Segmentation and Tracking of Ionospheric Storm Enhancements

Matthew P. Foster[a] and Adrian N. Evans[b]

Department of Electronic and Electrical Engineering, University of Bath, UK

ABSTRACT

A two-stage method for segmenting and estimating the motion of ionospheric electron density enhancements is presented. These enhancements occur during geomagnetic storms, when matter ejected from the sun enters the upper atmosphere. The first stage of the proposed method segments the enhancements in the images by using a tuned attribute morphology filter, based on feature contrast. A temporal feedback mechanism is also employed to improve the temporal stability of the outputs. The second stage makes use of shape contexts to calculate boundary matches. These matches are then used to estimate motion vectors and a post filter is applied to detect and correct anomalous vectors. Overall, the scheme produces data-driven estimates of the motion of ionospheric electron density enhancements that can be used to study the evolution of ionospheric storms. Illustrative results are presented for a sequence of images from the Halloween storm of 2003. The results of this technique can also be used to validate models and data from other instrumentation which images the ionosphere during storms.

1. INTRODUCTION

The ionosphere is the region of the atmosphere, extending from an altitude of approximately 50 km and to over 1000 km. In this region, free electrons can exist for short periods of time, and form an electrically conducting plasma. These electrons are liberated when extreme ultra violet (EUV) light from the Sun ionises neutral atoms in the atmosphere. The free electrons then form a plasma which has various effects on electromagnetic waves, including delaying or blocking their propagation at certain frequencies. The extent of the ionosphere is dependent on a variety of variables, most importantly the diurnal cycle, as the number of electrons is dependent on the amount of sunlight. During the day, layers of ionisation may be found as low at 50 km whereas at night there is less photo-ionisation and electrons recombine with positively charged ions, leading to a depletion in ionisation. At this time, the ionosphere can be considered as starting at over 200 km. The vertical extent of the ionosphere is marked by its merging with the solar wind at very high altitudes. One of the most important measurable quantities associated with the ionosphere is the number of electrons in a given volume – or the electron density. This density changes in response to many stimuli in addition to the diurnal cycle, and in particular to the injection of extra material from the Sun. When large amounts of material are ejected from the Sun, the Earth's magnetic field and the ionosphere are disturbed. This event is known as a geomagnetic storm.

Ionospheric enhancements are increases in electron density that occur during geomagnetic storms, such as the event during Halloween of 2003. During this, and other similar events, patches of enhanced electron density, known as storm enhanced density (SED), travel across the northern polar cap towards the night-side of the globe in a convective pattern known as a tongue of ionisation (TOI). These complex convective patterns are the subject of various scientific studies.[1,2]

The traditional approach to analysing convection patterns of this type is to use electric and magnetic fields models, such as that proposed by Weimer[3] and the IGRF.[4] However, model-based approaches are only of limited use in studying specific storms as, by nature, they tend to reflect normal, quiet conditions. This paper addresses this problem and proposes a data-driven approach for analysing the motion of SED which can be used to study the evolution of ionospheric enhancements during storm conditions.

Email [a]m.p.foster@bath.ac.uk and [b]a.n.evans@bath.ac.uk. M. P. Foster is funded by an EPSRC doctoral training award.

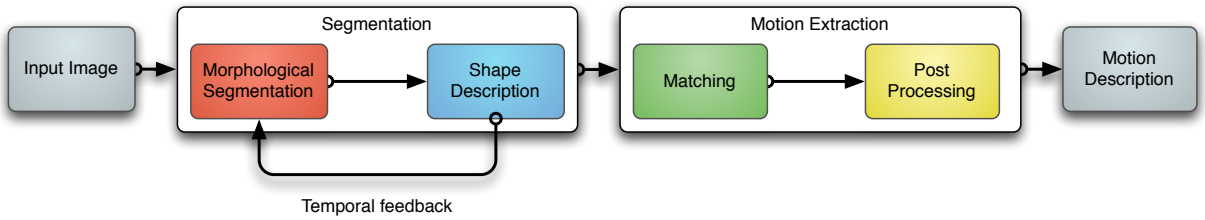

Figure 1: The processing pipeline for the TEC segmentation and matching scheme.

More specifically, we describe a two-stage technique that applies image processing methods to time-varying sequences of images of the total electron content (TEC) of the ionosphere, in which regions of SED can be identified as having a TEC value above the background level. The first stage of our approach identifies SED in the TEC images using tuned morphological attribute filters, followed by a watershed segmentation. Temporal feedback is also used to enforce a higher degree of consistency between the segmentation results for each frame in the sequence. The regions of SED are synonymous with the TOI and once their boundary has been identified the second stage uses the shape correspondences between successive frames to estimate the inter-frame motion and hence the underlying convection pattern. The processing pipeline for the complete process is presented in Fig. 1.

The format of the remainder of this paper is as follows. Section 2 describes the data used throughout this study, including the specific problems that result in very low-resolution images of the regions under study. Section 3 discusses the development of the segmentation technique, including information on how filters were tuned to improve performance using segmented images and a temporal feedback methodology. Section 4 presents the method used to infer inter-frame motion, by applying shape contexts to the segmented boundaries. Finally, section 5 discusses the conclusions that can be drawn from this study, as well as future work on enhancing performance and quality, and uses for the output vectors.

2. ELECTRON DENSITY DATA AND THE TONGUE OF IONISATION

The data used for this study consists of TEC maps which were produced using the MIDAS software from the University of Bath.[5] MIDAS performs tomographic inversions using path measurements of electron content between global positioning system (GPS) receivers and satellites. The maximum resolution of output field that can be produced by any given inversion algorithm is heavily dependent on the number of available path measurements, which means that the number of GPS receivers is an important factor in determining output resolution. There is a strong correspondence between population density and the number of GPS receivers in a given area, such that the resolution of the TEC images varies with geographic location. For example, over North America, where the population density is relatively high, there are over 500 GPS receivers operated by the IGS, whilst in more sparsely populated areas, such as the polar regions (considered here), the number of receivers is considerably lower and so the resolution of the TEC maps that can be produced is low.

The convective pattern to be analysed in this study is a region of SED known as the TOI. After its initial appearance at lower latitudes, the TOI soon moves into the northern polar region, so images of the TOI tend to be characterised by their low resolution. The images used in this study were formed by inverting GPS path measurements on to a thin shell at an altitude of 400 km to produce images of 25×25 pixels in size, corresponding to a grid covering a 100° square region, centred on the north pole. Fig. 2 shows the area covered by the image data, and the direction of the incident solar radiation. Coastlines and regions are marked to demonstrate how the grid is positioned. The thicker black lines mark the pixel extents and show how the images are warped when projected onto the globe. Fig. 2 also shows an example TEC image of the polar region in which the TOI can clearly be identified. This region consists of a three-dimensional volume with a higher than normal number of electrons. The overall aim is to track this feature as it moves through the image sequence. This motion is expected to begin in the bottom left, and move up to the top right over the course of the storm, corresponding to the convective pattern which has previously been described.[1,2]

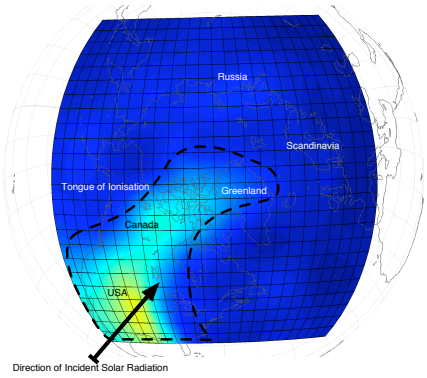

Figure 2: A schematic diagram showing how the displayed images relate to Earth's geography, and the direction of incident solar radiation. The direction of the sun will precess as the Earth rotates during the 3 hour period.

The data used to illustrate our scheme for segmenting and tracking the TOI consists of a sequence of 25×25 pixels images covering the period from 20:00 to 23:00 UTC on 30^{th} October 2003, with a temporal resolution of five minutes. Fig. 3 shows four frames from the sequence, with a time separation of 50 minutes between each pair. These clearly show a growing region of enhancement, in Fig. 3a—3c, which then splits off in Fig. 3d. Our aim is to study the evolution of this storm using the two-stage segmentation and motion analysis scheme described in Sections 3 and 4 respectively.

3. SEGMENTING THE TONGUE OF IONISATION

The TOI images can be considered to be composed of three main components: regions of SED corresponding to significant regional maxima; noise and artefacts corresponding to other maxima (and minima), and the remainder of the image which can be considered background. The aim of the segmentation stage is to correctly discriminate between the TOI and the other image components. Once the images have been segmented, it is possible to extract information about the enhancements, including their boundaries which can then be used for estimating the motion of the TOI, and hence studying its evolution.

A simple technique for separating objects from background or clutter is optimal thresholding.[6] This approach is appropriate when the object of interest has a different range of values than the background, as is the case with the TOI. It is also designed for use on data with a bi-modal intensity distribution. However, histograms of the TEC data for the four images from Fig.3, shown in Fig. 4, show that the TEC data is characterised by a distribution with only a single peak, which corresponds to the background value. The TEC values which correspond to the TOI occur in the tails (on the right hand side) of the illustrated distributions. This, coupled with the presence of noise, and the fact that the maximum TEC value changes over time means

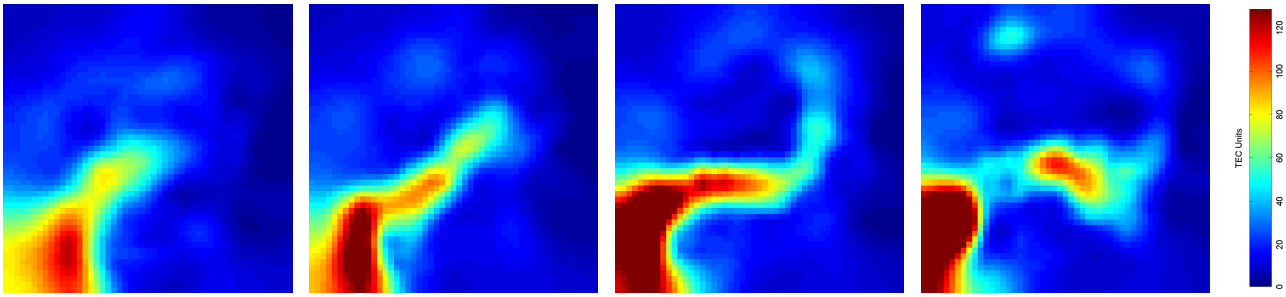

(a) 20:00 UTC (b) 20:50 UTC (c) 21:40 UTC (d) 22:30 UTC (e)

Figure 3: Example false-colour frames from the images sequence from Halloween 2003. Images (a)—(d) are each separated in time by 50 minutes, and are up-sampled by two. The colour-scale is shown in (e).

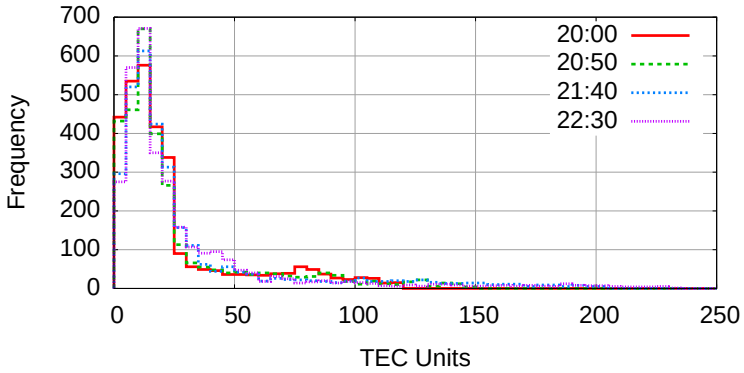

Figure 4: Histograms showing the distribution of TEC values in the data. The peaks on the left represent background data (low values), and the area on the right represents both noise and the TOI.

that simple thresholding is doomed to failure. Other segmentation methods, such as those based on image texture are, also inappropriate due to the low resolution and smoothness of the data.

To overcome these problems, a segmentation scheme based on mathematical morphology (MM) is proposed. Morphological filters have a number of advantages over traditional filters, most importantly shape preservation and idempotence. MM provides a number of tools for feature extraction and segmentation and many of these have been applied to remotely sensed data. For example, Soille and Pesaresi[7] (2002) provide a topical review of MM applications in remote sensing, ranging from segmentation to skeletonisation. Pesaresi and Benediktsson[8] (2001) propose a morphological segmentation method using the derivative of the morphological profile with residuals from opening and closing by reconstruction. This method is extended to multiple scales by using increasingly large structuring elements. More recently Akçay[9,10] (2007,2008) designed a hierarchical MM segmentation technique for remotely sensed imagery. This method creates morphological profiles by applying structuring elements of increasing size to each spectral band. Image patches extracted at different scales are then merged to create meaningful objects using spectral homogeneity and connectivity to optimise the groupings.

Whilst the segmentation schemes proposed in Refs. 7 to 10 were developed for application to very high-resolution (VHR) images, the images in this study are characterised by extremely low resolution. In these images, the lack of texture and low size mean that multi-scale segmentation methods are unsuitable. To help overcome these problems a segmentation scheme based on attribute morphology is proposed. Attribute morphology is the extension of traditional greyscale morphology to use more general feature information in the opening and closing operations. Typical attributes include area (the most commonly used feature), moments, orientation, contrast and many other measurable properties of image structure. Here, attribute morphology is advantageous when compared with standard morphology as it does not require the use of structuring elements, all of which would be large when compared to the dimensions of the images, and so introduce significant regions of noise. It also allows the filtering action to be finely controlled by attributes other than shape and size, producing segmentation results that can be tuned according to the properties of the features under consideration. Another useful property is that of shape preservation, a property shared with reconstruction based methods.

Contrast was chosen as the most suitable attribute for segmenting the TEC data due to its low sensitivity to gradient relative to other attributes, such as area. Whilst the peak SED value within the TOI shows considerable variation over the duration of the storm, the background TEC value is much more stable. Performing an attribute closing with a suitable contrast limit essentially flood-fills any basins with a contrast value less than the chosen attribute limit.

In order to choose a suitable contrast value with which to operate the morphological filter, ground-truths were generated by segmenting the images by hand. These were necessary because of the noisy nature of the frames which meant that inter-frame differences were not necessarily consistent and that the sudden

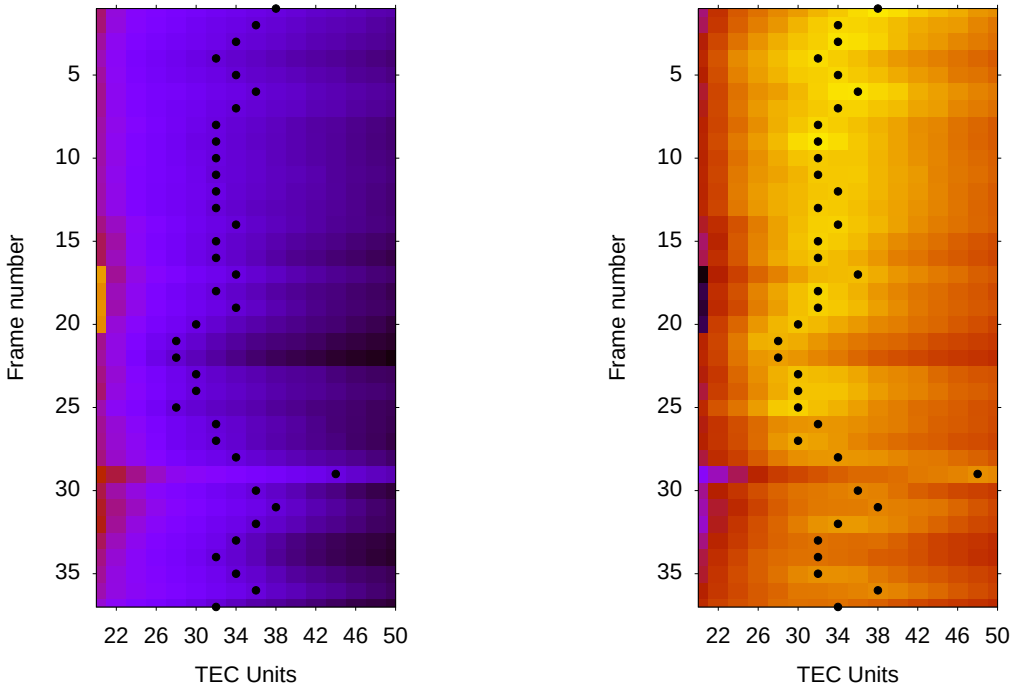

(a) Area Comparison (b) Correlation Comparison

Figure 5: Comparisons of area difference (a) and correlation (b) between frames segmented at various contrasts, and hand segmented frames. In both cases the contrast with the closest matches to the hand segmented frames was 33 TEC Units.

appearance and disappearance of small features was common. This process gave rise to a set of binary images, one for each frame, in which the salient regions were marked.

To find the contrast value for each frame which gave rise to the results most closely matching the corresponding ground-truth images, the data were filtered using a range of contrasts, and the outputs compared to the ground truths. The optimal value found in this fashion was then used as the starting point for the segmentations. Fig. 5 shows two image maps illustrating the best contrast parameter for each frame, using two different criteria – area (darker [around the dots] means a lower difference between the test and hand-segmented frames), and correlation (brighter means a higher correlation coefficient between the test and hand-segmented frames). In both cases, similar behaviour can be seen; on average, the contrast value which gave rise to results closest to the hand-segmentations was 33 TEC units. Fig. 5 shows that the performance of the attribute-based segmentation is relatively insensitive to the choice of contrast value. Therefore, although the contrast value here was selected using hand-segmented TEC images, the availability of more data would enable the starting contrast value to be estimated using training data with little overall loss in performance.

To improve the temporal consistency of the segmentation results over the image sequence, a temporal feedback system was used, along with the ground-truth derived contrast value. This system operated by performing an initial segmentation using a contrast closure, followed by a small-scale area-opening to remove small peaks which were considered clutter. This was then thresholded to give a binary mask identifying the TOI. The feedback system operated by calculating the change in foreground area between each frame and the frame preceding it, and then increasing or decreasing (as appropriate) the contrast level in small increments if the area differed by more than 10%. If the contrast was altered, the image was re-segmented. This process was repeated in an iterative fashion, until the segmented area was stable or until 10 iterations had occurred. This helped ensure the correspondence of objects between frames, leading to a smoother motion description.

The above process yielded results which are similar to the ground truth images whilst maintaining area

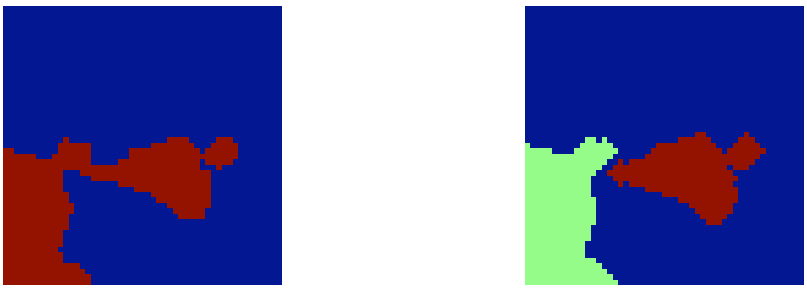

(a) (b)

Figure 6: Reducing under-segmentation of SED islands using the watershed transform. Example TEC image from 22:35 UTC on 30$^{\text{th}}$ August 2003 (a) before and (b) after watershed transform post-processing.

consistency between frames. However, in a few cases the segmentation was too generous, meaning that some areas which were separated in the hand-segmented images were left joined by the feedback process. This occurred when saddle-shaped structures joined features which would otherwise be separate. To mitigate this problem, a further stage was added to the process. By inverting the input images, masked by the segmented images, and then applying the watershed transform, thin saddle points joining areas were severed, as demonstrated in Fig. 6. The segmentation results for the images from Fig. 3 are shown in Fig 7. To obtain final TOI representations in a form which is smoothed and can be re-sampled, the image boundaries were traced and then converted to a smoothing-spline representation.

4. MOTION ESTIMATION BY BOUNDARY MATCHING

After boundaries have been extracted, the problem of estimating motion reduces to one of shape correspondence, i.e. given the boundaries for the TOI in two adjacent frames, the aim is to discover how one shape must be deformed in order to produce the other. The end result of this process is a series of vectors showing the correspondences between samples on each boundary, which can then be taken to represent cross-frame motion.

The method used to estimate the cross-frame motion was shape context (SC) matching, a technique for measuring similarity between shapes, using their boundaries.[11] Context matching consists of three steps, the first two of which can be used to estimate the relative motion of the two shapes. The entire SC process consists of: solving correspondences between points on the shape boundaries; using these correspondences to estimate a boundary alignment; summing the matching errors to calculate a similarity measure.

A SC is a histogram which describes the distribution of boundary points relative to an origin point (which is also part of the boundary). Shape contexts are constructed by picking a point on the boundary, and subtracting its coordinates from all others. These new values are then converted to polar form, and binned

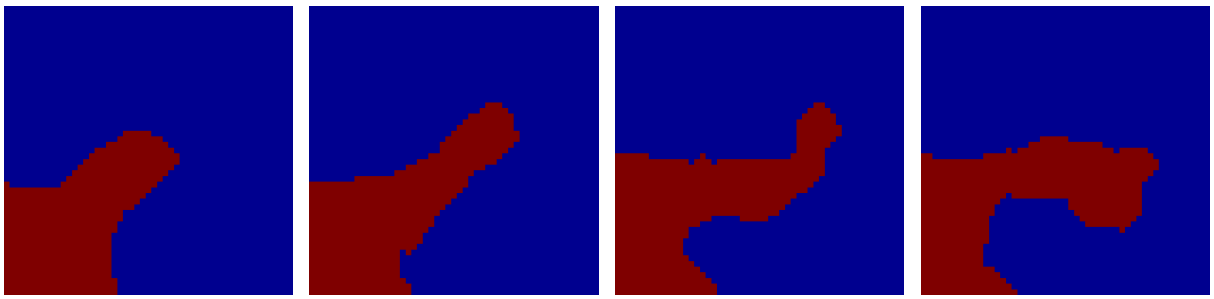

(a) 20:00 UTC (b) 20:50 UTC (c) 21:40 UTC (d) 22:30 UTC

Figure 7: Images from Fig. 3 segmented using morphological segmentation process described in the text.

to create a histogram. This process is repeated for each point on the boundary to produce a SC for each possible sample. Shape contexts must be created for both of the shapes being matched, so that a cost matrix $C_{ij} = C(p_i, p_j)$ can be computed for every possible pair of context origin points (p_i, p_j), where p_i is a point on the boundary of the first shape, and p_j is a point on the second. The cost of matching is calculated by comparing the context histograms and can be performed using a variety of different methods, such as the χ^2 cost, Earth mover's distance (EMD)[12] and diffusion distance (DD).[13] The latter two are particularly attractive as they allow 'leakage' across bins and thus provide more reliable costs. The DD measure was used in this study as it was found to be more robust than χ^2 and is also both conceptually and computationally cheaper than the EMD. It also has the added advantage of being physically relevant in this case, as it models anisotropic diffusion. The DD is based on a multi-scale Gaussian pyramid and is defined by:

$$K(h_i, h_j) = \sum_{l=0}^{L} |d_l(\mathbf{x})|, \tag{1}$$

where

$$d_0(\mathbf{x}) = h_1(\mathbf{x}) - h_2(\mathbf{x}) \tag{2}$$

$$d_1(\mathbf{x}) = [d_{l-1} * \phi(\mathbf{x}, \sigma)] \downarrow_2 \quad l = 1, \ldots, L \tag{3}$$

In the above equations, l represents the current layer of the Gaussian pyramid with L layers, $\downarrow_2$ denotes downsampling to half-size, σ is the standard deviation of the Gaussian filter $\phi(\cdot)$ and h_1 and h_2 are histograms. The metric used was the Euclidean distance, which means that the DD is a true histogram metric.

Constraints and extras costs can be added to this technique by modifying the matrix C. For example, a distance weighting could be added to help ensure that the minimisation only chooses points which are close in terms of boundary order. Measures relating to image properties could also be added in, although finding suitable measures is problematic as many such as curvature can be very noisy and introduce 'pits' into the cost matrix. These pits act to constrain the output matches by introducing local minima.

Once the cost matrix has been created, the problem becomes one of weighted bipartite matching, which can be solved using the Hungarian method, amongst others. This results in a permutation which describes the optimal mapping between shape boundary points and can be used to calculate the distances between corresponding boundary points on the two shapes. This is done by subtracting the coordinates of the first shape's boundary points from the corresponding points on the second. The result of this process is the set of vectors required to warp the first shape into the second; their relative motion.

As described in section 1, the image sequence corresponds to a patch over the polar cap. A consequence of this is that one side of the image is always in sunlight, and one is in darkness. Because of the nature of the ionosphere, the side in sunlight is continually injected with electrons and ions while the side in darkness undergoes depletion due to recombination effects. This results in the TOI being drained at its tip, in a fashion similar to the snout of a glacier, where ice melts and water drains away. If the rate of melting increases, the snout will appear to retreat whilst water will always flow away. If the position of the snout alone was being measured, the glacier would appear to be flowing backwards, which can clearly never be the case.

When recombination effects cause the tip of the TOI to retreat, motion estimation using boundary matching is unable to capture the true motion. This problem occurs towards the end of the image sequence under examination. Fig. 8b shows an example of this problem occurring due to a slight drop in value at the very tip of the TOI. Unfortunately, these retreating boundaries cannot be detected by simply examining the vector directions relative to other vectors in the same frame, as frames where this problem manifests itself tend to have very noisy vectors in general. For this reason, several techniques to detect and replace these retreating vectors were evaluated for suitability and effectiveness. The methods of detection included simply thresholding vectors (a method which fails to capture anomalous vectors, because it does not examine adjacent frames), subtracting each frame's vectors from the previous frame's vectors, interpolated to the correct positions, and then thresholding using magnitude, angle, or mean and standard deviation. In practice, using a combination of mean and standard deviation thresholding was found to be the most effective.

After detection, the marked vectors must be replaced. This can be done in several different ways, including:

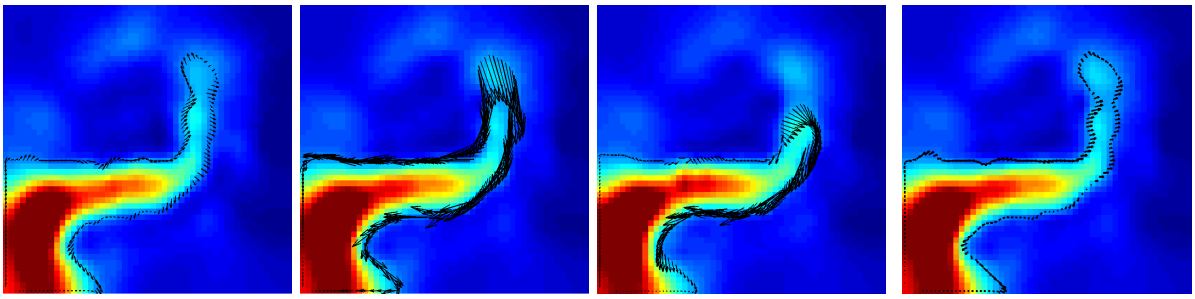

(a) 21:25 UTC (b) 21:30 UTC (c) 21:35 UTC (d) 21:30 UTC

Figure 8: Example frames illustrating the retreating snout problem. (a)–(c) Frames from before, during and after the problem occurs, respectively, and (d) the result of detecting the depletion and replacing all the vectors with vectors interpolated from those from the previous frame.

- Replacing the marked vectors with the vector-median[14] of the remaining vectors. As the frames containing the retreating boundary problem tend to contain noisy vectors, this was found to be unpredictable, and leads to unrealistic jumps in vector magnitude at the edges of regions of vectors being replaced.
- Replacing the detected vectors with new vectors generated using an affine model of the previous frame's vectors. Whilst this tends to give more suitable magnitudes than the previous method, the affine fit often gives vectors which point a seemingly *wrong* direction. It also leads to jumps in vectors at the edges of replacement areas.
- Replacing the detected vectors with new vectors generated by interpolating the previous frame's vectors onto the marked vector positions. This was carried out using both standard interpolation and regularisation. This gives good results provided the previous frame's vectors are reasonable. For this reason, processing should be done on a frame by frame basis.
- Replacing all of the vectors in the frame containing marked vectors with the previous frame's interpolated vectors. This is an effective way of getting smooth vectors over the entire boundary, provided the previous frame's vectors are smooth. The result of the replacing all vectors in affected frames in this fashion can be seen in Fig. 8d.

Of the detection and replacement methods tested, replacing *all* vectors based on thresholding vector magnitude statistics was found to be most effective due to the overall smoothness of the output. From an aesthetic point of view, it is advantageous because replacing the entire frame does not result in the discontinuities which are present in the outputs created using the vector median and affine methods.

Fig. 9 shows some examples of output frames with overlaid vectors. These vectors have been smoothed using the total replacement method described above, and represent the current level of output quality of the system described here. However, several problems still manifest in these outputs, the most important of which is *spurious border rotation*. This phenomenon occurs because of changes in boundary length, which means that samples across frames tend to lie slightly to one side of one another. This leads to a tendency for the output vectors to appear to be slightly rotated. One way to try and mitigate this is the use of up-sampling of output fields to increase their density by fitting warping surfaces.[11] This is carried out by separating the Cartesian components of the motion vectors, and interpolating these onto two separate surfaces.[15] One surface will then describe the magnitude of vectors in each axis, and vectors at new points can be found by taking values from each surface at the correct position. Using regularisation instead of interpolation will provide smoother outputs, which should be largely free of noise.

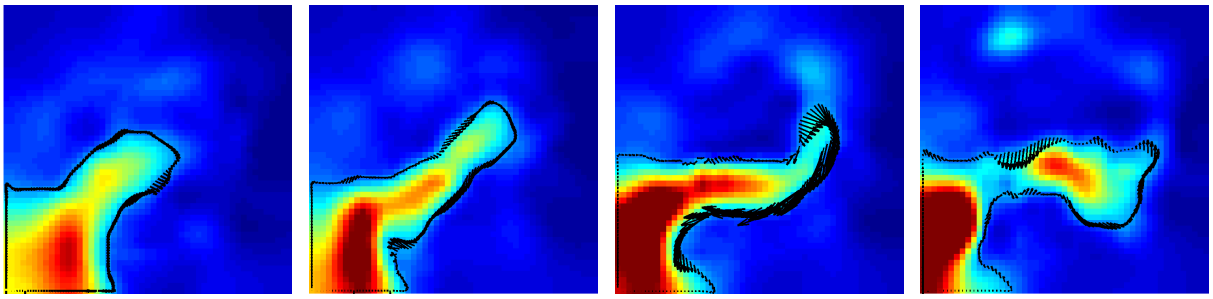

(a) 20:00 UTC (b) 20:50 UTC (c) 21:40 UTC (d) 22:30 UTC

Figure 9: Example false-colour frames from Fig. 3. Vectors are shown after post-processing by replacing all vectors in frames with detected retreating areas.

5. CONCLUSIONS

A new method for segmenting and tracking features in TEC images of the polar ionosphere during storm conditions has been presented. Unlike other applications of greyscale morphology to remotely sensed data, where the current trend is towards using VHR imagery, the application here is characterised by very low resolution images. Other major challenges result from the non-rigid motion of the TOI and the recombination effects. To overcome these problems an attribute morphology-based segmentation scheme has been proposed. This method uses the contrast attribute and produces segmentation results that are robust to changes in the maximum SED value. In addition, temporal feedback is used to improve the consistency of the segmentation results by automatically varying the attribute limit through the sequence. The results achieved by this approach are temporally consistent and show a good correspondence with hand-segmented ground truths.

To estimate the motion of the TOI an approach was developed which applied shape contexts to track boundary changes in the segmented TOI between successive frames. This method makes use of the shape contexts of smoothed boundaries (with equal numbers of samples) to estimate correspondences which can then be converted into motion vectors. Although this was found to be an effective approach when the TOI feature was stable, it gave erroneous results when the TOI boundary appears to retreat as a result of electron depletion. To mitigate this effect, a post-processing step was proposed in which the mean magnitude and standard deviations of a frame's vectors were compared with those from the previous frame. When certain thresholds were exceeded, the frame's vectors were entirely replaced by vectors from the previous frame, interpolated to the correct positions.

The sequences of TOI images and accompanying motion vectors produced by the segmentation and tracking scheme produce a useful source of information that, when combined with data from other instruments and models, can provide a more complete picture of the motion of SED during ionospheric storms.

Future work will focus on the extension of these techniques into an unsupervised system which does not require manual creation of ground truths. This could be facilitated through the building of a library of pre-segmented training data. Independent data sources such as vector data from the SuperDARN or EISCAT radars could also be used both for validation of our results and perhaps more interestingly, the production of improved motion fields by combining them with our estimates of boundary positions.

6. ACKNOWLEDGEMENTS

We gratefully acknowledge P. Yin at the University of Bath for supplying the TEC data.

REFERENCES

[1] Foster, J. C., Coster, A. J., Erickson, P. J., Holt, J. M., Lind, F. D., Rideout, W., McCready, M., van Eyken, A., Barnes, R., Greenwald, R. A., et al., "Multiradar observations of the polar tongue of ionization," *Journal of Geophysical Research* **110** (2005).

[2] Stolle, C., Schluter, S., Heise, S., Jacobi, C., Jakowski, N., Friedel, S., Kurschner, D., and Luhr, H., "GPS ionospheric imaging of the north polar ionosphere on 30 october 2003," *Advances in Space Research* **36**(11), 2201–2206 (2005).

[3] Weimer, D. R., "Models of high-latitude electric potentials derived with a least error fit of spherical harmonic coefficients," *J. Geophys. Res* **100**(19,595) (1995).

[4] Macmillan, S., Maus, S., Bondar, T., Chambodut, A., Golovkov, V., Holme, R., Langlais, B., Lesur, V., Lowes, F., Luhr, H., et al., "The 9th-Generation International Geomagnetic Reference Field," *Geophysical Journal International* **155**(3), 1051–1056 (2003).

[5] Mitchell, C. N. and Spencer, P. S. J., "A three-dimensional time-dependent algorithm for ionospheric imaging using GPS," *Annals of Geophysics* **46**(4), 687–696 (2003).

[6] Nixon, M. and Aguado, A., [*Feature extraction & image processing*], Academic Press, 2 ed. (2008).

[7] Soille, P. and Pesaresi, M., "Advances in mathematical morphology applied to geoscience and remote sensing," *Geoscience and Remote Sensing, IEEE Transactions on* **40**, 2042–2055 (Sep 2002).

[8] Pesaresi, M. and Benediktsson, J., "A new approach for the morphological segmentation of high-resolution satellite imagery," *Geoscience and Remote Sensing, IEEE Transactions on* **39**, 309–320 (Feb 2001).

[9] Akçay, H. and Aksoy, S., "Morphological segmentation of urban structures," *Urban Remote Sensing Joint Event, 2007* , 1–6 (April 2007).

[10] Akçay, H. G. and Aksoy, S., "Automatic detection of geospatial objects using multiple hierarchical segmentations," *Geoscience and Remote Sensing, IEEE Transactions on* **46**, 2097–2111 (July 2008).

[11] Belongie, S., Malik, J., and Puzicha, J., "Shape matching and object recognition using shape contexts," *Transactions on Pattern Analysis and Machine Intelligence* **24**(4), 509–522 (Apr 2002).

[12] Ling, H. and Okada, K., "An efficient earth mover's distance algorithm for robust histogram comparison," *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **29**(5), 840–853 (May 2007).

[13] Ling, H. and Okada, K., "Diffusion distance for histogram comparison," in [*Proc. CVPR*], 246–253, IEEE Computer Society, Washington, DC, USA (2006).

[14] Astola, J., Haavisto, P., and Neuvo, Y., "Vector median filters," *Proceedings of the IEEE* **78**(4) (1990).

[15] Foster, M. P. and Evans, A. N., "An evaluation of interpolation techniques for reconstructing ionospheric TEC maps," *Geoscience and Remote Sensing, IEEE Transactions on* **46**, 2153–2164 (July 2008).