Automatic classification of auroral images in substorm studies

M. T. Syrjäsuo, E. F. Donovan, X. Qin, and Y.-H. Yang

Abstract: Millions of auroral images are captured every year by ground-based imagers. Even though the auroral appearance or "type" yields relevant information about the physical processes in the ionosphere and the magnetosphere, qualitative descriptions of auroras are typically used. Modern methods including those widely used in computer vision research can, however, make it possible to use objective and quantitative measures in analysing auroral appearance. We are currently developing techniques for automated auroral image analysis. In order to numerically compare auroral objects, we can either describe individual auroral shapes — such as arcs — or use statistical appearance models (texture). We demonstrate how one can use Fourier Descriptors to compare shapes extracted from auroral images. Also, using a recently developed texture analysis technique, we show how texture measurements can be used in classifying auroral type in a timeseries.

Key words: Aurora, pattern recognition, computer vision.

1. Introduction

Imaging the aurora by using ground-based optical instruments has long traditions in substorm research. The spectral, temporal and spatial resolution have been increasing, and now we are facing a problem of data: large imaging networks such as MIRACLE [10] and NORSTAR [3] produce millions of all-sky images annually. The situation will become worse with the launch of Time History of Events and Macroscale Interactions during Substorms (THEMIS) programme, which will produce over 100 million images every year.

The traditional data analysis in substorm studies uses actual measurements of physical properties (e.g. solar wind speed, electron density). Also, derived quantities are commonly used in order to understand the plasma processes in the magnetosphere and the ionosphere. Regardless of this quantitative information, the auroral image data are studied by using qualitative descriptors such as "bright auroral arcs" or "patchy aurora". Undoubtedly the "type" of the aurora yields relevant information about the physics. More importantly, a self-consistent global model should be able to predict this auroral type: otherwise our understanding, upon which the model is based, is not accurate.

Computer vision is a branch of computer science in which techniques for automated image analysis and processing are studied and developed. Automating the analysis makes it possible to browse through vast image sets, extract information and learn and recognise patterns. In [11], we used automated processing to obtain diurnal auroral occurrence statistics. While the actual results were not new, the analysis was: the statistics were based on 350,000 auroral images, from which an automated routine extracted information about whether an image contained aurora and if so what was its type. The type of aurora was learned from 258 manually labelled example images, after which the computer could provide a classification to all remaining images.

While we chose to use a three distinct auroral types (arcs, patchy aurora and Omega-bands) in [11], there is, of course, more variation in auroral types. In fact, the automatic classifier could only provide a clear category for 12% of the images that contained auroras. Also, the detection of Omega-bands was quite inaccurate, possibly because contrary to arcs and patchy auroras, we had a small number of examples for training the classifier.

There clearly is a need for more accurate mathematical treatment of auroral image contents. We can use example images for training a classifier to recognise certain types of auroras, but obviously learning the type categories from actual data would provide more objective type definitions. In the rest of this paper, we concentrate on describing the shape of an auroral object by using mathematical methods and demonstrate the use of texture measures for classication of auroral images.

2. Shape analysis

Extracting shapes is one of the most intensively studied problems in computer vision. Of course, there is no algorithm that works well for all applications. For our purposes, we have used a modified version of the isolabel-contour map algorithm from medical imaging [9]. The algorithm consists of four steps: 1) extract contours; 2) detect strongest edges; 3) score individual contours based on their overlap on edges; 4) choose non-overlapping contours with the highest scores. An example of detected auroral shapes can be seen in Fig. 1 and selection of extracted shapes is shown in Fig. 2. Details and practical applications of this algorithm can be found in [12, 13].

Once an auroral object is outlined, we can form a mathematical expression for the shape. Let x_i and y_i be the pixel coordinates $i = 1, \dots, N$ on the outline. Now, we can define a centroid shape signature

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

Received 11 May 2006.

M. T. Syrjäsuo and E. F. Donovan. Institute for Space Research, University of Calgary, Alberta, Canada.
X. Qin and Y.-H. Yang. Department of Computing Science, University of Alberta, Alberta, Canada.

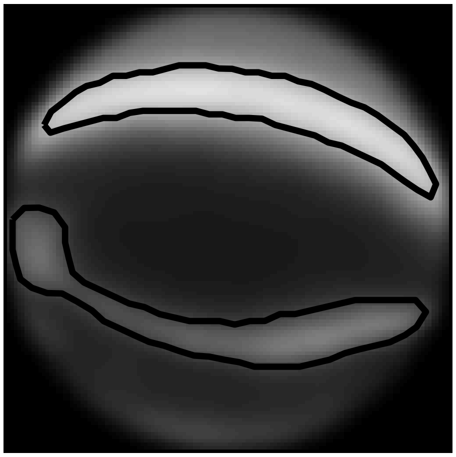

Fig. 1. Two salient auroral objects detected and outlined by the shape extraction algorithm.

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

$$a(k) = \sum_{i=0}^{N-1} r_i e^{-j2\pi(k-1)(i-1)/N}, \quad k = 0 \dots N-1, \tag{2}$$

where j is the imaginary unit. The use of centroid provides a translation invariant representation. We can represent the shape by first defining the Fourier Descriptors (FD) of the shape

$$\mathrm{FD}(k) = \left| \frac{a(k)}{a(0)} \right|, \quad k = 0 \dots N-1, \tag{3}$$

and then using a subset of these FDs to provide a more compact approximation of the shape:

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

where $M = 16$ has experimentally been found to provide a good approximation of the shape for comparison purposes.

Given two auroral shapes and their FD-representation, we can measure their similarity by using the Euclidian distance:

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

where $\mathbf{f}_1$ and $\mathbf{f}_2$ correspond to the two different shapes being compared and $||\cdot||_2$ denotes an L_2-norm. We further assert that small distances correspond to more similar shapes. Similarity, of course, is a complex human concept, but for practical purposes, this definition works surprisingly well. A selection of shapes was organised by their mathematical representation in Fig. 3 illustrating how well this approach captures the apparent similarity.

We have implemented a content-based image retrieval system based on the Fourier Descriptors. The system lets the user choose an initial search shape, after which images which contain similar shapes are returned. The system currently contains 20,000 extracted shapes and can be experimented with at `http://aurora.phys.ucalgary.ca/cbir/`.

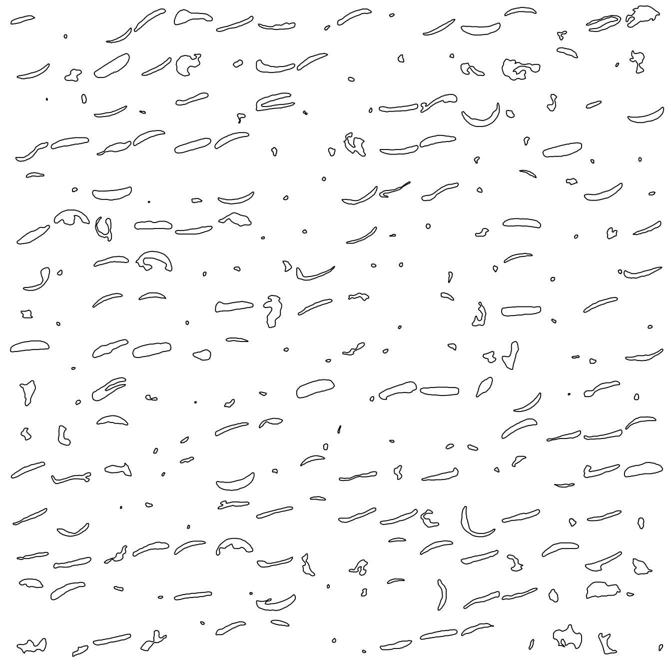

Fig. 2. A selection of extracted auroral shapes [13]. While the arcs are most obvious shapes, there are many other irregular shapes in different orientations and sizes.

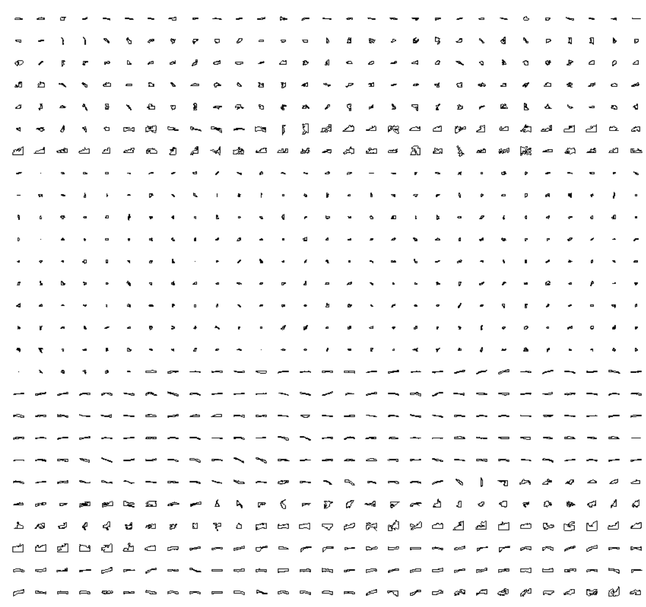

Fig. 3. Extracted shapes sorted by using their numerical representation.

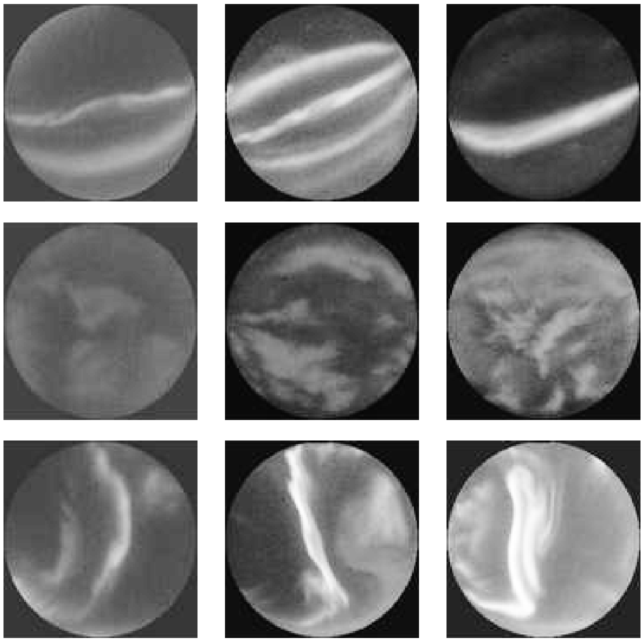

Fig. 4. Top row to bottom row: auroral arcs, patchy auroras and north-south structures. These all-sky images show north at the top and east on the right; the circular field-of-view captures the whole sky.

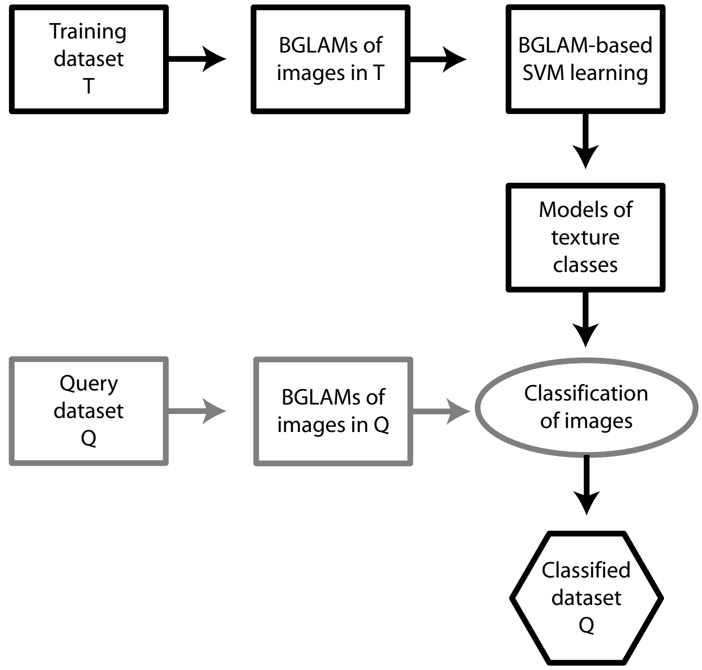

Fig. 5. An overview of the BGLAM-based algorithm for texture classification.

3. Aurora content as a texture

3.1. Gray level aura matrices

Texture can be defined as a characteristic property of any object or image. In the case of aurora, the patterns that appear in images share perceived similarities even if the individual shapes are not the same. Some of the commonly used terminology relies on texture: for example, patchy aurora usually refers to irregular auroral "blobs" whose characteristic sizes and blob distributions create a patchy appearance.

Fig. 4 shows three auroral categories — arcs, patchy aurora and north-south aligned auroral features. While the elements of arcs and north-south structures are similar, their orientation is different. Also their pattern is significantly different from patchy auroras.

There are mathematical methods which can be used to extract information about the texture. As with shapes, a texture distance measure can be utilised to classify auroras based on their general appearance in the images. One of these methods is based on gray level aurora matrices (GLAM).

As a generalisation of gray level co-occurrence matrix [2, 15], gray level aura matrix [4] has been used as a powerful tool for texture analysis, synthesis, segmentation and classification [5, 6, 7]. Among all the GLAMs, basic GLAMs (BGLAMs) are particularly important. In fact, BGLAMs are a basis of GLAMs and two images are the same if and only if their corresponding BGLAMs are the same — for the proof, see [8]. In other words, an image can be uniquely represented by and then faithfully reconstructed from its BGLAMs.

Based on the above BGLAM theory, we can use a BGLAM-based distance function for quantitatively measuring the similarity between texture images. The new distance function satisfies the important properties of non-negativity, symmetry, and triangle inequality, and thus is metric. Furthermore, one unique property of the BGLAM-based distance function is that it is one-to-one. Namely, a zero value of the distance measure will guarantee that the two images are identical. Since the distance function is continuous, the one-to-one property implies that if the distance of image Y from image X gradually changes (i.e. converges) to zero, image Y will gradually get close (i.e. converge) to X. For texture images, this one-to-one property guarantees that the smaller the distance value, the more similar the two texture images are. A distance measure without the one-to-one property cannot guarantee this.

3.2. Auroral textures

Texture classification can be done using a BGLAM-based approach (Fig. 5). Given an unseen texture image, the approach classifies it into one of the pre-learned classes. There are two states in the algorithm: a learning stage and a classification stage. In the first stage, models of texture classes are learned from the BGLAMs of training examples using the Support Vector Machine [14], and in the second stage, a given texture image is classified into one of the pre-learned classes, to which the image has the largest signed distance.

We used a small number of sample all-sky images images with varying contents: 401 arcs, 355 patchy auroras, 52 north-south structures, 76 Omega-bands and 113 cloudy skies. These images were used in training a classifier to recognise the image contents. The accuracy of the classifier in the training set was very good (over 90% correct classification).

The classifier was then used in determining the class for all auroral images acquired during one night. In order to compare the classifications of these previously unseen images, we also

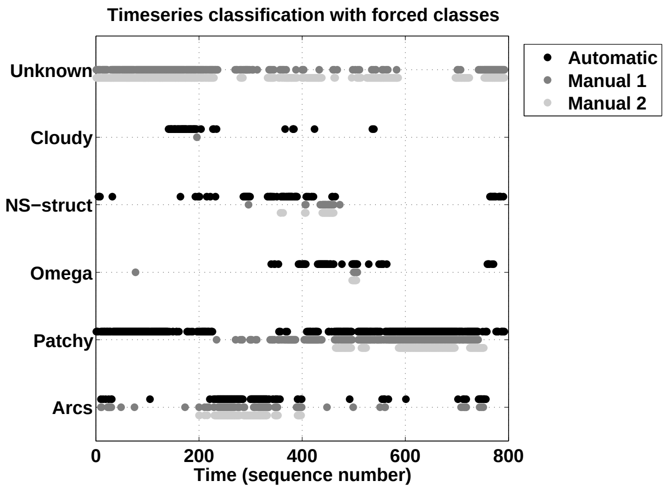

Fig. 6. One night of images as classified by the automatic method. For comparison, two manual classifications performed by two auroral image experts are provided. The auroral images were classified into "Cloudy", "North-south structures", "Omega-bands", "Patchy auroras" and "Arcs". The automatic classifier had no option to choose the "Unknown" auroral class.

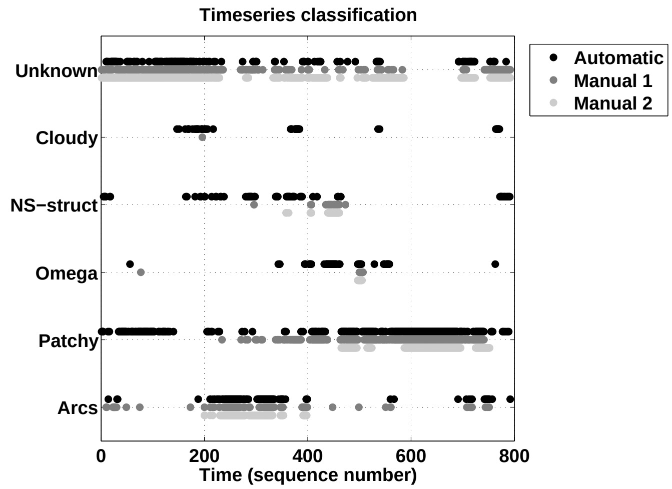

Fig. 7. The same time-series of auroral images as in Fig. 6. This time, however, the automatic classifier could use the "Unknown" auroral class.

provided manual classification for each image. The manual classification was performed by two auroral experts (Syrjäsuo and Donovan) who examined each of the images in random order to guarantee an independent auroral type for each image. Because the complexity of the image contents, the experts utilised a special category ("unknown") for images whose contents could not be classified unambiguously.

We performed two automatic classifications runs. In the first run, the classifier was forced to choose one auroral category, whereas the second run included an "unknown" category also in the automated classification.

The first run results are shown in Fig. 6. The overall agreements with the two experts were 42% ("Manual 1") and 34% ("Manual 2"). Not surprisingly, when including the unknown class in the automated classification, the results (Fig. 7) are noticeably better with 53% and 50% agreement.

While there are differences between the classifier's and experts' classes, there are two important observations: (1) the experts agreed on the class in about 70% of the images and (2) the experts chose the unknown class in almost 50% of all images. With those images that the experts did not classify as unknown, the automatic method is much more accurate with 72% and 81% correct classification.

4. Discussion

In this paper, we have presented some recent results from our efforts to develop automatic classification algorithms for auroral images. The essence of our approach is a common feature of all computer vision techniques: we use a training set and algorithms that map unclassified images into a hyperspace. The automatic algorithms classify images as similar if they are close together. The effectiveness of the algorithm is assessed by whether or not images that are close together in that space are actually similar in a meaningful way. This is not as subjective as it sounds: the idea is that a truly successful auroral image classification technique would group images of aurora caused by some common underlying physical process close together in the appropriate space. So, ideally, inverted-V arcs would be near other inverted-V arcs, polar cap patches near other polar cap patches, etc.

In particular, we have very strong motivations for carrying out this work. On the practical side, we are creating hundreds of millions of auroral images and are in the very beginnings of developing an auroral virtual observatory. We want to be able to attach content descriptors to every image in our data set, likely including that information in the overarching data base and meta data structures. This would greatly facilitate studies of auroral physics with these large cumbersome data sets. That being said, however, we have a much more important motivation, namely to use these classification algorithms to help us better understand the physics of the aurora and geospace dynamics.

For example, in this paper we have presented what is to our knowledge the first creation of time series of such classifications. This was done using texture analysis, applied to one night of data. If one examines Figs. 6 and 7, we see that there is an evolution through the night as the aurora evolves from patchy, to arcs, then NS-structures and Omega-bands, and then to patchy again. Qualitatively this evolution is well known in the literature as a common diurnal variation, as evidenced as far back as in Akasofu's early work (see eg., Fig. 1 of [11] which is a modified version of an earlier figure from [1]). What is new here is capturing this variation over the course of a typical night quantitatively. Further, we can see hints that the automatic algorithm is responding to transition between types in some meaningful way. In particular, the experts classified only a few images as Omega bands around image number 500. The automatic method classified images leading to that time also as Omega bands. Our idea is that the classifier is seeing some Omega-like features in the preceding images and responding to those.

The hope is that we will be able to create time series of quantitative and — more importantly — physically meaningful classifications of auroral images. In reality, the transition from arc, to NS-structures and Omega-bands, and ultimately patchy aurora is a repeatable consequence of the typical substorm looked at through one all-sky imager. This transition, then, reflects physics that we as a community are struggling to understand. These time series of auroral classification should prove to be an excellent tool when used in an assimilative way with global geospace models. For example, the above mentioned evolution reflects both the magnetospheric evolution in the substorm and the changing magnetosphere-ionosphere coupling. These changes in the system, if properly understood, should allow us to predict the changes in the aurora. These quantitative time series of image classification will be an essential ingredient in testing the output of global models.

References

1. Akasofu, S.-I., The auroral oval, the auroral substorm, and their relations with the internal structure of the magnetosphere, —em Planet. Space Sci., *14*, 587?-595, 1966.
2. Davis, L. S., Johns, S. A. and Aggarwal, J.K., Texture analysis using generalized cooccurrence matrices, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1(3), 251–259, 1979.
3. Donovan, E. F., Trondsen, T. S., Cogger, L. L. and Jackel, B. J., Auroral imaging in Canadian CANOPUS and NORSTAR programs, *Proc. of Atmospheric Studies by Optical Methods*, 109–112, 2003.
4. Elfadel, I. M. and Picard, R. W., Gibbs random fields, cooccurrences, and texture modeling, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(1), 24–37, 1994.
5. , Elfadel, I. M. and Picard, R. W., Miscibility matrices explain the behavior of grayscale textures generated by Gibbs random fields, *SPIE Conference on Intelligent Robots and Computer Vision IX*, 524–535, 1990
6. Picard, R. W. and Elfadel, I. M., Structure of aura and co-occurrence matrices for the Gibbs texture model, *Journal of Mathematical Imaging & Vision*, 2, 5–25, 1992.
7. Qin, X. and Yang, Y.-H., Similarity measure and learning with gray level aura matrices (GLAM) for texture image retrieval, *Proc. IEEE Conf. Computer Vision and Pattern Recognition, CVPR-04*, 326-333, 2004.
8. Qin, X. and Yang, Y.-H., Basic gray level aura matrices: theory and its application to texture synthesis, *IEEE Int. Conf. Computer Vision*, 128–135, 2005.
9. Shiffman, S., Rubin, G. D. and Napel, S., Medical image segmentation using analysis of isolabel-contour maps, *IEEE Transactions on Medical Imaging*, 19(11), 1064–1074, 2000.
10. Syrjäsuo, M. T., Pulkkinen, T. I., Janhunen, P., Viljanen, A., Pellinen, R. J., Kauristie, K., Opgenoorth, H. J., Wallman, S., Eglitis, P. , Karlsson, P., Amm, O., Nielsen,E. and Thomas, C., Observations of substorm electrodynamics by using the MIRACLE network, *Proc. the Fourth Int. Conf. Substorms, ICS-4*, 111–114, 1998
11. Syrjäsuo, M. T. and Donovan, E. F., Diurnal auroral occurrence statistics obtained via machine vision, *Annales Geophysicae*, 22(4), 1103–1113, 2004.
12. Syrjäsuo, M. T. and Donovan, E. F., Content-based retrieval of auroral images — thousands of irregular shapes, *Proc. IASTED Int. Conf. Visualization, Imaging, and Image Processing, VIIP 2004*, 224–228, 2004.
13. Syrjäsuo, M. T. and Donovan, E. F., Using relevance feedback in retrieving auroral images, *Proc. IASTED Computational Intelligence, CI 2005*, 420– 425, 2005.
14. Vapnik, V. N., Statistical learning theory, John Wiley & Sons, New York, 1998.
15. Zucker, S. W., Finding structure in co-occurrence matrices for texture analysis, *Proc. Computer Vision and Graphic Image Processing*, 12, 286–308, 1980