Proceedings of the Fourth IASTED International Conference
COMPUTATIONAL INTELLIGENCE
July 4-6, 2005, Calgary, Alberta, Canada

USING RELEVANCE FEEDBACK IN RETRIEVING AURORAL IMAGES

Mikko T. Syrjäsuo and Eric F. Donovan
Institute for Space Research
Department of Physics and Astronomy
University of Calgary
Calgary, Alberta, T2N 1N4
email: mikko,eric@phys.ucalgary.ca

ABSTRACT
In modern space physics research, digital imagers are widely utilised in studies of the near-Earth space environment. The physical process being directly observed is the aurora, and millions of auroral images are acquired annually. These data sets provide a wealth of opportunities for developing and testing content-based image retrieval (CBIR) techniques with the irregular natural shapes occurring in auroral displays. Our CBIR implementation with relevance feedback was used in searching for one rare auroral form ("North-South structure") that is a manifestation of an important physical process of general interest to space physics researchers today. We finish with a brief discussion of important benefits of anticipated application of this technique to multi-terabyte multi-million auroral image data sets.

KEY WORDS
Geosciences, pattern recognition, machine vision, content-based image retrieval

1 Introduction

Space science research focuses on understanding physical processes in the near-Earth environment. These processes arise from the complex interaction between the solar wind plasma, the terrestrial magnetic field and the Earth's atmosphere. Energy extracted from the solar wind drives numerous processes, including the aurora which is commonly referred to as the northern (and/or southern) lights. The mechanisms which give rise to temporal and spatial structure in the aurora are the subject of intense study, and so optical observations of the aurora are of fundamental importance in this field.

The aurora is a global phenomenon with important structure at scale sizes down to well below one kilometre. This fact has led to the development of optical observations capable of high spatial resolution over large geographical areas. The auroral distribution changes rapidly as well, so the rate at which these observations are collected is correspondingly high. Furthermore, these observations complement long-term satellite programs, so they are carried out for years. The end result is that auroral optical observations yield an enormous quantity of images.

The large auroral image data sets offer an excellent opportunity for development and testing of computer vision techniques. While the data sets are large, there are only a limited number of "types" of aurora, and so image classification, while difficult, is almost certainly a tractable problem.

Despite the fact that the research of computer vision in general and content-based image retrieval (CBIR) in particular has been very active in the recent years, there are few reported applications for analysis of auroral images. Before one can proceed to use any of the "standard tools" in literature [1, 2], there is an unsolved problem concerning auroral shape extraction: (1) the accurate shape of aurora is difficult to model, (2) finding edges in slightly transparent objects is non-trivial, (3) the images are typically photon-limited and hence noisy; and 4) one image can contain zero or more auroral objects as well as other objects such as lights from nearby settlements.

Our immediate objective is to create tools for auroral scientists that make it practical to browse and search image datasets of millions of images. We have initiated our CBIR experiments by using a relatively small data set consisting of 340,000 raw images from a Canadian auroral imager. The number of images is small enough that manual browsing of summary plots and even individual images in searching for certain auroral shapes is accomplishable though cumbersome. On the other hand, the number of images is large enough for testing the performance of automated image analysis in a challenging real environment: robust detection and extraction of irregular auroral shapes captured in varying imaging conditions is crucial.

2 Auroral imaging

2.1 The all-sky imager

We have used image data from a CANOPUS (Canadian Auroral Network for the Open Program Unified Study) all-sky imager (ASI), which was operated in Gillam (Manitoba, Canada) from 1986 to 2001 [3]. The imager consists of a fish-eye lens providing a field-of-view (FOV) of 160° degrees — almost the whole sky in one image — followed by a filter wheel, an image intensifier and a charge-coupled device (CCD) camera. We used all images

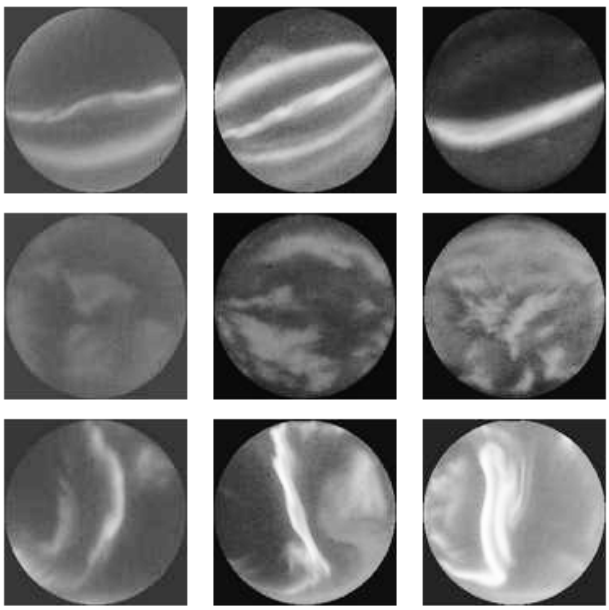

Figure 1. Top row to bottom row: auroral arcs, patchy auroras and north-south structures. These all-sky images show north at the top and east on the right; the circular field-of-view captures the whole sky.

acquired through a 557.7-nm narrow bandpass optical filter, which captures the most commonly observed yellowish green colour of auroras. Note that using an optical filter essentially removes all colour information which would enormously help in segmenting the image based on contrasting colours of the aurora and dark sky.

2.2 Typically observed auroral shapes

Fig. 1 shows a selection of auroral images demonstrating the variety of auroral forms. From a subjective point of view, *arcs* are the most unambiguous of auroral shapes. They are relatively bright elongated shapes that cross the FOV in approximately east-west direction. *Patchy auroras* are randomly distributed irregular auroral blobs. *North-south structures* are shapes that are aligned more or less along the meridian (ie., the north-south direction). Arcs and patches are relatively long-lived forms. A single imager will see arcs within its FOV for tens of minutes to several hours. The same is true for patches. North-south structures are more transient shapes that last no more than a few minutes. One consequence of their short lifetime is that they are exceedingly difficult to locate using summary plots. Both arcs and north-south structures are considered *discrete* shapes as opposed to patches which are referred to as *diffuse* auroras and do not have clearly defined shapes: this study concentrates on discrete auroras whose shape can be extracted much more easily.

3 Shape extraction

3.1 Image preprocessing

The raw data were stored as grayscale images at 200×200 pixel resolution (256 gray levels). We preprocessed the images by using a Gaussian low-pass filter ($\sigma = 4.5$, window size 25×25). We also determined the maximum $B_{\max}$ and mean B_{mean} brightnesses of image.

The number of images was 340,000. This includes all images of cloudy skies as well as those obtained during clear viewing conditions but when there was no visible aurora. Our first task is to remove the ones without visible auroras from further processing. Based on our earlier work [4, 5], we can use a simple linear classifier to separate images with salient auroras from those containing dim or no auroras. Let

$$y = 1.14B_{\mathrm{mean}} - B_{\max} + 15. \tag{1}$$

Now, based on the values of y we classify the contents of an image:

$$\text{If } y \begin{cases} < 0 & \text{the image contains salient aurora} \\ \geq 0 & \text{the image has no salient aurora} \end{cases}$$

This simple classifier vastly reduces the number of images requiring further analysis: the final shape extraction was performed to only 75,000 images.

3.2 Multilevel shape analysis

The ambiguous nature of auroral shapes forces us to use heuristics in extracting their outlines. We used a trimmed version of the isolabel-contour map algorithm from medical imaging [6]. There are four steps in this algorithm: (1) extracting contours, (2) detecting the strongest edges, (3) scoring the individual contours based on their overlap on edges, and (4) choosing non-overlapping contours with the highest scores to represent the most salient auroral objects in the image. An example of detected auroral shapes is given in Fig. 2. A more detailed description of the algorithm can be found in [5].

We used this algorithm to extract 93,000 objects from the 75,000 auroral images containing salient aurora. For each object, the maximum and mean brightnesses within the object were also determined.

4 Representation of an auroral object

4.1 Shape

We used a Fourier Descriptor (FD) based shape representation with the centroid distance as the shape signature. Let x_i and y_i be the pixel coordinates $i = 1, \ldots, N$ of the (closed) contour curve surrounding the detected auroral object. The shape signature is

$$r_i = \left((x_i - x_c)^2 + (y_i - y_c)^2\right)^{\frac{1}{2}}, \tag{2}$$

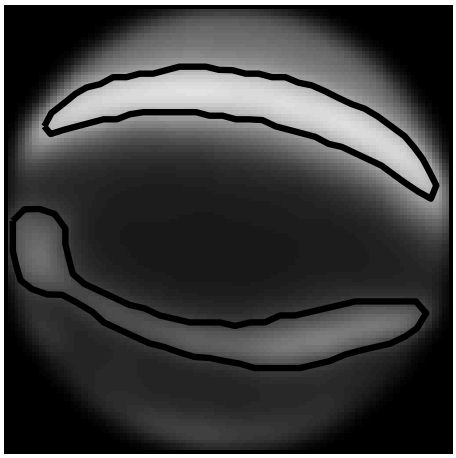

Figure 2. Results of object boundary extraction: there are two salient auroral objects at different brightness levels.

where (x_c,y_c) is the centroid of the shape. Now the Fourier coefficients of the signature are

$$a(k) = \sum_{i=0}^{N-1} r_i e^{-j2\pi(k-1)(i-1)/N}, \quad k = 0 \ldots N-1, \quad (3)$$

where j is the imaginary unit. Using the centroid distance provides a translation invariant representation. When comparing auroral shapes, the size and orientation can be relevant. We decided to separate the size and orientation comparison from the boundary shape comparison by using translation, rotation and scale invariant FDs:

$$\mathrm{FD}(k) = \left| \frac{a(k)}{a(0)} \right|, \quad k = 0 \ldots N-1.$$

A reasonable approximation of the shape can be formed by using M FDs, where $M < N$. However, as the normalisation forces $\mathrm{FD}(1) = 1$ for all shapes, that component was dropped from the descriptor

$$\mathbf{f} = [\mathrm{FD}(2)\ \mathrm{FD}(3)\ \ldots\ \mathrm{FD}(M+1)], \quad (4)$$

where $M = 16$ in the experiments of this paper.

The similarity is measured as the Euclidean distance:

$$d_{\mathrm{FD}}(\mathbf{f}_1, \mathbf{f}_2) = ||\mathbf{f}_1 - \mathbf{f}_2||_2, \quad (5)$$

where $\mathbf{f}_1$ and $\mathbf{f}_2$ correspond to the two different shapes being compared and $||\cdot||_2$ denotes an L_2-norm. Small distances are associated with more similar shapes.

4.2 Other object features

We used the dc-component of the Fourier transform coefficient (ie. $a(1)$ from (3)) as the measure of size. Similarly, the location of the object can be associated with the centroid of the contour. In order to compare the orientation, we used the contour points to determine the direction of the major axis. This results in five additional features:

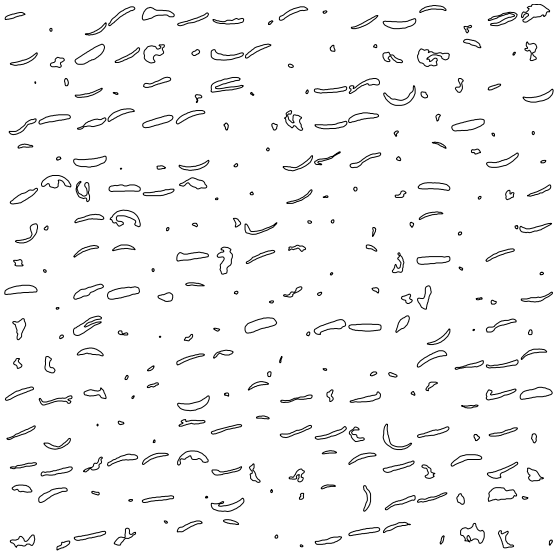

Figure 3. A selection of extracted auroral shapes: the most distinctive shape is the auroral arc (horizontally elongated shape). However, there are many other irregular shapes in different orientations and scales.

$$\begin{aligned} s_{\mathbf{f}} &= a(0) && (6)\\ \mathbf{c}_{\mathbf{f}} &= [x_c\ y_c] && (7)\\ B_{\mathbf{f}}^{\max} &= B_{max} && (8)\\ B_{\mathbf{f}}^{\mathrm{mean}} &= B_{mean} && (9)\\ \mathbf{v}_{\mathbf{f}} &= \mathbf{v}_{\mathrm{major}}, && (10) \end{aligned}$$

where $B_{\max}$ and B_{mean} are the maximum and mean brightnesses within the auroral object, and $\mathbf{v}_{\mathrm{major}}$ is the orientation of the major axis (with 180° ambiguity).

The final similarity between two objects $\mathbf{f}_1$ and $\mathbf{f}_2$ is calculated as a weighted sum:

$$\begin{aligned} d(\mathbf{f}_1, \mathbf{f}_2) = w_1 \cdot d_{\mathrm{FD}}(\mathbf{f}_1, \mathbf{f}_2) &+ \\ w_2 \cdot |s_{\mathbf{f}_1} - s_{\mathbf{f}_2}| &+ \\ w_3 \cdot ||c_{\mathbf{f}_1} - c_{\mathbf{f}_2}||_2 &+ \\ w_4 \cdot |B_{\mathbf{f}_1}^{\max} - B_{\mathbf{f}_2}^{\max}| &+ \\ w_5 \cdot |B_{\mathbf{f}_1}^{\mathrm{mean}} - B_{\mathbf{f}_2}^{\mathrm{mean}}| &+ \\ w_6 \cdot \sqrt{1 - |\mathbf{v}_{\mathbf{f}_1} \cdot \mathbf{v}_{\mathbf{f}_1}|} & \quad , \end{aligned} \quad (11)$$

where w_i represent the weights of each feature. Again, a smaller similarity value indicates more similar objects. We used a nonlinear measure for the the orientation to emphasise large orientation dissimilarities.

5 Shape retrieval

5.1 Normalisation

The individual feature distances in (11) may have vastly different values, so the distances need to be normalised in order to not let any single feature overshadow the other ones. We used the Gaussian normalisation instead of normalising based on the minimum and maximum values. This reduces the bias that would be introduced by a few unusually large distances.

For all auroral object pairs f_i and f_j, $\forall i \neq j$, we compute the similarities $d(\mathbf{f}_i, \mathbf{f}_j)$ of each individual feature such as (5). Then we compute the mean μ_{ij} and standard deviation σ_{ij} for each feature. Now the normalisation

$$d'(\mathbf{f}_i, \mathbf{f}_j) = \frac{d(\mathbf{f}_i, \mathbf{f}_j) - \mu_{ij}}{3\sigma_{ij}} \tag{12}$$

and

$$d''(\mathbf{f}_i, \mathbf{f}_j) = \frac{d'(\mathbf{f}_i, \mathbf{f}_j) + 1}{2} \tag{13}$$

brings 99% of all $d''(\mathbf{f}_i, \mathbf{f}_j)$ in the range of $[0, 1]$. The out-of-range values are mapped to either 0 or 1, which does not introduce large errors: very dissimilar objects should be disregarded from the retrieval process in any case.

As the number of objects increases, forming all pairs for computing μ_{ij} and σ_{ij} becomes a formidable effort. One should note that our aim is to be able to search among millions of auroral images. We estimated $\hat{\mu}_{ij}$ and $\hat{\sigma}_{ij}$ for each feature by using 11 sets of 100,000 randomly picked object pairs. Fig. 4 shows the behaviour of average orientation distance between two auroral objects: initial fluctuations converge to $\hat{\mu}_{ij} \approx 0.1584 \pm 0.0005$ at 100,000 pairs. The estimated error figure is one standard deviation of the values from the 11 runs. The other normalisation coefficients were determined similarly.

5.2 Updating the weights

The relevance feedback is based on the user's contribution in assessing whether the retrieved objects are relevant. This relevance is reflected by updating the weights in (11). We gave the user the choice of marking an object either "relevant" or "no comment". Those objects marked relevant — including the original query object— are used in updating the weights in (11). We used standard-deviation-based approach as in [7].

Intuitively, if an object feature provides a good representation of the object's appearance, the values in all relevant object shapes are close to each other. Respectively, large differences indicate poor representation. We computed the pair-wise distances of all relevant objects for each feature and analysed the standard deviations. The new weights are simply:

$$w_i' = \frac{1}{\sigma_i}, \tag{14}$$

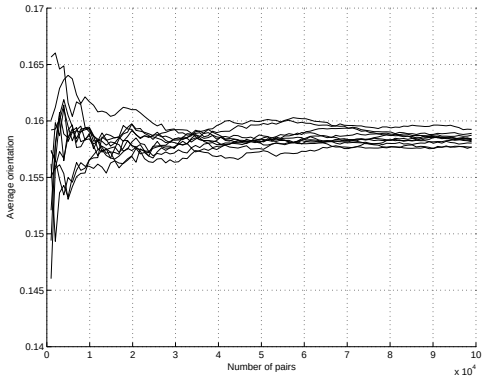

Figure 4. The behaviour of 1,100,000 randomly selected object pairs distances: the average distance in orientation settles down to approximately 0.158 in all 11 trials, which provides us confidence of the estimated means and standard deviations used in the normalisation process. Each curve present one trial.

where σ_i is the standard deviation of feature i. The updated weights are then normalised:

$$w_i'' = \frac{w_i'}{w_T'}, \tag{15}$$

where $w_T' = \sum_i w_i'$.

5.3 Implementation

We used Matlab for extracting the shapes from the auroral images. Each of the 340,000 auroral images was first analysed to determine whether it contained aurora or not by using (1). Those with aurora were then processed to extract one or more auroral objects. This phase took approximately 34 hours on a standard PC-computer with a 1.5GHz clock frequency. The process of determining the parameters for feature normalisation took approximately 22 hours. The object information was then stored into a PostgreSQL relational database.

In our implementation, a web browser provides the graphical interface while the PostgreSQL database management system performs the actual retrieval calculations. One (retrieval) iteration takes approximately one minute if the similarity of every object in the database is evaluated. As the auroral scientists are often interested in the occurrence of a shape at a particular time of day, we also provided one-hour time slots to speed the retrieval by limiting the number of objects to be compared. This reduces the retrieval time to approximately ten seconds.

5.4 Flow of the retrieval process

The user can initiate the retrieval by using one of eight manually drawn "interesting" auroral shapes and specify

whether the time of day is irrelevant or provide a one-hour window in which to search for similar shapes. Alternatively, the user can browse randomly selected shapes and choose one as the query image. The first query uses equal weights w_i and the objects in the database are ranked by their total similarity as determined by (11). The 20 most similar objects are returned.

In the next step, the user marks the relevant objects and restarts the retrieval. Based on the marked objects and the original query object a new weighting is determined using (14)-(15). Then the objects in the database are ranked again using (11), and as before the 20 most similar objects are returned. All already shown relevant and non-relevant objects are rejected from the retrieval process, although they are used in determining new weights for the retrieval.

6 Results

The images in the database represent auroral activity in Gillam between 1986 and 2001. The actual contents of the images are not known: summary plots generalise the data and cannot be used to locate specific shapes. This is especially the case with short-lived phenomena such as north-south aligned auroral structures. However, this is a perfect opportunity to test the retrieval capability with a useful application to auroral research: an auroral scientist used the CBIR to locate north-south structures. After one iteration all retrieved images were judged to be relevant (ie., containing shapes that were subjectively classified as north-south aligned auroras) with the orientation (40%) and FD-shape (44%) being the dominating features.

Fig. 5 shows the initial results after querying for a hand-drawn shape of the letter Y. No auroral theory predicts a shape like this, and, not surprisingly, there were no close matches. The closest shape was another artificial shape (hand-drawn letter T). However, there were a small number of objects that shared one property with the letter Y: three distinct brightness branches. Fig. 6 shows the retrieval after relevance feedback: there are clearly more shapes with this property.

7 Discussion

The lack of true content classification makes it very difficult to objectively evaluate the performance of the retrieval system. In our previous work [5] we experimented with a small set of manually classified auroral shapes (350 objects) with good retrieval precision when using only the FD-based shape representation. Nevertheless, our initial tests in the larger data set with artificial shapes appear very promising. We also find the relevance feedback mechanism useful in assessing different features' suitability for computer vision applications in auroral research: the features consistently having larger weights in (11) after training are strong candidates for, e.g., data mining, shape modelling and analysing the temporal behaviour of auroral shapes. In fact, the feature $B_{\mathbf{f}}^{\max}$ representing the maximum brightness value within an object appears to provide little additional information for retrieval purposes.

The CBIR system for auroral shapes presented in this paper is far from complete. The user's feedback interface needs improvement and additional features such as the ability to choose initial weighting for the different features. As this is the only existing implementation of this kind we know of, a common practise for describing the query auroral shape has not yet been formed. One obvious extension to the current feature selection would be the number of objects in the image: this would make it possible to locate patchy auroras, which consist of a number of irregular auroral blobs. Our simple implementation could also benefit from more sophisticated similarity learning which utilises the underlying shape distribution such as in [8]. Also, less heuristic weight adjustment approaches should be studied as in [9].

While we consider our auroral CBIR implementation a first step, the results will eventually teach us which features are the most useful ones for analysing auroral images. Furthermore, the shape information can be used to quantitatively describe the time-series of occurring auroral forms. At present, researchers can only qualitatively describe the temporal evolution of auroras, noting for example that the aurora evolved from one type to another using very subjective criteria. By having a numeric representation that was successfully used in retrieving similar shapes, we now have the opportunity to create time series that quantitatively reflect the temporal evolution of form. This in turn will allow better combined use of the observations — the currently existing as well as future multi-million image datasets — theory, and simulations to explore the underlying physics.

Acknowledgements

CANOPUS is operated by the Canadian Space Agency. This work was financially supported by the Alberta Ingenuity Fund (Canada) and the Natural Sciences and Engineering Research Council (Canada).

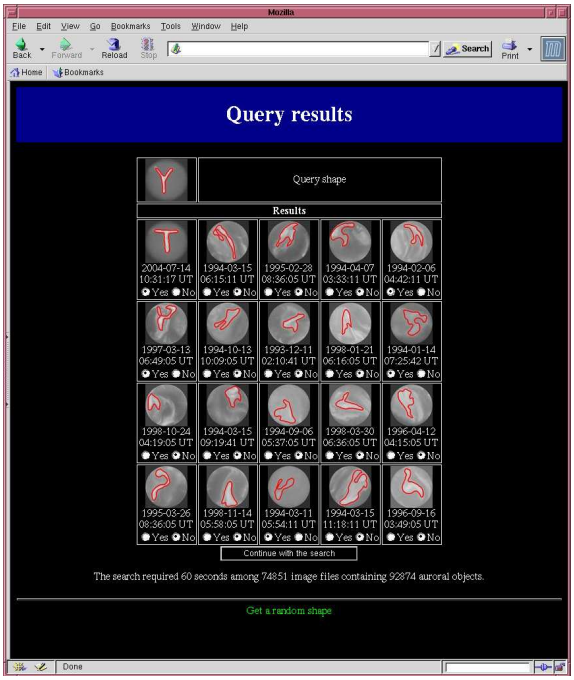

Figure 5. A screen grab of the results of the initial query. The query object was a shape resembling the letter Y. One of the retrieved objects is another artificial shape resembling the letter T. There are a few other shapes that could be considered distorted version of Y and have been marked relevant.

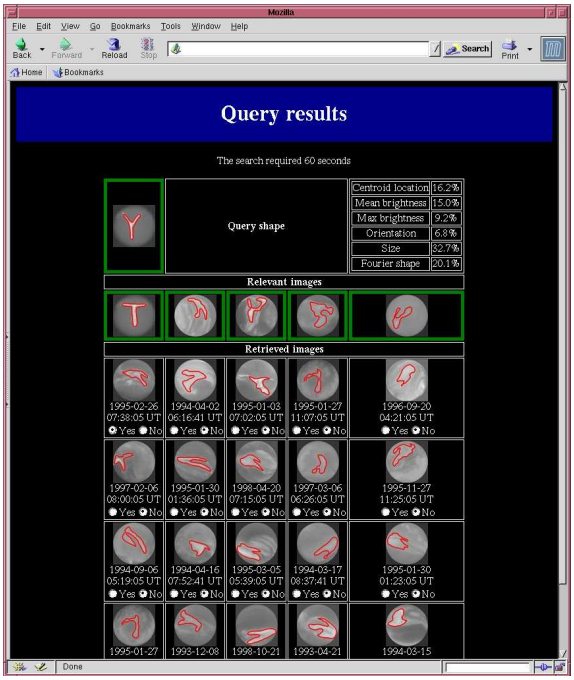

Figure 6. Based on the relevance feedback, the next retrieval uses updated weights and retrieves more "distorted" Y-shapes. The features can be ordered into an order of decreasing weight: size (33%, FD-shape (20%), location (16%), mean brightness (15%), maximum brightness (9%) and orientation (7%).

References

[1] Y. Rui, T.S. Huang and S.-F. Chang, Image retrieval: current techniques, promising directions and open issues, *Journal of Visual Communication and Image Representation*, 10, 1999, 39–62.

[2] A.W.W. Smeulders, M. Worring, S. Santini, A. Gupta, A. and R. Jain, Content-based image retrieval at the end of the early years, *IEEE Trans. Pattern Anal. Machine Intell.*, 22(12), 2000, 1349–1380.

[3] G. Rostoker, J.C. Samson, F. Creutzberg, T.J. Hughes, D.R. McDiarmid, A.G. McNamara, A. Vallance-Jones, D.D. Wallis and L.L. Cogger, CANOPUS - A ground-based instrument array for remote sensing of the high-latitude ionosphere during the ISTP/GGS Program, *Space Science Review*, 71, 1995, 743–760.

[4] M.T. Syrjäsuo and E.F. Donovan, Analysis of auroral images: detection and tracking, *Geophysica*, 38(1–2), 2002, 3–14.

[5] M.T. Syrjäsuo and E.F. Donovan, Content-based retrieval of auroral images — thousands of irregular shapes, *Proc. IASTED Int. Conf. on Visualization, Imaging, and Image Processing, VIIP 2004*, Marbella, Spain, 2004, 224–228.

[6] S. Shiffman, G.D. Rubin and S. Napel, Medical Image Segmentation Using Analysis of Isolabel-Contour Maps, *IEEE Trans. Med. Imag.*, 19(11), 2000, 1064–1074.

[7] Y. Rui, T.S. Huang, M. Ortega and S. Mehrotra, Relevance feedback: a power tool for interactive content-based image retrieval, *IEEE Trans. Circuits Syst. Video Technol.*, 8(5), 1998, 664–655.

[8] G.-D. Guo, A.K. Jain, W.-Y. Ma and H.-J. Zhang, Learning similarity measure for natural image retrieval with relevance feedback, *IEEE Trans. Neural Network*, 13(4), 2002, 811-820.

[9] T.S. Huang and X.S. Zhou, Imager retrieval with relevance feedback: from heuristic weight adjustment to optimal learning methods, *Proc. IEEE Int. Conf. on Image Processing*, 3, 2001, 2–5.