric for comparing a new shape with a library of prototypical instances. This approach can be very successful and is widely used for image analysis. However, since the metric is usually only computed after all the features have been observed in the image to be classified, the variations found cannot be used to aid the search for image features.

We deal here with images which are variable, so that we adopt the alternative where the model is adapted as features are located. The variation is now represented in the model itself, thus defining a generic model which describes all expected instances of a particular shape class. If a particular object reconstruction can be described by this model then the object is an instance of the shape class. Several computer vision systems attempt to provide model knowledge in order to guide the analysis.

An example of this approach for machine vision is ACRONYM, a model-driven system developed by Brooks for finding airplanes in satellite images [Brooks 1981]. This system used generic models defined by structured assemblies of generalized cylinders as described by [Binford 1971]. By placing ranges on the parameters describing the generalized cylinders it was possible to describe the expected range of variation for various classes of airplanes, thereby building up a classification tree which at the root described all airplanes, and which became progressively more specific as the tree was traversed towards the leaves. The advantage of this approach is that the system could initially be told it was looking for some kind of airplane. Once it had found a few parts of an actual airplane on the image, it could use the dimensions of these parts to specialize the generic model to only those airplanes that were consistent with the image. The cylinders were generalized to cones for organ desciptions [Soroka 1981]. The specialized model generates expectations about where to search for additional parts of the airplane on the image, thereby providing a top-down approach to searching the image based on a-priori knowledge about the contents of the image.

1.2 Obtaining Secondary Parameters

Image analysis based on parameterized models can be expanded to cover a wide variety of objects and features by defining more general models, and appropriate adaptation schemes. The operations needed for matching form the basis for descriptive parameters to augment the database.

To enable subsequent search, at a useful level of abstraction, requires further processing. The findings must be aggregated to form meaningful secondary parameters. Detected features can be used to classify objects found into types and subtypes. For three-dimensional objects volume and aspect ratios can be produced which may be more significant than their lengths, widths, and heights. These secondary parameters become candidates for indexing the database, so that images containing certain object types can be rapidly retrieved [Wiederhold 1987].

In this paper we present two application, from quite different domains, which use abstract models for image analysis. Parameters derived from adapting the model to the image become the indexable attributes to be placed into the database. We believe that such techniques have great promise, and that their application will broaden the usefulness of image databases.

1.3 Problems with Image Analysis.

Images obtained from scanning devices are rarely complete, and also have artifacts. In the work described in this paper we have images of objects, but because of physical constraints the observations on an image are incomplete. Having a model helps to ignore artifacts, but we still must obtain information about the object proper.

If parts of the border of an object are missing or obscure, edge-following algorithms can easily get lost, and reconstruction must use images obtained in a later scan with different viewing angles and illumination. The models we use are intended to overcome this problem in two application domains. In the first, a medical application, the shape model is constrained only by geometric first and second-order continuity. Multiple images, taken at different angles, reduce the uncertainty. In the second, a space sciences application, we use a physical analogy to constrain the model transformations.

Ultrasound Images One of the two applications presented in this paper deals with images obtained in pre-natal medical care, using ultrasound. Ultrasound is one of the least invasive of imaging techniques. It is hence appropriate for delicate tasks, as the imaging of a human fetus, so that abnormalities, typically slow growth, can be diagnosed early. Each ultrasound scan produces a two-dimensional image plane; many such images must be combined to obtain an adequate three-dimensional view of the fetus. The successive images are not evenly spaced, nor all at the same angle, since the scanning head must follow the contour of the maternal body. The components of the scans are also incomplete; in particular, surfaces which are not close to perpendicular to a beam will not be recorded. Also, ultrasound does not pass through bone or air, obscuring portions of the image. Similar limitations exist for the imaging of other organs.

For analysis of medical images this model-based approach is very appealing because it is already known what to expect in the images, and there are only a finite (albeit large) number of useful abnormal shape categories representing known pathologies. The major problem is how to capture the essential shape and range of variation in the model. Relatively simple geometric models such as generalized cylinders are unlikely to accomplish this.

Satellite Images The aurora borealis is one of the most spectacular phenomena of nature. It occurs in a roughly circular or oval region around the magnetic poles of the earth. Auroral displays are the result of processes which transfer energy from the solar wind to the Earth's magnetosphere and ionosphere. Many of these processes cause the energization and scattering of charged particles which precipitate along magnetic field lines into the high latitude upper atmosphere. There they ionize atmospheric constituents to produce visible and UV emissions. Auroral phenomena are present most of the time, though they vary greatly in dynamic behavior and intensity.

In earth surface scanning by satellites in visible wave lengths cloud cover can obscure views of the earth. This is not a problem for our auroral images, but here solar illumination obscures the dayside portion of the auroral oval. The auroral images gathered by the DE-1 satellite are simpler to model than the medical images presented above. To a first-order approximation, the auroral emissions form an oval or an elliptical region in the images. Even for this simple representation, there are practical problems to be solved in finding the oval. One of the difficulties is that the auroral region does not always appear continuous. The images sometimes show gaps in the auroral oval. The physical processes associated with the formation of the auroral oval, however, suggest that one should interpolate across gaps to define a continuous model for the oval.

1.4 Object Parameterization

Humans seem to utilize a mental model to compensate for incomplete or obscured data in the image. The parameterized geometric shape models presented earlier seem too rigid for our task. Biologic and other natural objects are not easily captured by this representation. Objects such as the kidney or heart are not describable by analytic formulas that maintain accuracy while capturing essential shape parameters and range of variation. Still, for breadth of application it is wise to use general models; for the experiments directed towards fetal imaging we used balloons of various shapes in a water bath.

2. Model-based Medical Image Analysis

In recent years there has been a proliferation of medical imaging techniques. There is also a trend towards generating digital (as opposed to analog) images, both as two-dimensional slices and as complete three-dimensional reconstructions. With the proliferation of digital images has come the development of picture archiving and communications systems (PACS) in radiology departments, which are being built to allow efficient storage, retrieval, and display of these images. With the development of PACS and the ever-increasing number of digital images there is a growing need for efficient and useful indexing schemes.

2.1 Indexing of Images

As in other imaging fields, the prime methods for indexing are manual input of primary parameters such as patient name, type of exam, etc. Diagnosis codes are also often included, following interpretation by the radiologist. Even with these relatively simple indexes, use of a PACS greatly increases the accessibility and usefulness of medical images. Without PACS the filmed images are placed into manila-paper jackets, identified with a sequence number, and shelved. A PACS can retrieve in one request all images of a specific patient and can also cross-reference images by patient diagnosis. Rapid availability of multiple images via a computer network increases the physicians' expectations for new retrieval options.

This setting motivates use of the computer for automatic or semi-automatic interpretation and classification of images. Although a completely automated analysis is a very distant and questionable goal, semi-automatic analysis and classification will be useful to reduce the tedium of scanning the ever-increasing number of images that are being produced. Such systems will allow improved quantification of organ parameters such as shape and volume, which are useful as subtle indicators of disease and response to therapy, and will be useful secondary parameters. Quantification in particular is very tedious and error-prone when performed manually, since it usually requires many detailed measurements on the images.

2.2 Current Medical Image Analysis Techniques

Image enhancement techniques are based on work done in other areas, such as analysis of satellite images. The next phase is the application of image analysis techniques for classifying images. For medical imaging it is not sufficient to simply match a complete image against a set of stored templates because the number of possible variations is too great. First the image must be segmented into meaningful regions corresponding to anatomic structures, then parameters of these regions (such as area, average grey level, etc.), can be measured. Measurements are used to establish features, and these features are used to classify the individual biologic structures.

Most methods utilize low-level image processing techniques such as region-growing, and edge-detection [Yachida 1980] to separate the regions of interest. There is continuing development of techniques for image understanding using such concepts as mathematical morphology [Brady 1982].

In several cases these simple low-level techniques perform well, as long as there is an unambiguous means for separating regions or edges based on a statistical classifier such as a discriminant function. A particularly simple and useful example is the segmentation of bone from the other parts of the images in 3D reconstructions from computed tomography (CT) scans. In this case the CT number stored in each 3D pixel or *voxel* of the 3D image representing bone differs from the CT numbers representing soft tissues so that a simple thresholding technique may be used, and this is the basis for several commercial systems. However, for separating soft tissues in CT or magnetic resonance imaging (MRI) scans no simple thresholding technique is adequate.

2.3 Model-based Image Segmentation of Ultrasound Images

An example of a medical imaging modality where simple thresholding of regions is grossly inadequate is ultrasound. Ultrasound instruments produce two-dimensional real-time slices by transmitting pulses of high-frequency sound waves into the tissues, and then creating an image from the returned echoes. Brightness discontinuities occur where impedance mismatches between adjacent tissues produce returned echoes. However, ultrasound does not pass through bone or air, and the ultrasound beam must be more or less at right angles to tissue interfaces, so edges at other angles representing organs are often not continuous.

Although the images from ultrasound do not have the resolution of X-ray images, ultrasound has the advantage that it uses no ionizing radiation, is painless, low-cost, and produces images in real-time. Thus, it is useful for fetal imaging, where it is not appropriate to use ionizing radiation, and for heart imaging, where motion is a problem for slower imaging methods, such as X-ray tomography.

Because of the difficulty of using thresholding techniques in processing ultrasound images, segmentation is done by hand tracing of the borders. However, since organs are three-dimensional, three-dimensional reconstructions are needed to give accurate volume and shape estimates. In these cases hand tracing of all the borders would be hopelessly inefficient and error-prone.

A radiologist is able to segment ultrasound images because knowledge of anatomy compensates for missing or ambiguous edge data. This is evidenced by noting that a novice looking at an ultrasound image can usually make nothing out of it, while after training in anatomy and experience with ultrasound images, the novice is able to see things not seen before. Thus, one of the main goals of this work is to give the computer knowledge of anatomy that mimics the radiologist's knowledge.

This knowledge is in the form of a general model of both normal and abnormal anatomy, and forms the basis both for aiding the analysis and segmentation of images. Parameters describing the specific model adaptation are then used for indexing these images in a visual database. The model should eventually include many different features and would form the basis for a knowledge base of human anatomy, but current work concentrates on the representation of shape.

2.4 Representing Anatomic Shape

The difficulty with developing a shape model of anatomy is that it is very difficult to represent the essential shape of an organ as well as the range of variation. This problem is not as acute in machine vision for man-made objects, since these objects are often describable by geometric shapes used during the design of these objects. For a biologic object such as a kidney, no combination of simple shapes is able to capture the essential shape as well as the range of variation.

In our experiments with ultrasound we attempted to develop a generic model for anatomical objects. This model is only applicable to objects which could be described as distortions of a sphere, which includes, for example, the left ventricle of the heart or the kidney. The model was tested on three-dimensional ultrasound reconstructions of two balloon shape classes (round and long-thin) ultrasonically imaged in a water bath. Several reconstructions from each shape class were used as a training set to establish the generic models for round and long-thin balloons. These models were then used to aid the analysis of images from new round or long-thin balloons that the computer had not seen before. The use of the model allowed fewer images of one object to be examined while still maintaining accuracy of volume and shape. It also showed the feasibility of model guided search for edges in the images. The following summarizes our issues based on results reported in [Brinkley 1985].

2.5 Model Representation

Figure 1 shows the basic model structure, which consisted of a set of fixed radials emanating from the origin of an object-based coordinate system. Three variable vertices along each radial, when connected via triangular surface patches to nearest neighbors, defined an inner and outer uncertainty surface, and a middle *bestguess* surface which was halfway between the inner and outer uncertainty surfaces. The inner and outer uncertainty surfaces defined the uncertainty volume within which the organ was expected to lie, while the bestguess surface was the computer's best guess, at any particular time in the processing, as to where the actual surface was, consistent with the data it had seen so far. Initially, the inner and outer vertices were set to small and large values respectively, thus ensuring that the initial model enclosed any possible organ instance.

2.6 Interaction of the Shape Knowledge with the Data

This initial model was superimposed on actual 3D image data by manual input of the 3D coordinates of the two balloon endpoints, thus bypassing the very difficult problem of aligning the model with the data. This approach was considered reasonable in an interactive system because it is relatively simple for a human operator to indicate a few landmarks in the data. For large volume processing it will be desirable to automate this step as well.

Given the two endpoints for a new balloon the program superimposed the initial model on the data, then generated a smaller uncertainty volume by moving the inner and outer vertices closer together. The smaller uncertainty volume was generated via interaction of learned shape knowledge contained in the generic model with the two endpoints representing the initial data from the image. The shape knowledge was represented as constraints on the range of possible slopes of each of the edges connecting adjacent radii. The ranges were obtained from the training set of similarly-shaped balloons. The constraints on edges, together with the model vertices, formed a constraint network [Brinkley 1987]. The hypothesis was that the collection of local constraints when interacting together in a relaxation process, were enough to generate the global shape and range of variation for the object.

The two endpoints of the balloon each defined the best guess and uncertainty vertices for the two radii at the poles of the model. This information was then propagated throughout the network of radial vertices by utilizing the learned slope constraints. Thus, for the radials adjacent to the pole vertices the slope constraints allowed the uncertainty to be reduced from the initial large value, even though no direct data was obtained at these radials. The reduced uncertainty at these radials in turn reduced the uncertainty at their neighbors, and so on throughout the network. The process was similar to the progress of a wave traveling over a globe covered with water. Waves of information began at the poles where data was obtained, and traveled outward towards the equator. The further the wave traveled away from the data, the more it became attenuated, so less information was known at the equator than at the poles.

Figure 2 shows the uncertainty volume for the generic model of long-thin balloons, after it had been superimposed on the 3D ultrasound data by indicating the two balloon endpoints manually. The uncertainty surfaces are those generated by interaction of the generic slope constraints with the two endpoints. The fact that the data was acquired at the poles is indicated by a narrow uncertainty volume in these regions. Figure 3 shows the initial bestguess model superimposed on the actual data, which in this case consisted of 3D coordinates of balloon edges manually input with a light pen. Note that the generated model is a reasonable depiction of what might be considered the *essential* shape of a long-thin balloon, even though no prototype was ever created.

Figure 2 also shows the intersection between an ultrasound slice and the uncertainty model. The ultrasound data for a balloon was obtained from a series of 2D slices related to each other by a position locating device. Given the position information and the superposition of the model on the data it was not difficult to determine the relationship of each slice to the model. The intersection of the model with each slice then produced a 2D tolerance region on each slice which provided a region within which to search for edges in the ultrasound image. These edges could then be used to update the model in the same manner as the initial balloon endpoints updated the model.

Thus, Figure 4 shows the result of updating the uncertainty volume after taking into account edge information from the first selected ultrasound slice, and Figure 5 shows the resulting bestguess surface. Note that now the uncertainty in the model is much narrower, and the bestguess surface is a much better fit to the data.

The model-fitting process was allowed to continue in this fashion. A particular ultrasound slice was selected for examination, the edges on that slice were used to update the model, and the process was repeated. Since numeric uncertainty and bestguess volumes were always available the procedure could terminate when uncertainty volume did not change much between successive scans. Thus, the computer *knew* when it had seen enough data. For the balloons this usually occurred after about one-third of the scans had been examined. We demonstrated that volume accuracy was the same at early termination as if all the scans were examined.

2.7 Feature Extraction

The generality of the model causes that many parameters, among others measures along all the radials, are generated. These are not useful for indexing. although they do define the actual shape much more precisely than would be possible by relying only on measurements from selected two-dimensional observations. However, now the volumes of the shapes and subshapes can be computed. Volume is a useful secondary parameter for indexing the database. Since volume is a direct corollary of fetal weight an eventual operational system will be able to collect in the database the most critical parameters of fetal growth.

3. Modeling Satellite Images of the Aurora

In this section, we present a model used in ongoing scientific investigations of the aurora. The model is based on assumptions regarding the representation of the aurora in images obtained by high altitude polar satellites. One current assumption, still to be validated, is that the aurora is at a constant altitude above the poles. This permits us to work with a 2-D model. Assumptions about the physics of the aurora provide techniques for finding features in the images. We show the results of preliminary demonstrations of the techniques using auroral images obtained from the Dynamics Explorer 1 (DE-1) satellite.

Investigations of auroral phenomena during the 1960's and 1970's utilized arrays of all-sky cameras to study the development of auroras all around the polar cap throughout the night- and day-side portions of the earth. Polar-based global imaging of the aurora utilizing satellites was initiated over 7 years ago with the Dynamics Explorer-1 (DE-1) auroral imaging photometers at visible and UV wavelengths. More than 500,000 images have been acquired to date by DE-1 and upcoming satellites having higher temporal and spatial resolutions will produce over 600 images per day.

3.1 Information Content of Auroral Images

Images of the aurora contain significant information pertinent to many scientific investigations of magnetospheric and auroral physics. This information ranges from the location of the aurora (latitudinal boundaries as a function of position around the earth) to the intensity of the aurora at different wavelength emissions, to structural features within the aurora and the time evolution of auroral disturbances. A number of important classification parameters for auroral images can be computed once the auroral boundaries are determined. Such secondary parameters include the total integrated intensity of the auroral oval, width of the auroral oval as a function of position around the earth, intensity of the aurora as a function of position around the earth, and the area of the dark region inside the auroral oval.

3.2 Modeling Auroral Images

As indicated in the introduction, auroral images are incomplete. There is a need, therefore, to model the oval region using closed, continuous curves which fill the gaps in the visible portions of the aurora. The set of techniques which we are using for identifying the auroral oval are based on recent computer vision work by Kass [Kass 1988]. Their method is related to fitting spline curves under tension to the image data.

Before trying to identify the auroral oval, one must eliminate the effect of dayside illumination which causes problems for the intensity-based modeling algorithms. The technique used for removing the dayside illumination involves a *shaded sphere* calculation, a technique used in the computer graphics community to render realistic images. The sun is assumed to be an infinitely distant point source of illumination. The reflectance from the atmospheric surface at ultraviolet wavelengths is approximated to follow a Lambertian law. These two approximations, together with information about location and position of the earth and sun provided by the DE-1 satellite ephemeris data, allows the calculation of the vector dot product between the sun's position vector and the local surface normal vector. With the additional assumption that the satellite is far from the earth, the illumination under the Lambertian law is proportional to the cosine of the incident angle between the sun vector and the surface normal. The dayside illumination is removed by subtracting from each pixel in the original image a value proportional to the cosine of the incident angle. Once the dayside illumination is removed, the curves describing the auroral oval may be found. The results shown below for finding the oval ignore the illumination compensation step since the auroral oval for the images shown are in the winter hemisphere and therefore appear in the dark portion of the earth.

The technique for finding the auroral oval is based on curve fitting with splines. For standard, one dimensional splines, there is a relationship between the fitting of the splines to the data and the bending energy of small deflections of a rod. The curves used here, developed first by Kass, are analogous to splines with an additional tension term used to fit two-dimensional data. That is, the curves are analogous to elastic materials which resist both bending and stretching.

3.3 The Basis for the Model

In this section, we discuss the model used for finding various features in the auroral images. Formally, the model simulates the statics of an elastic curve which is constrained to lie in the plane of the image [Kass 1988, Samadani 1989]. The image exerts forces on the elastic curve. Strings, thin lengths of rubber or thin metallic rod are examples of the materials which are simulated. The feature detection process starts by positioning the elastic curve on the image in some initial configuration. Starting from the initial configuration, a local minimum of the energy of the elastic curve is found. The configuration of the elastic curve at equilibrium provides the geometric description of the feature.

The total energy, E, of the elastic curve is given by [Terzopoulos 1988] as

$$E = \int_{\Omega} \alpha(s - s^0)^2 + \beta(k - k^0)^2 + P(x, y)da$$

where the integral is over the curve coordinate a. The natural arc length and the natural curvature of the elastic curve, given by s^0 and k^0, respectively, determine the preferred shape for the material. The actual arc length and the actual curvature of the elastic curve are given by s and k. The first term of the integral, therefore, describes the energy added to the system by perturbing the length of the material and the second term of the integral describes the energy added to the system by perturbing the curvature of the material. The constants α and β determine the characteristics of the material. For example, materials with large α will not stretch or shrink easily and materials with large β will not bend easily.

A potential, $P(x,y)$, derived from the image, applies external forces to the elastic curve. This potential can be modified to allow the tracking of various features. Feature extraction is defined as finding an extremum of the integral equation for the energy. An iterative method is used to step from an initial configuration of the elastic curve to an equilibrium configuration by choosing new configurations of the curve which lowers its energy. At the equilibrium, the shape and location of the elastic curve describes the feature. This equilibrium results in a smooth fit to the image data since the elastic components of the energy are lower for smooth functions.

3.4 Implementation

We have developed an interactive software system to allow experimentation with the use of elastic curves to find auroral oval features. The software was developed using the Sun 3/260 computer and the SunView* windowing system. The software system was developed to allow users to gain experience with the elastic materials. Preliminary results show that the techniques used in the system are successful in finding the curves passing through the most intense locations of the oval and in outlining the inner and outer boundaries of the oval. The system may also be used to track the changes of the auroral oval through time.

The system assumes default values for various of its parameters. While the system is running, it allows the user to interactively modify the default parameters by using a graphical interface of software-simulated sliders and buttons. The parameters under user control include the elastic constants which control the material's resistance to bending, resistance to stretching and the preferred rest length of the material. The user may also modify the step size of the gradient descent iteration and the boundary conditions which determine whether the elastic material forms a closed loop or is open at the ends. The user can select new images and also control the elastic material's positions. The elastic material may either be input from a file or the user may use a mouse to draw the initial configuration of the elastic material.

3.5 Interactive Image Analysis

By choosing different functions of image intensities, one may find various features. The first, simplest application is finding a curve which goes through the most intense parts of the oval. To do this, one chooses a potential field directly proportional to the image intensities. Then, one first draws the elastic curve outside of the actual oval.

The initial configuration of the elastic curve is currently input manually, but this step may be automated by using latitude and longitude information from the satellite to define an ellipse which is certain to enclose the actual auroral oval. The length at rest of the elastic material is chosen to correspond to a natural perimeter which is smaller than the perimeter of the actual oval. After setting of the initial parameters, the iteration process is started.

The technique for finding the curve through the maximum intensity areas of the oval was applied to a section of a DE-1 image containing the auroral oval. This image is shown in grayscale in Figure 6. The initial and final configurations of the curve are shown in Figure 7, superimposed on the auroral intensities. In this figure, the initial, irregular hand-drawn elastic curve is found outside of the oval. Also shown is the final location of the elastic curve at equilibrium. The auroral intensities shown in Figure 7 have been decreased to make the curves more visible. The curve at equilibrium fits smoothly the

* SunView is a trademark of Sun Microsystems, Inc., Palo Alto CA

maximum intensity points of the oval. Figure 8 shows the intermediate results of the iterative algorithm. Appearing in the figure from left to right, and top to bottom are the various images which depict snapshots of the configuration of the elastic curve as it proceeds from its initial configuration to its equilibrium configuration.

The equilibrium configuration of the elastic curve through the most intense points could be used to study bright spots in the aurora. An example of the secondary parameters that could be obtained here are the image intensities along the elastic curve. As an alternative, one could store the average intensity in small neighborhoods of pixels along the curve. This information could be used to search for bright spots above some threshold.

Another application of the system is finding the inner and outer boundaries of the oval. This is done by using two independent elastic curves, one for each boundary. The potential field used in this case is directly proportional to the negative values of the image intensities. To find the outer boundary, one draws an elastic curve outside of the actual oval and sets its rest length to correspond to a natural perimeter for the curve which is smaller than the perimeter of the oval. During the iterative process, this elastic curve shrinks to fit the outer boundary of the oval. It stops at the outer boundary since further shrinking would result in a higher energy due to the potential field term derived from the negative intensities of the image. To find the inner boundary, one draws an elastic curve inside the oval and sets its rest length to correspond to a natural perimeter which is larger than the perimeter for the actual oval. During the iteration process, this elastic curve will grow to fit the inner boundary of the oval.

The technique for finding the inner and outer boundaries of the oval was applied to the image in Figure 6. Figure 9a shows the initial configurations of the two elastic curves which are used to find the boundaries, superimposed on the auroral intensities. The auroral intensities have been decreased in the image to make the curves more visible. Figure 9b shows the equilibrium configurations of the two elastic curves, which are found to fit the inner and outer boundary of the superimposed aurora. Figure 10a shows the intensity values of the image before processing. Figure 10b shows the intensity values of the image, but only at those locations that correspond to auroral features between the inner and outer elastic curves. Thus, a comparison of Figures 10a and 10b demonstrates that the preceding steps are successful in extracting meaningful information from the original image data.

The configurations of the inner and outer curves at their equilibrium could be used as the basis for secondary parameters. One can compute from these inner and outer boundaries the area enclosed between the inner and outer boundaries, which reflects the extent of the aurora. The area inside the inner boundary is also a useful parameter. Other secondary parameters, such as the average integrated auroral intensity between the boundaries could be obtained. The equilibrium configurations of the inner and outer curves allows the extraction of each of the parameters mentioned.

3.6 Tracking of Changes through Time

In addition to finding features, one may use the elastic curves for tracking features through time. This can be done by applying the elastic curves to a time sequence of images. For example, consider the tracking of the brighter areas of the oval. Start with the first image in the sequence and find the bright regions of the oval in this image by using the technique described above. Now, use the equilibrium configuration of the elastic curve in the first image as the initial configuration of a second curve applied to the second image in the sequence. Solving, as in the first image, the equilibrium configuration is found for this, second image. This procedure is repeated for each image in the sequence, resulting in a sequence of curves whose equilibrium configuration changes with time. The parameters obtained from one image can be used to initialize the model for the next image.

The motion tracking technique was applied to the image sequence shown in the top row of Figure 11. Because of satellite scanning speed, there is roughly a 25 minute interval between consecutive images in the sequence. The bottom row of the figure shows the sequence of equilibrium configurations of the elastic curves corresponding to the images above them. The changes in the configurations of the elastic curves tracks the changes of the bright areas of the oval through time. The same technique could be used to track inner and outer boundaries or other features.

3.7 Feature Extraction

The techniques described above can be used to find various features in the images. The examples shown have used the configurations of elastic curves to find the positions of the highest intensity regions of the oval, the positions of the inner and outer boundaries of the oval, and to track the changes in the oval through time. Other applications may be developed by changing the definition of the potential field derived from the image intensities.

For all applications, the results of the feature extraction is a description of the two dimensional coordinates of discrete points along the curves. These parameters are suitable for indexing an image database. They provide a compact description of the interesting features in the images. Also, the data structure is one-dimensional, simplifying its storage. Even though the data structure is one dimensional, useful two-dimensional geometrical parameters can be derived from this information, with far less effort than analysis of the original images. For example, perimeter and centroid are readily calculated in a time linear with respect to the number of points on the curve. The changes in perimeter and centroid may also be easily followed through time. Other parameters, such as local motion of a segment of a curve, requires further processing, but methods for such calculations are known [Hildreth 84]. The curves can be used to derive the geometrical quantities to be stored in the database, or the curve parameters themselves could be stored and queries made directly on them. The choice between the two will be easier to make once we have more experience with the processing speed requirements of the application.

4. Summary

We have shown that similar principles can be applied to widely differing applications of image analysis. By using a high-level adaptive model and establishing parameters which describe the adaptation of the model we obtain information of a type which low-level image analyses, characterized by image-enhancement techniques, cannot provide. Adaptive-model-based image analysis extracts image features to augment the database and index the images for query and further selection.

The features extracted require much less storage than the original images. The database system may be designed to use different media and access capabilities for the image storage and the searchable portion of the database.

Medical Image Information The representation used here to describe general non-structured biologic objects is a set of geometric distortions of a sphere. The system uses this model to extract three-dimensional organ reconstructions from a series of arbitrarily-oriented ultrasound slices. A training set of ultrasonic reconstructions of similarly-shaped objects is used to give the computer generic knowledge of a given shape class.

An hypothesize-verify paradigm is employed to alternately request new data and to update the tolerance region and bestguess surface. This enabled the analysis program to build an accurate three-dimensional representation from a series of arbitrarily-oriented ultrasound image slices of the object.

Auroral Image Information The model used to describe the auroral shape is based on an analogy with elastic material. The properties of the material are parameterized. The system allows a user to interactively choose parameters which vary the material's rest length and its resistance to either stretching or bending. For example, to find the outer boundary of the aurora oval, elastic material is placed in a circle of arbitrary size which encloses the aurora. The initial rest length of the elastic material is chosen to guarantee that it is less than the perimeter of the outer boundary. The equilibrium configuration of the elastic material then describes the outer boundary. The inner boundary is found symmetrically.

4.1 Indexing Parameters

Parameters for database obtained for the medical shapes are volume measures. Such a secondary parameter is useful to indicate growth, where it desirable, as in a fetus, or not desirable, as in cancers and some cardiac conditions. For the aurora borealis the model provides the perimeter, centroid, and area enclosed within the boundary. These parameters are made available for querying by the scientist, who can now conveniently track series of images. The environment planned for by [Crehange et al 1984] has that flavor.

Once an object has been recognized and measured, then further parameters, indicating normality or abnormality, might be useful. Abnormalities might be unexpected length-to-width ratios, or large local variations in a surface. Analysis of multiple objects in an image can provide measures of relative size, placement, and orientation.

4.2 Limitations and Extensions

The models we used will also have to be extended to allow more complex structured objects. A larger class of objects is possible if the vertices are allowed to move in arbitrary directions rather than only along their radials. A development of this approach showed partial success in the representation of proteins by constraint networks [Brinkley 1987].

Due to the large number of images which are becoming available, it is best to develop feature extraction methods which will involve little or no human interaction. The techniques we show are currently used, at least partially, interactively, but full automation of the processing is the eventual goal. There are two aspects where interaction is currently important:

1 Defining some initial fiducial points so that the program does not have to search for an anchorpoint for the model. This is difficult if the image can contain multiple candidate objects.

2 Guiding the program if an incorrect object boundary is accepted early on. Such an error will unfairly bias the resulting model.

The problem of guidance occurs in most areas of science, where partial data may suggest an incorrect hypothesis, which then must be corrected to account for new data.

Approaches to be considered for dealing with this problem are to use more knowledge, such as probabilistic descriptions of the constraints, allowing backtracking, and to use least-commitment concepts by allowing a set of possible edges in defining each boundary, rather than just one.

4.3 Parallel Image Analysis

The heavy computing demands engendered by image processing drive us to consider parallel processing. Most considerations of parallelism to date have focused on pixel-oriented, synchronous, analysis of images, and that will remain an important preliminary phase for many images [Yamaguchi 1982].

Model-based interpretation of images presents an opportunity for extending asynchronous parallel processing. Currently some projects exploit parallelism to deal with multiple objects in a scene. We think that multiple candidate objects and models can be investigated in parallel, adding one more dimension of parallelism. An independent demon may abort some of the model processes which have located the same object. A measure of matching success attained by these independently operating processes can be the criterion for letting them continue. Even left alone, successful processes would tend to terminate earlier than problematical ones. The final collection of objects and models matched may be pruned, or we might populate the database with multiple *maybe* candidates.

4.4 Incomplete Processing

In any case, if we avoid manual interaction and refinement, we may not be able to resolve the image contents determination to a single, highly probable set of objects. But we can permit having multiple indices to one image. Some false images will be retrieved in response to queries on secondary parameters. We then must leave the final selection of relevant images to the eventual user of the image information. In this approach the secondary parameters provide a relatively coarser filter than we now seek to obtain with our interactive analyses. However, the real objective of having secondary parameters would *not* be violated by this approach, since the concept of using secondary parameters to aid the user by providing a manageable set of images is still being achieved.

5. Finale

We have used adaptive model-based techniques to obtain secondary parameters for indexing images stored in databases. While these applications differed greatly they also show important commonalities. The top-down approach employed permits the use of generic shape knowledge to compensate for limitations of the imaging modality, and can reduce the analysis time by requiring less data to be acquired.

This techniques are not adequate for full automation of real-world image processing. It appears, however, that they are sufficiently effective to be used as an augmentation for manual indexing. If employed in that manner, the experience gained will help us in developing the technology for more complete automation.

If model-based image processing develops as we would hope then we can envisage specialized image knowledge-bases. For instance, a knowledge-base of human anatomy would represent both normal and pathologic shape variation in different shape classes. When classifying a new set of three-dimensional image data defining a likely object one or a set of ranked hypotheses could be suggested by the matching parameters vis-a-vis the database.

Not being able to reduce, during analysis, the findings to a single object hypothesis need not discourage work in automated image analysis. Any significant reduction in search space will be helpful to the users of large image databases.

Acknowledgements

Knowledge-acquisition research was partially supported through the RX project, funded by the National Center for Health Services Research (HS-3650 and HS-4389) and the National Library of Medicine (LM-4334). Research towards fetal imaging was supported by NIH under grants HD-12327, RR-0785, and by a fellowship (F32 GM08092) to J. Brinkley. The scans were made in the department of Gynecology and Obstetrics of Stanford University, and we thank the chairman, faculty — especially Desmond McCallum, and staff for their assistance. Figures 1–5 came from [Brinkley 1984].

The work on auroral images has been supported by the Center for Aeronautics and Space Information Systems (CASIS) under NASA grant NAGW419. We thank Domingo Mihovilovic for providing most of the code for the interactive elastic curve system on the Sun. We thank Drs. L.A. Frank and J.D. Craven at the University of Iowa for use of the DE-1 auroral images examined in this research. Michael Walker participated in the conceptualization of the work on auroral feature extraction and measurement; this research is continuing through a grant from the NASA Center for Excellence in Space Data and Information Systems (CESDIS), further focusing on interaction with large scale data management.

Basic research on the interaction of knowledge and data is supported by the Defense Advanced Research Projects Agency under DARPA N39-84-C-211.

References

Arvidson, Raymond, Gio Wiederhold, et al: *Issues and Recommendations Associated with Distributed Computation and Data Management Systems for the Space Sciences*, volume 2; Committee on Data Management and Computation, Space Sciences Board, National Academy of Sciences, National Academy Press, 1986.

Binford, T.O.: "Visual Perception by Computer"; Invited talk at the IEEE Conference on Systems, Science and Cybernetics, Miami, 1971.

Brady, M.; "Computational Approaches to Image Understanding"; *ACM Computing Surveys*, Vol. 14 No. 1, pp. 3–71, 1982.

Brinkley, J.F., W.E. Moritz, and D.W. Baker: "Ultrasonic Three-dimensional Imaging and Volume from a Series of Arbitrary Sector Scans"; *Ultrasound Med. Biol.*, Vol. 4, pp. 317–327, 1978.

Brinkley, James F.: Ultrasonic Three-Dimensional Organ Modelling; PhD thesis, Committee on Graduate Studies, Stanford Unversity 1984.

Brinkley, J.F.: "Knowledge-driven Ultrasonic Three-Dimensional Organ Modelling"; *IEEE Transactions on Pattern Analysis and Machine Intelligence*; Vol. PAMI-7 No. 4, pp. 431–441, 1985.

Brinkley, J.F., B.G. Buchanan, R.B. Altman, B.S. Duncan, and C.W. Cornelius: "A Heuristic Refinement Method for Spatial Constraint Satisfaction Problems"; Stanford University, STAN-CS-87-1142, KSL-87-05, Jan. 1987.

Brooks, R.A. and T.O. Binford: "Geometric Modeling in Vision for Manufacturing"; *SPIE Proceedings* 281, 1981.

Chang, S.K., C.W. Yan, D.C. Dimitroff, and T. Arndt: "An Intelligent Image Database System"; *IEEE Transactions on Software Engineering*, Vol.14 No.5, May 1988, pp. 681-688.

Crehange, M., A. AitHaddou, M. Boukaiou, J.M. David, O. Foucaut, and J. Maroldt: "EXPRIM, an Expert System to Aid in Progressive Retrieval from a Pictorial and Descriptive Analysis"; in *New Applications of Databases*, Gardarin and Gelenbe(eds), Academic Press, 1984.

Hildreth, Ellen C.: *The Measurement of Visual Motion*; MIT Press, Cambridge MA, 1987.

Kass, M.A., A. Witkin and D. Terzopoulos: "Snakes: Active Contour Models"; *International Journal of Computer Vision*, Vol. 1, pp. 321–331, 1988.

Kasturi, R. and J. Alemany: "Information Extraction from Images of Paper-Based Maps"; *IEEE Transactions on Software Engineering*, Vol.14 No.5, May 1988, pp.671-675.

Roussopoulos, N., C. Faloutsos, and T. Sellis: "An Efficient Pictorial Database System for PSQL"; *IEEE Transactions on Software Engineering*, Vol.14 No.5, May 1988, pp. 639-650.

Samadani, Ramin: "Changes in Connectivity in Active Contour Models"; to appear in IEEE Workshop on Visual Motion, March 1989, Irvine, CA.

Soroka, B.I.: "Generalized Cones From Serial Sections"; *Computer Graphics and Image Processing*, Vol. 15, pp. 154–166, 1981.

Terzopoulos, Demetri, John Platt, Alan Barr and Kurt Fleischer: "Elastically Deformable Models"; *ACM Transactions on Computer Graphics*, Vol. 21, No. 4, July 1987, pp. 205–214.

Wiederhold, Gio: *File Organization for Database Design*; McGraw-Hill Book Company, New York, NY, March 1987.

Yachida, M., M. Ikeda, and S. Tsuji: "A Plan-guided Analysis of Cineangiograms for Measurement of Dynamic Behavior of Heart Wall"; *IEEE Transactions on Pattern Recognition and Machine Intelligence*; Vol. PAMI-2 No. 6, pp. 537, 1980.

Yamaguchi, K. and T.L. Kunii: "PICCOLO, Logic for a Picture Database Computer and Its Implementation"; *IEEE Trans. Computers*, Vol.C-31 No.10, Oct.1982, pp.983–996.

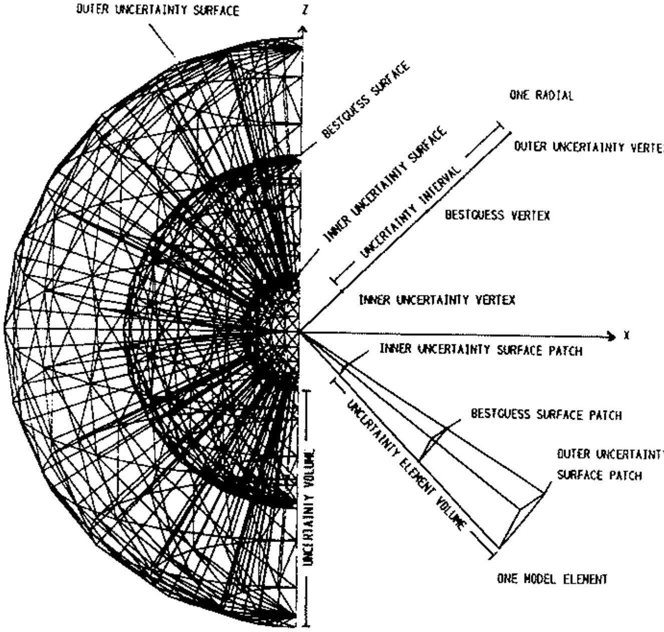

Figure 1: Model structure.

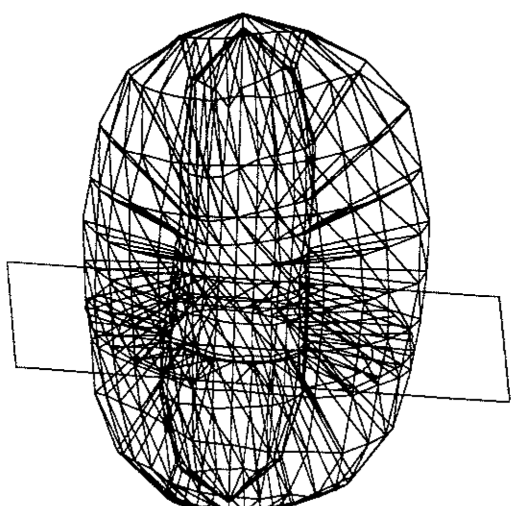

Figure 2: Initital uncertainty volume and first selected slice, long-thin ballon.

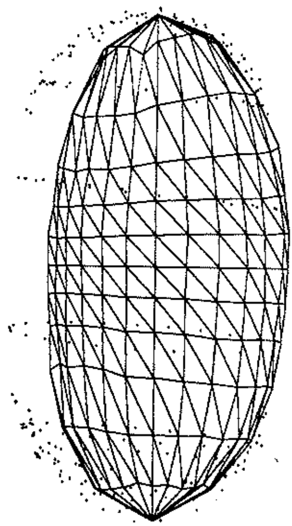

Figure 3: Initital bestguess surface superimposed on data, long-thin ballon.

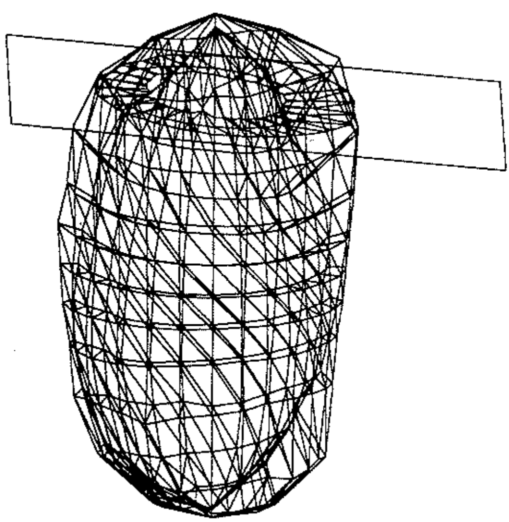

Figure 4: Uncertainty volume after global updating by first selected slice, second selected slice, long-thin ballon.

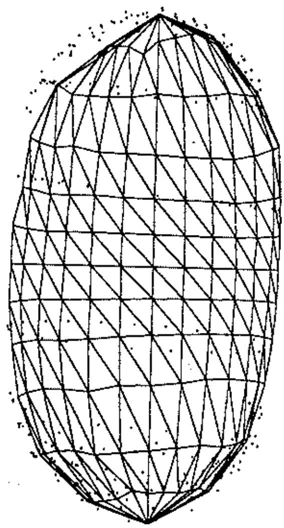

Figure 5. Bestguess surface after first selected slice, superimposed on data, long-thin ballon.

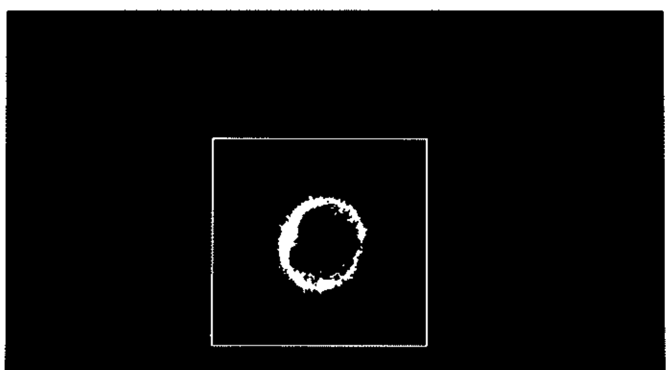

Figure 6. Graylevel coded intensity values for the portion of a DE-1 satellite image containing a winter hemisphere auroral oval.

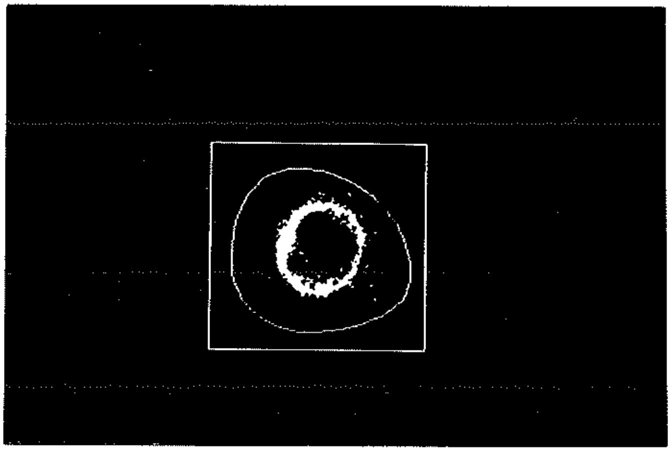

Figure 7. The initial and equilibrium configuration of an elastic curve superimposed on a darker grayscale representation of the auroral oval of Figure 6. The curve is used to find the most intense boundary of the oval.

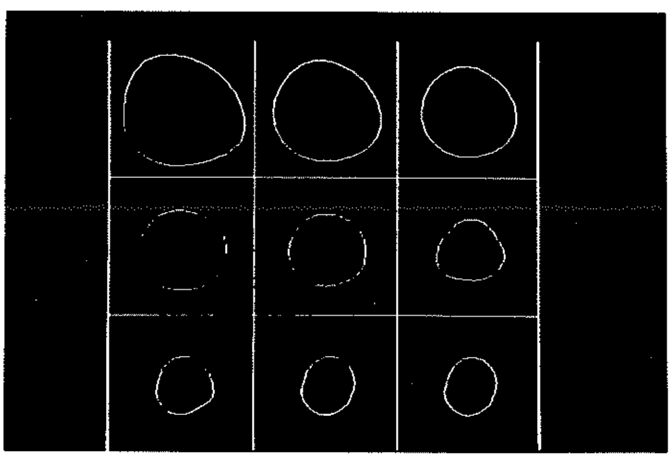

Figure 8. Intermediate steps of the iterative solution, beginning with the initial configuration and continuing to the equilibrium configuration of the elastic curve used for finding the brightest region of the oval.

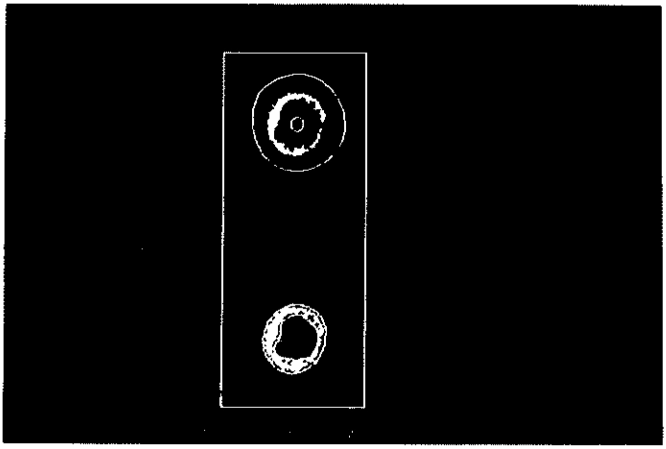

Figure 10. Graylevel coded intensity values of the DE-1 image shown in Figure 7, but only for the region between the inner and outer boundaries found by the elastic curves in Figure 9b.

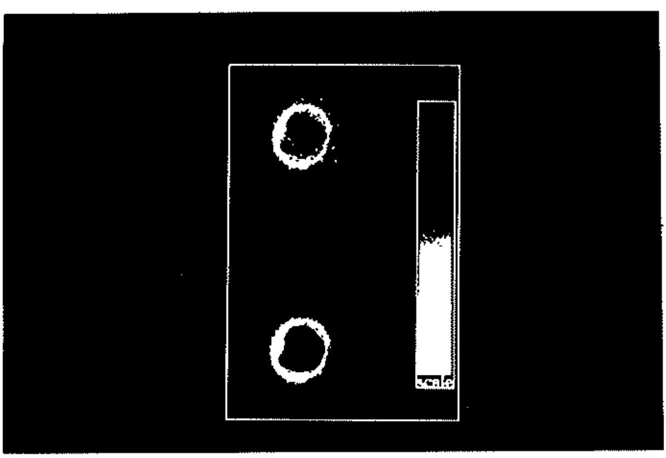

Figure 9. a) Initial configurations of the elastic curves used to find the inner and outer boundaries of the auroral oval. The elastic curve inside tends to grows and the elastic curve outside tends to shrink. b) The equilibrium configurations of the elastic curves which fit the boundaries of the aurora.

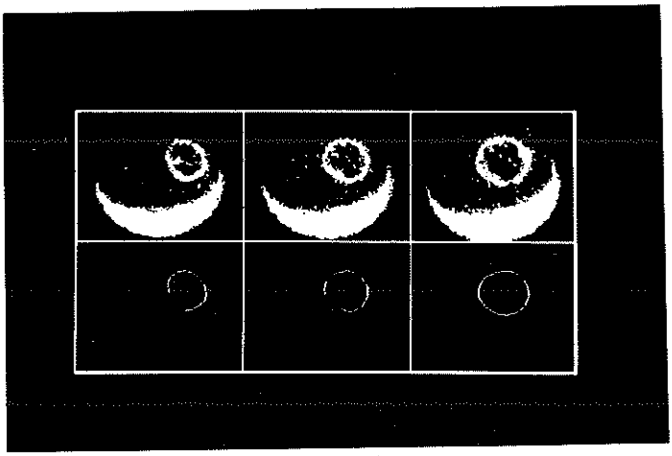

Figure 11. a) The three images in the top row show a sequence of DE-1 satellite images taken within one hour. b) The three images in the bottom row show the configurations of the elastic curves as they follow the motion of the maximum intensity areas of the auroral oval.