Automatic Identification of the Main Ionospheric Trough in Total Electron Content Images

Gregory Starr[1], Sebastijan Mrak[1], Yukitoshi (Toshi) Nishimura[1], Michael Hirsch[2], Prakash Ishwar[1], and Joshua L. Semeter[1]

[1]Boston University
[2]SciVision, Inc.

November 24, 2022

Abstract

The main ionospheric trough (MIT) is a salient density feature in the mid-latitude ionosphere and characterizing its structure is important for understanding GPS and HF signal propagation, and identifying geospace phenomena such as the plasmapause boundary layer. While a number of previous studies have statistically investigated the properties of the MIT utilizing low-altitude satellite observations, they have been limited to latitudinal cross sections, and have not considered the inherent two-dimensional structure of the trough. In this work, we develop a regularized inversion method for identifying the two dimensional structure of the trough in Total Electron Content (TEC) maps. Because no ground truth labels exist for the MIT, we extensively characterize the behavior of the algorithm by comparing it to the method developed by \citeA{aa-2020}. We show that statistics computed on the resulting labels are robust to our choice of algorithm parameters and that we are able to match the results of \citeA{aa-2020} with a particular selection of the parameters. Without ground truth, these two properties provide much stronger verification than a comparison using a single parameter setting. In addition to enabling fundamentally different studies, our MIT labels are able to provide statistical MIT properties with higher resolution.

Automatic Identification of the Main Ionospheric Trough in Total Electron Content Images

Gregory Starr[1,3], Sebastijan Mrak[1,2,4], Yukitoshi Nishimura[1,2], Michael Hirsch[1,2], Prakash Ishwar[1], Joshua Semeter[1,2]

[1]Department of Electrical and Computer Engineering, Boston University
[2]Center for Space Physics, Boston University
[3]Now at JHU APL
[4]Now at Space Weather Technology, Research and Education Center, University of Colorado Boulder

Key Points:

- A method for identifying the main ionospheric trough in global total electron content (TEC) images is developed and validated.
- Statistics extracted from a 10-year dataset are found to be in good agreement with a SWARM-derived satellite model.
- Our technique enables statistical analyses of two-dimensional (lat, lon) trough properties with high resolution.

Corresponding author: Gregory Starr, `gstarr@bu.edu`

Abstract

The main ionospheric trough (MIT) is a salient density feature in the mid-latitude ionosphere and characterizing its structure is important for understanding GPS and HF signal propagation, and identifying geospace phenomena such as the plasmapause boundary layer. While a number of previous studies have statistically investigated the properties of the MIT utilizing low-altitude satellite observations, they have been limited to latitudinal cross sections, and have not considered the inherent two-dimensional structure of the trough. In this work, we develop a regularized inversion method for identifying the two dimensional structure of the trough in Total Electron Content (TEC) maps. Because no ground truth labels exist for the MIT, we extensively characterize the behavior of the algorithm by comparing it to the method developed by Aa et al. (2020). We show that statistics computed on the resulting labels are robust to our choice of algorithm parameters and that we are able to match the results of Aa et al. (2020) with a particular selection of the parameters. Without ground truth, these two properties provide much stronger verification than a comparison using a single parameter setting. In addition to enabling fundamentally different studies, our MIT labels are able to provide statistical MIT properties with higher resolution.

1 Introduction

As a partially conducting medium, the Earth's ionosphere affects electromagnetic waves in the radio frequency range (< 2 GHz). Crucial portions of our infrastructure rely on communication with satellites including weather forecasting, navigation, and recently, even internet. One particularly disruptive phenomenon is the main ionospheric trough (MIT). The MIT is a longitudinally elongated band of low electron density that forms just equatorward of the high latitudes characterized by particle precipitation (i.e., aurora). The trough affects radio signals due to the large electron density gradients and small-scale irregularities which form within it (Rodger et al., 1992; Kintner et al., 2007; Le et al., 2017; Mishin et al., 2003). Understanding the MIT's occurrence, geolocation, and electron density features will help us predict where interruptions might occur and allow us to mitigate them. Additionally, the trough is a major part of our geospace (ionosphere, magnetosphere), whereby the trough could be used as a tracer for geospace features such as the plasmasphere boundary layer (PBL) (Yizengaw & Moldwin, 2005). Therefore, a comprehensive global picture of the MIT traces geospace features imprinted to the ionosphere and has a potential to advance our understanding of this tightly coupled system.

The MIT has traditionally been characterized by three elements: the poleward wall, equatorward wall and the minimum in between. The poleward wall is associated with the equatorward boundary of the auroral precipitation region (Rodger, 2008; Rodger et al., 1992), and the equatorward wall is associated with the ionospheric footprint of the plasmapause (Zou et al., 2011; Rodger et al., 1992; Pedatella & Larson, 2010). Figure 1 shows an example of the three components of the MIT. The meridional (north-south) electron density gradient is typically much steeper at the poleward wall than at the equatorward wall (Spiro et al., 1978). While the MIT does not have a concrete definition, it is distinguished from other trough-like features in the ionosphere (e.g. high latitude trough, ring trough (Karpachev, 2019), light ion trough) based on its location (Rodger et al., 1992; Rodger, 2008). The MIT is observed most often in darkness and has an average width of about 5 to 10 degrees in latitude (Aa et al., 2020; Yang et al., 2015; Collis & Häggström, 1988). Observations in TEC maps (Zou et al., 2011) and sequential radar scans (Nilsson et al., 2005) show that the MIT is longitudinally elongated. Since it mainly occurs in darkness, its spread in local time is strongly correlated with season (Rodger, 2008), though this has not yet been directly quantified.

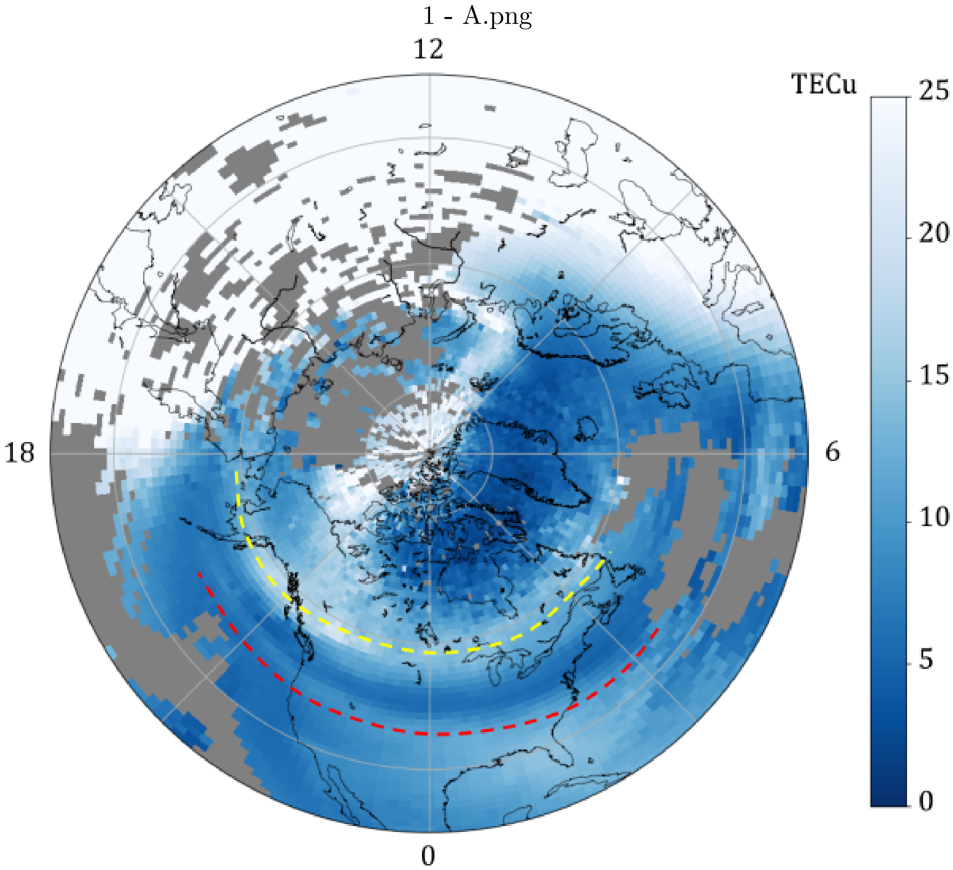

Figure 1. An example of the MIT in TEC measurements at 07:00 UT on 2014-02-19. It appears here as a dark band spanning the entire night-side mid-latitude ionosphere. Yellow and red lines show approximate locations of the poleward and equatorward walls of the MIT, respectively. Gray color indicates missing data.

The primary mechanisms that create and maintain the trough have been identified in past studies. The most commonly mentioned is called the stagnation mechanism, whereby the corotation and convection flows in the evening side ionosphere cancel each other out, forming a region where the plasma recombines during its long residence in darkness. This theory was analyzed, compared with measurements and determined to be plausible by Spiro et al. (1978) and Nilsson et al. (2005). Collis and Häggström (1988) found that the trough minimum typically occurs in a region bounded equatorward by the transition from corotating to convecting flow, and poleward by the electron precipitation boundary. A more general convection-based theory, of which the stagnation mechanism is a special case was explained by Quegan et al. (1989). They emphasized that the electron density at any location is due to the production and loss along the path on which that plasma travelled. Given that the convection-corotation pattern is very complex, two paths with very different histories can be brought close together, forming a trough. While substantial progress has been made on understanding the MIT formation mechanisms, the relative importance of the various mechanisms has not been fully established, and some of the most popular ionospheric models do not adequately reproduce its behavior (Yang et al., 2015). Additionally, adequate statistics of how the MIT behaves during heightened periods of geomagnetic activity have not been obtained and the full spectrum of MIT behavior, including what edge cases exist, is not yet known.

Algorithmic approaches for MIT identification have so far only been developed for one-dimensional data. Existing methods follow the same basic steps: (1) estimate a background value for their measurements and (2) threshold the ratio between the measurements and their background value. If the data is not already one-dimensional, it is first reduced to a latitudinal profiles. For latitude-altitude measurements from the European Incoherent Scatter radar (EISCAT), Ishida et al. (2014) averaged the electron density along magnetic field lines between 300 and 350km. For background estimation, they took the median of the upper half of the sorted electron density values, then they found troughs where the electron density fell to 20% below the background. A similar approach was taken by Voiculescu et al. (2006) for latitude-altitude measurements of electron density estimated by tomography. They averaged the electron density between 200 and 400km, then looked for regions where it dropped below 50% of the "outside value". Yang et al. (2015) used the same TEC dataset as us, but they computed latitudinal profiles by averaging TEC over the course of a day in two hour Magnetic Local Time (MLT) bins. They computed the background as the mean TEC between magnetic latitude 45 and 70 then determined the trough minimum from the minimum of each profile, i.e. they assumed the trough was present in every profile. A similar approach was taken by Pryse et al. (2006) except using TEC computed from tomography data. Aa et al. (2020) computed the background electron density measured by the Swarm satellites, then used a threshold of 50% to identify the trough. Finally, one exception is the work by Pedatella and Larson (2010), in which the authors defined the equatorward wall of the MIT as the location in a latitudinal TEC profile equatorward of the minimum where the latitudinal TEC gradient is -0.1 TECu / degree.

Utilizing one-dimensional data (satellites / latitudinal cross sections) rather than data with two or three dimensions has several limitations. The first is that it inherently has less data. For example, Aa et al. (2020) analyzed data from the Swarm constellation which has 3 satellites, each of which have an orbital period of roughly 1.5 hours. Over 10 years, they make roughly 350,000 measurements of the trough region in the northern hemisphere.

$$\frac{16\text{ orbits}}{1\text{ day}} \cdot \frac{2\text{ measurements}}{1\text{ orbit}} \cdot \frac{3,650\text{ days}}{10\text{ years}} \cdot 3\text{ satellites} = 350,400 \tag{1}$$

Two-dimensional data, for example Total Electron Content (TEC) maps, contains measurements at most local times. With a reasonable estimate of 90 sampled local times and 24 maps per day, this results in almost 8,000,000 measurements of the trough.

$$\frac{24\text{ maps}}{1\text{ day}} \cdot \frac{90\text{ measurements}}{1\text{ map}} \cdot \frac{3,650\text{ days}}{10\text{ years}} = 7,884,000 \tag{2}$$

Another benefit of two-dimensional data is that it allows for the determination of the covariances of the trough's parameters across local times. One-dimensional data inherently can only measure marginal distributions of the trough's parameters across local times. Finally, and perhaps most importantly, one-dimensional data is not suitable for observing the relationship between plasma convection and the trough. Since convection has already been established as a key mechanism involved in MIT formation, a two-dimensional dataset is crucial to further advancing our understanding of the trough.

In this work, we developed a methodology that automatically identifies MIT in TEC maps using tools from regularized inverse problems and image analysis. In this paper we develop the model then validate it by showing its agreement with the Aa et al. (2020) results as well as its robustness to parameter perturbations.

2 Methodology

Many parameterizations of the MIT are possible, each has strengths and weaknesses. We choose to utilize the image segmentation framework, in which each pixel of an input image is assigned a class from a discrete set. In the context of this work, the input is an image made up of TEC measurements discretized onto a latitude-longitude grid and the output is a binary image of the same size with the positive class corresponding to the identified MIT. Using this approach, the MIT has several advantages. It makes no assumptions about the shape of the trough, allowing for the MIT to be identified at multiple latitudes at the same local time. To ease identification, previous studies only looked for a single trough in any latitudinal profile. However, there is no guarantee that the MIT can only have one contiguous section at any given local time. Another advantage of this parameterization is that it results in a very flexible and widely useful final data product. With a small amount of additional processing, many different measurements of the MIT can be extracted, allowing the end user to adapt the dataset for their specific needs.

2.1 Data Preparation

The Madrigal GPS TEC dataset consists of over 20 years of TEC maps from 1998 to 2021. The line-of-sight TEC measurements from a global network of GPS receivers are converted from slant-TEC to vertical-TEC and collected into grid cells (1 degree latitude x 1 degree longitude x 5 minutes). The vertical TEC value at each cell is given by the median of all the measurements within it. For more information about the Madrigal GPS TEC dataset, see papers by Rideout and Coster (2006) and Vierinen et al. (2016).

Since the trough's location and behavior are influenced by the magnetic field, to prepare the TEC images, we convert the geographic coordinates used by Madrigal to the magnetic apex coordinate system (Richmond, 1995; Emmert et al., 2010). In the magnetic apex coordinate system, all points along a magnetic field line are mapped to the same coordinates. During this process, we also average several TEC maps together to improve their coverage. To form a single TEC image, we first convert each latitude-longitude grid point in 12 consecutive Madrigal TEC maps to apex latitude-MLT. We collect the converted points into cells on a regular magnetic apex latitude-MLT grid, and the values in each cell are averaged. We choose a grid cell size of (1 degree latitude x $2 \cdot \frac{24}{360}$ hours MLT x 1 hour UT). Larger amounts of time-averaging result in higher coverage in each map, but less time resolution. Finally, because the northern hemisphere has better coverage in the Madrigal dataset, we choose to limit our TEC images to magnetic apex latitudes above 30° north. The resulting shape of each image is (60 x 180) pixels corresponding to latitude rows and MLT columns. For the same reason we only use data from the years 2010-2020. Figure 1 shows an example of a single TEC map from our dataset.

2.2 Algorithm Description

Our method consists of 4 steps: (1) preprocessing, (2) scoring, (3) thresholding and (4) postprocessing. The preprocessing step involves spatially filtering the TEC data to remove structures that do not contribute to the MIT. In the scoring step, we assign a score to each pixel based on the preprocessed image and on the expected characteristics of the trough. In the third and fourth steps, we make initial decisions on each pixel via thresholding then clean up the thresholded image with heuristics. The output from each of these four steps is shown in Figure 2. We find that for two-dimensional data, the preprocessing by itself is insufficient to separate the trough from the rest of the ionosphere which is why we add additional processing steps. While two dimensional data may be more complicated to process, it allows a more accurate determination of the MIT because of the MIT's longitudinal coherence. For example, a small dip in a latitudinal profile of TEC might be a trough or might not, but if it is part of a longitudinally extended region of low TEC, then that indicates more strongly that it is part of the MIT.

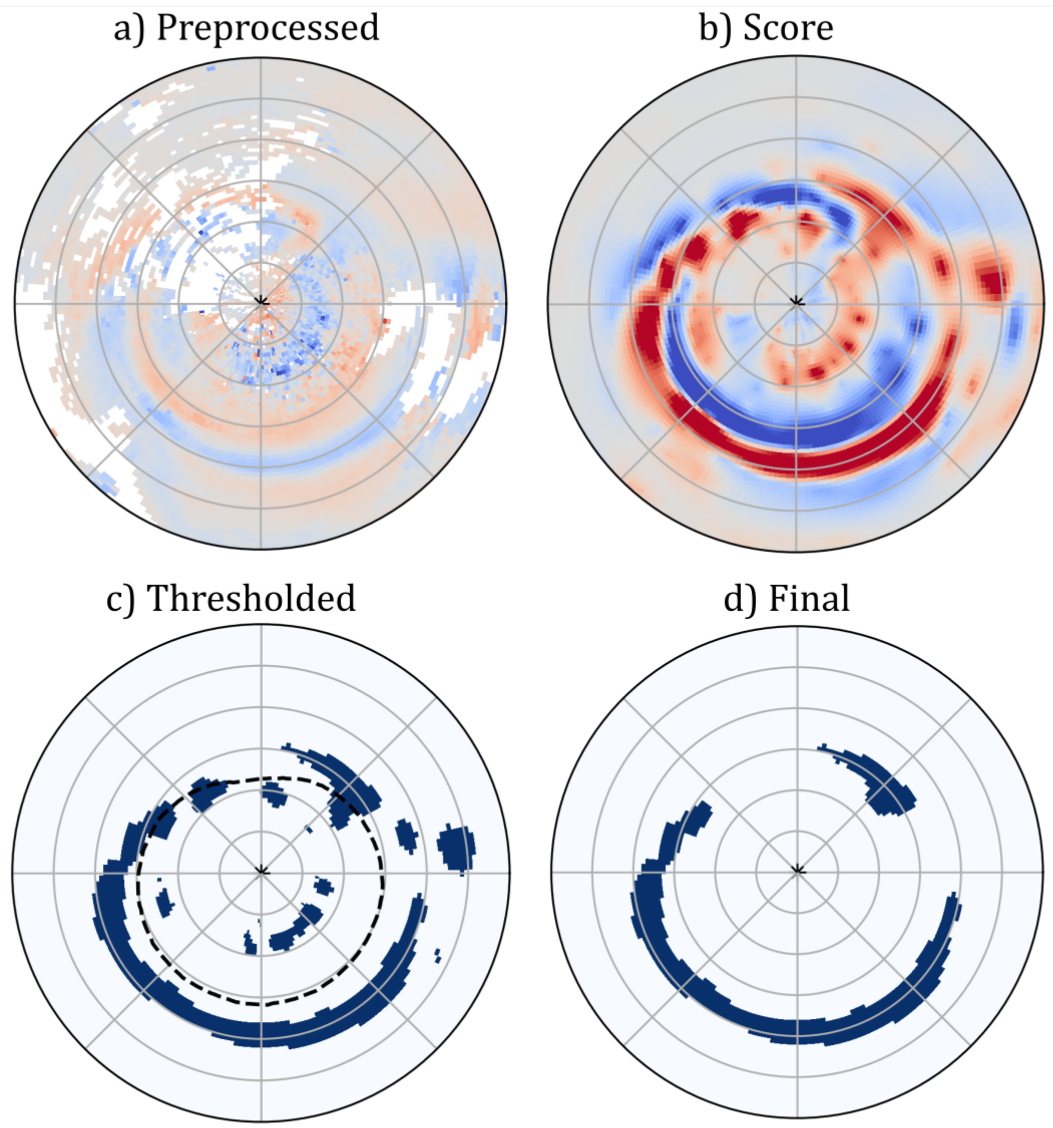

2.png

Figure 2. Algorithm overview. Input seen in Figure 1. In all panels, MLT noon is at the top, midnight is at the bottom, 6 is on the right and 18 is on the left. a) Output of preprocessing stage (section 2.2.1), red is high local TEC, blue is low local TEC. b) Output of scoring stage (section 2.2.2), red is high score, blue is low score. c) Thresholded score image, the equatorward boundary of the auroral oval is indicated with a dashed line. d) Output of postprocessing stage (section 2.2.3).

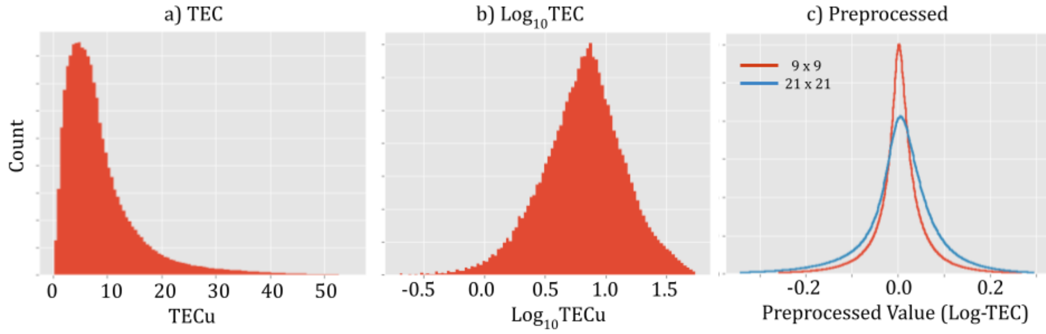

3.png

Figure 3. a) TEC distribution. b) Log TEC distribution. c) Preprocessed image pixel value distribution for two different selections of background filter size.

2.2.1 Preprocessing

The preprocessing stage takes a (60 × 180) pixel TEC image as input. First, we discard all TEC pixels with value below zero or above 150 TECu. Next, we convert TEC to $\log_{10}$TEC. Following Aa et al. (2020), we estimate the background using a sliding window average. Finally, we subtract this background from the $\log_{10}$TEC image. The sliding window size used by Aa et al. (2020) corresponds to about 17° of latitude and we find that window sizes between 15° and 19° work well. Figures 3a and 3b show the TEC and $\log_{10}$TEC distributions in our dataset. Figure 3c shows the distribution of preprocessed image pixel values for two selections of background estimation filter size. This panel demonstrates the effect the background estimation filter size has on the preprocessed image pixel value distribution. We chose to show a very small filter and a very large filter to accentuate the effect of the filter size, but in practice we use medium-sized filters.

The purpose of this stage is to filter out variations in the data that are not helpful for identifying the MIT. The highest amplitude of those variations is typically large-scale TEC structures including seasonal TEC variations and day-night variations. This can be seen in Figure 1 where there is a large difference between the TEC on the day side versus on the night side. In Figure 2a, which shows the output of the preprocessing stage, large-scale TEC structures are successfully removed. If the window size is too large, then variations such as day-night are allowed to pass and if the window size is too small, then large troughs will be filtered out. If our data were higher resolution, then we might have considered also using a low pass filter, effectively creating a bandpass filter. The highpass portion of the filter essentially sets the maximum size for the trough and a lowpass filter would set a minimum size. In our case, we would like to be able to detect troughs that are only one pixel (one degree) wide in latitude and so lowpass filtering is not needed.

One final note on the preprocessing stage is that converting to $\log_{10}$TEC places importance on the relative decrease in TEC in the trough rather than absolute decrease. Although this could exaggerate the occurrence rate of the trough during the winter when TEC values are lower in general, using the logarithmic scale more closely aligns our definition of the MIT with previous work. Additionally, the distribution of $\log_{10}$TEC is more symmetric than the distribution of TEC which is generally beneficial for analysis and machine learning.

2.2.2 Scoring

The reason that thresholding the preprocessed image performs poorly is that any relative low region will be labeled as trough such as isolated noisy pixels, low regions within the auroral oval, or low regions at low latitudes. The purpose of the scoring stage is to produce an image which is more suitable for thresholding, i.e. it more closely approximates our level of confidence in each pixel being part of the MIT. To accomplish this, we model the score image's contribution to the preprocessed image (forward model), then invert the model to find a score image given a preprocessed image. We incorporate additional prior information about the MIT into the inversion problem via regularization. An example of a preprocessed image and the corresponding scored image are shown in Figures 2a and 2b respectively. Note how the preprocessed image has some non-MIT low regions at both higher and lower latitude than we would expect the trough, e.g. near MLT 22 and MLT 6. Thresholding the preprocessed image directly would result in these regions being mislabeled. In the score image these regions are not scored as highly as the pixels within the MIT.

We model each preprocessed image as a linear combination of Gaussian radial basis functions (RBFs) (Bishop, 2006) with the weights given by the score image values. The i-th preprocessed image $\mathbf{x}_i$ is modeled as:

$$\mathbf{x}_i = -A_i\mathbf{u}_i + \epsilon_i \tag{3}$$

where $\mathbf{u}_i$ is the i-th score image and ϵ is noise. Each column of the matrix A_i contains an RBF centered on a pixel of the grid. The minus sign indicates that higher score values will result in lower $\log_{10}$TEC values in the preprocessed image. Nominally, all the A_i's are identical, but because $\mathbf{x}_i$ is always missing data, A_i refers to the full basis matrix with rows dropped corresponding to the elements that are missing in $\mathbf{x}$. If $\mathcal{G}_i$ is the set of indices where $\mathbf{x}_i$ has data, then the size of A_i is $(|\mathcal{G}_i|\times 10800)$. The inverse problem is expressed as:

$$\mathbf{u}_i^* = \arg\min_{\mathbf{u}_i} \mathbf{x}_i^T A_i\mathbf{u}_i + \alpha\|W_i\mathbf{u}_i\|_2^2 + \beta R(\mathbf{u}_i) \tag{4}$$

where $\mathbf{u}_i^*$ is the fitted score image, W_i is a diagonal weighting matrix and R is the total variation (TV) regularizer. α and β are nonnegative coefficients which weight the relative importance of the three components.

The first term maximizes the forward model in the direction of the data. Using RBFs causes the score image elements to affect all the elements in the corresponding neighborhood of the preprocessed image. This property is desirable because the MIT is a large scale structure, so a single low pixel in the preprocessed image does not constitute the trough. Only when a neighborhood of pixels is low should the region be considered part of the MIT.

The second term is weighted *L2* regularization which serves two purposes: the first is to prevent $\mathbf{u}_i$ from going to infinity where $\mathbf{x}_i$ is negative and the second is to prevent $\mathbf{u}_i$ from taking high values far away from the expected location of the trough. The diagonal weighting matrix W_i increases regularization on pixels which are far from the model developed by Deminov and Shubin (2018). The diagonal elements of W_i are given by:

$$w(\lambda, \phi) = c|\lambda - m(\phi)| + 1 \tag{5}$$

where $w(\lambda, \phi)$ is the weight at MLat λ and MLT ϕ, $m(\phi)$ is the latitude of the prior MIT model at MLT ϕ and c is a scalar which sets the maximum weight.

The third term is TV regularization which increases the sparsity of the solution's gradients. R is given by:

$$R(\mathbf{u}_i) = \sum_{j=1}^{M} \Big| \sum_{k\in\mathcal{N}_j} u_{ij} - u_{ik} \Big| \tag{6}$$

The inner sum computes the sum of the differences between the j-th element of $\mathbf{u}_i$ and all of its neighbors and the outer sum computes the sum of all the absolute values of the inner sum. The set of indices of the 4 neighbors of the j-th pixel in the grid is denoted $\mathcal{N}_j$. Minimizing the *L1* norm of the inner sum tends to make it sparse, which means the gradients of the score image will be sparse. This encourages the score image to have larger contiguous regions of the same value. The goal of using this is to create score images which are less influenced by noise and missing values, instead tending to have larger contiguous trough regions.

The inverse problem is designed to capture three assumptions about the MIT as it appears in TEC images. The first is that the trough is a large-scale structure which should contain regions of low TEC. The second is that the MIT should appear in the subauroral / midlatitude regions of the ionosphere, i.e. close to the Deminov and Shubin (2018) empirical model. Lastly, the trough should appear as one large contiguous region of low TEC, not as many smaller regions scattered around. The first assumption is incorporated in the forward model via the RBF basis, and the others are implemented in the inversion using regularizers.

2.2.3 Postprocessing

The purpose of postprocessing is to clean up specific errors which remain after thresholding the score image. These errors include small patches classified as MIT due to noise in the preprocessed image and trough-classified pixels within the auroral oval. Examples of the input and output to the postprocessing step are shown Figures 2c and 2d respectively. Positive pixels within the auroral oval are removed because density depletions within the auroral oval are "high latitude troughs" and not the MIT. For the auroral boundary (shown as the dashed line in Figure 2c), we use the fitted measurements from the Special Sensor Ultraviolet Spectrographic Imager (SSUSI) aboard the Defense Meteorological Satellite Program (DMSP) satellites (Paxton et al., 2003). Finally, the output of this step is the label image. Figure 2d demonstrates that the algorithm can label the two-dimensional structure of the trough at a given time. Even when the trough is not continuous in longitude, the algorithm identifies the trough where it appears in the data.

2.3 Verification

The MIT has no ground truth labels, therefore verifying an algorithm for identifying it is a major challenge. In other detection or classification settings we would compare our algorithm's output with a ground truth dataset, but no ground truth exists for the MIT, making this impossible. For the sake of simplicity most researchers set up their task as a binary detection problem, but in the reality, the MIT does not appear discretely. For this reason, any discretization of the MIT involves an arbitrary choice of a threshold.

Given this issue, it is insufficient to show that our labels match those from a previous study. Instead, we show that our algorithm has two properties. The first property is that within a reasonable range, our choice of parameters does not negatively affect the resulting labeled dataset. In cases where the parameter choice does affect the output, it should happen in an expected and reasonable way, e.g. making a threshold more selective should decrease the measured MIT probability. This is important because slight perturbations our parameter values should not change the measured statistics of the trough. The second desired property is that there exists a selection of parameters which can have our labels match those from previous studies with good accuracy. This property is important because it means that our method is identifying the same types of troughs as other methods, only for two-dimensional data.

There are two ways we compare with another algorithm. The first is "instance comparison", in which we match up our detections with those from another algorithm, instance by instance. Because this involves comparing a one-dimensional data product with a two-dimensional one, this requires some processing, which we describe below. The other way we can compare is at the dataset level, where we compute the same statistics on both datasets and see how well they agree.

2.3.1 Instance Comparison Methodology

To perform the "instance comparison", we randomly select $N = 200$ days from the 7 year span of SWARM data to run it on. The Aa et al. (2020) algorithm splits the Swarm orbits into segments between 45° and 75° MLat, through which each satellite passes twice in the northern hemisphere per orbit. The segments in which the satellites' latitude is increasing, we call "ascending" segments and the other segments we call "descending" segments. Each of our label images corresponds to one hour, and we choose the ascending segment and the descending segment which are closest in time. This time alignment is illustrated in Figure 4a. The dashed black line indicates a Swarm satellite's MLat over time and the vertical blue lines indicate the time limits of a series of TEC maps. The red and green lines show the ascending and descending orbital segments respectively. The arrows show which Swarm orbital segments are associated with each TEC map. Note that the TEC map marked "TEC_{t+1}" is being compared to Swarm measurements from two different orbits. A spatial view of the comparisons is shown in Figure 4b. The inner circle and middle circles represent orbital segment latitude boundaries as defined by Aa et al. (2020) (75° and 45° respectively) and the satellite's orbital segments are colored the same as in Figure 4a. The blue shape is the trough estimated by our algorithm.

We then extract a 3-pixel-wide path of our label images under the SWARM satellite orbital segments and mark the highest and lowest MLats where our labels are positive for the poleward and equatorward walls of the MIT respectively. We compare these latitudes to ones from the Aa et al. (2020) algorithm. In the case of multiple trough candidates within a single orbital segment, the Aa et al. (2020) algorithm chooses the one with the lowest MLat. For our comparison, we instead choose the one which best agrees with our labels. This results in better ground truth labels because rather than somewhat arbitrarily choosing the lower MLat troughs, we choose troughs that appear in two separate datasets. To allow for better agreement between the two sources of data, we raise the threshold of the Aa et al. (2020) algorithm from -0.3 to -0.2. A threshold of -0.2 defines troughs as a 36% decrease in electron density from the background value. This is still a reasonable choice as other authors have chosen their thresholds even higher (Ishida et al., 2014).

Aligning the Aa et al. (2020) labels with ours nominally results in 6 comparisons per TEC map (3 satellites, 2 orbital segments, but often less due to missing Swarm data) and 6 variables per comparison. There are two binary variables, one for SWARM and one for TEC, indicating whether any trough is detected in the comparison. There are also four continuous variables indicating the latitudes of the poleward and equatorward walls for each of the two data sources. If no trough is detected in one of the data sources, then there are no values for these latitudes. In a true positive case, when a trough is detected by both our algorithm and the Aa et al. (2020) algorithm, we compute the errors of the wall latitudes as:

$$E_P = \lambda_{TP} - \lambda_{SP} \tag{7}$$

$$E_E = \lambda_{TE} - \lambda_{SE} \tag{8}$$

Where E_P is the poleward wall error, E_E is the equatorward wall error, λ_{SP} is the poleward wall latitude from SWARM, λ_{TP} is the poleward wall latitude from TEC, etc. The two continuous errors, E_P and E_E, are shown in Figure 4b and 4c. With 6 comparisons per TEC map and 24 TEC maps per day, we get a total of $144N = 48,800$ comparisons, from which we compute accuracy, rates for the binary error types, and statistics (mean / standard deviation) for the continuous errors.

4.png

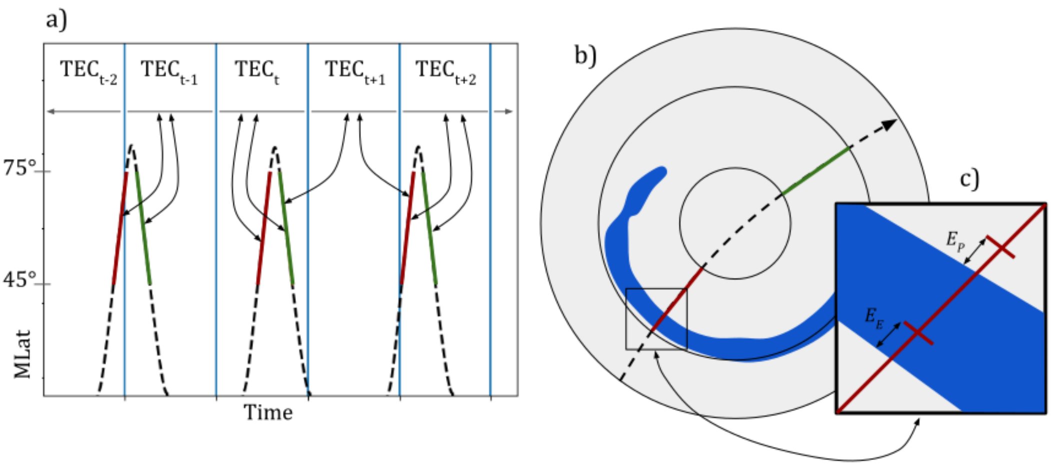

Figure 4. a) Time alignment of satellite orbital segments with TEC maps. Blue vertical lines show TEC map intervals. Black dashed lines show satellite MLat. Red and green lines mark ascending and descending segments respectively. Arrows indicate which segments are compared with each TEC map. b) Spatial view of SWARM - TEC comparison. Dashed, red and green lines are the same as in (a). Blue region is the trough estimated by our algorithm. c) Poleward and Equatorward wall continuous errors.

3 Results

To demonstrate that our algorithm performs well over a broad range of parameter values, we focus on the effects of the *L2* regularization weight parameter (α in equation 4). In our experiments we find that α has the greatest impact on the score images and resulting labeled dataset. A higher setting of α forces the score image values to be lower, especially for pixels far from the Deminov and Shubin (2018) empirical model. On an instance level, the TV regularization weight parameter (α in equation 4) might have a somewhat different effect, but at the statistical level we find that changing β has an effect similar to changing α but less pronounced. Ultimately, several of the parameters have the effect of broadly increasing or decreasing the score image pixel values including α, β and the background estimation filter size. Generally, increasing the score pixel values might not change the overall binary agreement with the Aa et al. (2020) detections, but would rather increase both true positive and false positive rates. Among the parameters that have this effect, we find that α is the most significant and so this is the parameter we focus on in this study.

3.1 Instance Comparison

Figure 5 shows the results of the instance-level comparison. In each panel the different colored lines represent different settings of α. Figure 5a shows how the binary agreement between our detections and those of Aa et al. (2020) varies over a range of thresholds. Binary agreement is the percentage of comparisons which agree on MIT presence. At lower settings of α, a higher threshold is required to achieve the same agreement level. This is expected because a lower α setting should increase all score image pixel values. For these curves, α appears to scale the x-axis. For this reason, lower settings of α make the output less sensitive to the choice of threshold. One detail to notice is that all of the curves achieve similar agreement with Aa et al. (2020) at their maximum. This is important because it means our level of agreement with Aa et al. (2020) is not sensitive

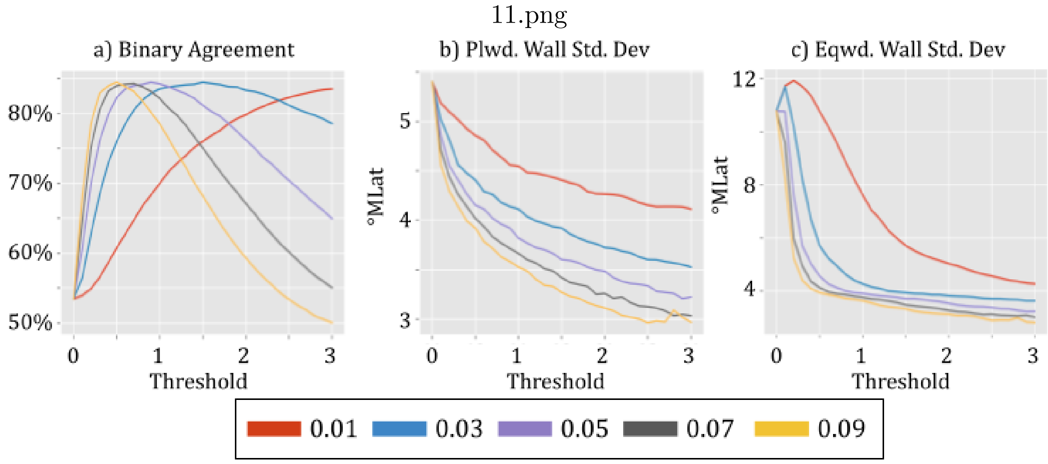

Figure 5. Results of the instance-level comparison between this work and Aa et al. (2020) for different settings of α and threshold. Colored lines indicate setting of α. a) Binary agreement: percentage of instances where the labels agree on presence of MIT. b) Standard deviation between estimates of the latitude of the poleward wall. c) Standard deviation between estimates of the latitude of the equatorward wall.

to our setting of α. This figure demonstrates both of the previously mentioned desirable properties of our algorithm because the maximum agreement does not depend on α and because we can achieve a high level of agreement by tuning the threshold.

Figures 5b and 5c describe how the continuous error distributions (equations 7 and 8) are affected by the choice of α and threshold. In both plots, the y-axis shows the error standard deviation measured in degrees latitude. Figure 5b shows the error standard deviation for the poleward wall latitude and 5c shows the error standard deviation for the equatorward wall latitude. In both plots, for all settings of α, increasing the threshold generally decreases the error standard deviation. This effect is due to the fact that the error standard deviation can only be computed for comparisons where both algorithms detect the MIT. Increasing the threshold causes our algorithm to reject shallow troughs which are less likely to appear in both datasets. Only deep troughs will remain which are more likely to be detectable in both datasets. This reduces the chance of the Aa et al. (2020) algorithm and ours selecting different troughs, which reduces the error standard deviation. Another detail to note is that equatorward wall error standard deviations at low threshold values are much higher than the poleward wall values. Our post-processing operations most likely prevent the poleward wall estimate from being too far off, but our equatorward wall estimate is less constrained. In general, the equatorward wall of the MIT is less well-defined than the poleward wall due to shallower gradients, and other researchers have found it difficult to identify with confidence (Prölss, 2007). Figures 5b and 5c also help verify our results because for each setting of α, near their peak agreement, they achieve a similar error standard deviation value.

3.2 Statistical Comparison

In this section, we compare statistical relationships exhibited by the MIT computed over our dataset and the Aa et al. (2020) dataset. In Figure 6 we show variations of the MIT with level of magnetic activity (as quantified by Kp index). The left column panels (a, c, e, g, i) show the MIT occurrence rate vs Kp and the right column panels (b, d, f, h, j) show the latitude of the MIT minimum (within 5 hours MLT of midnight) vs

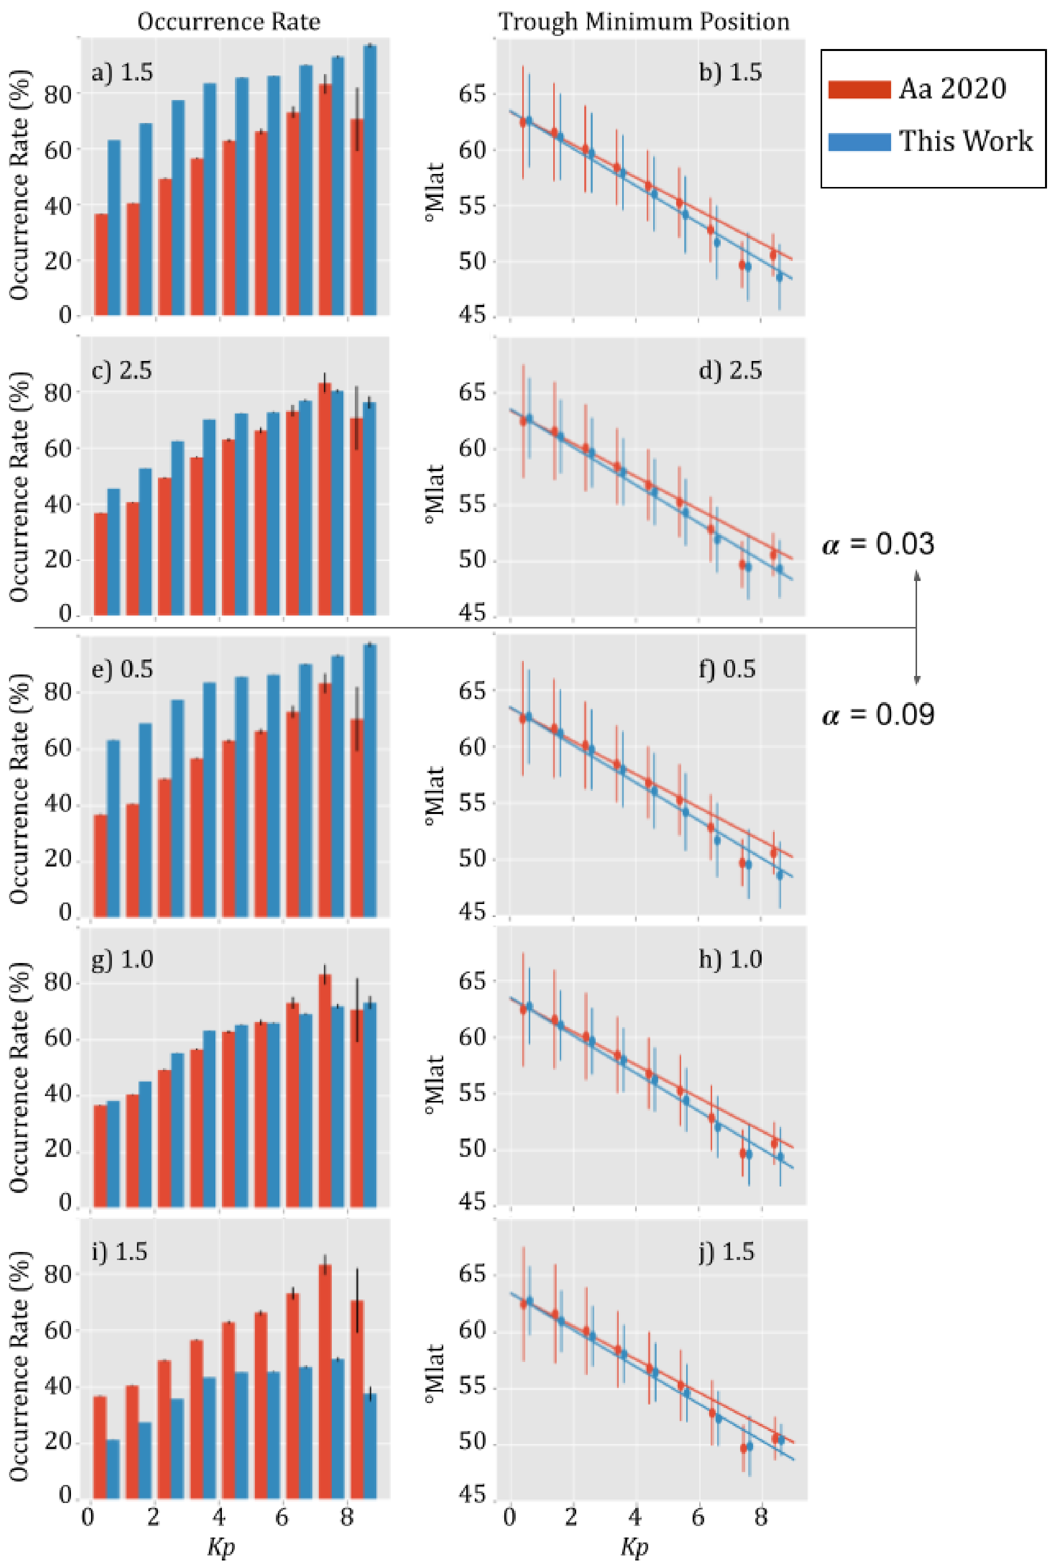

13.png

Figure 6. Results of statistical level comparison between this work and Aa et al. (2020) showing MIT variation with level of geomagnetic activity. Left column: occurrence rate. Right column: latitude of MIT minimum. Each row is computed for a different combination of α and threshold.

Kp. In all panels of Figure 6, our results are shown in blue and the results of Aa et al. (2020) are shown in red. Each row shows the results for a different combination of α and threshold. The threshold for each row is printed next to the panel label. The top two rows (a, b, c, d) are computed with α set to 0.03 and the bottom three rows (e, f, g, h, i, j) are computed with α set to 0.09. To compute the occurrence rates in Figure 6, we filter out latitudinal profiles which have less than 50% of their data or that are outside the MLT range $[-5, 5]$. Then we divide the number of profiles which have any positive labels by the total number of profiles. For the SWARM computation, we eliminate segments outside of the $[-5, 5]$ MLT range and those in the southern hemisphere. To compute the latitude of the MIT minimum, we use our labels to mask out non-trough pixels of the TEC images images, then we search for the latitude which achieves the minimum value at each MLT. Both of these computations are binned by Kp.

The left column panels of Figure 6 show the occurrence rate of the MIT at different levels of magnetic activity. The black error bars are the same as used by Aa et al. (2020):

$$\sigma = \sqrt{\frac{f \times (1 - f)}{N - 1}} \tag{9}$$

where σ is the uncertainty, f is the occurrence rate, and N is the total number of examples for the given bin. As the threshold increases, the occurrence rate decreases at all levels of Kp uniformly. This behavior is expected and is an example of the first desirable property. With an appropriate choice of threshold, for both settings of α, we can make our occurrence rates match those of Aa et al. (2020), which demonstrates the second desirable property.

The right column panels of Figure 6 show the latitude of the MIT minimum at different levels of magnetic activity. It is clear that settings of α and the threshold have essentially no effect whatsoever on this statistical relationship. This fact helps verify our method because we should not expect this statistic to be affected by our selectivity of troughs. Both properties are demonstrated in this figure because the result is insensitive to α and threshold and because it matches Aa et al. (2020) closely.

One of the most interesting contributions from Aa et al. (2020) was their detailed maps of seasonal MIT occurrence rate. Since they utilized data from Swarm satellites which, over the course of the dataset, cover all latitudes and local times, they were able to improve the field of view and detail over the earlier maps of Ishida et al. (2014). These occurrence rates are restricted to $Kp \leq 3$. Copied in Figure 7a is their winter map. We compute the same map by counting the number of times the MIT is observed and dividing by the number of times we have TEC data in each grid cell. We use the same winter group as Aa et al. (2020) which is November–February. The results for various settings of α and threshold are shown in Figures 7b - 7f. The two numbers labeling each panel of Figure 7 are the settings of α and threshold respectively.

The first thing to notice about Figure 7 is that each of panels b-f match the general distribution computed by Aa et al. (2020), i.e. two modes, one around 3-4 MLT and one around 16-17 MLT. The most obvious difference between the panels is that the occurrence rate decreases as the threshold increases. This observation is additional evidence that statistical results computed with our dataset are not very sensitive to our choice of parameters, and that the effects our parameter choice does have is expected. The decrease in occurrence rate is not perfectly uniform though, and at the highest threshold, there is a much higher occurrence rate in the two modes (4 and 17 MLT) than at other locations. This indicates that more deep troughs occur in our dataset at the two modes than at other locations. Figure 7f appears to be the closest match to the plot from Aa et al. (2020) shown in Figure 7a. Finally, Figure 7 demonstrates that with our labels, we can compute distributions with increased spatial resolution.

12.png

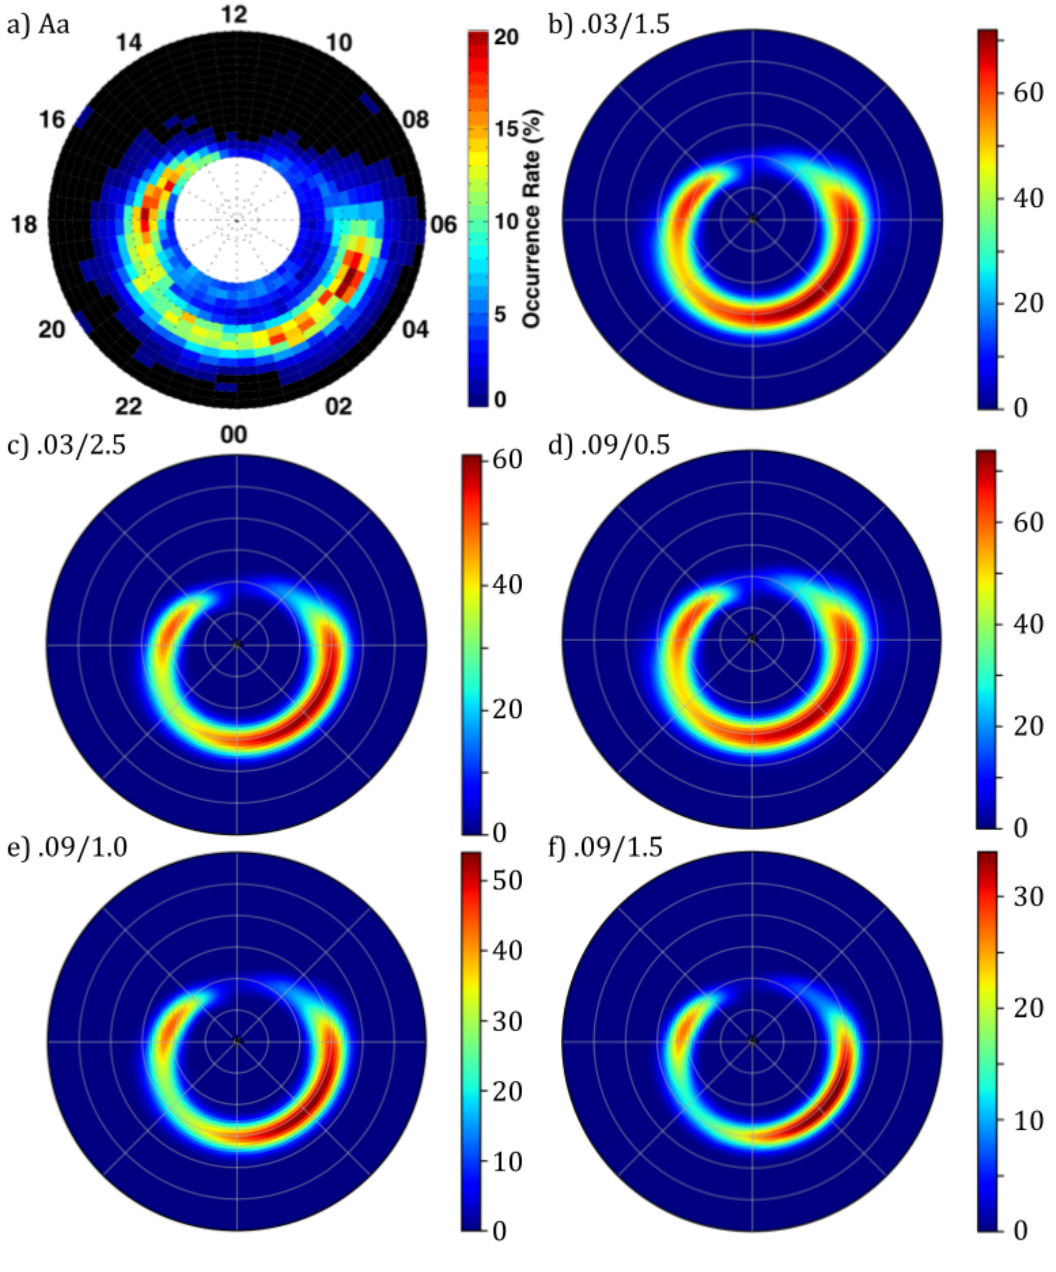

Figure 7. Results of statistical level comparison between this work and Aa et al. (2020) showing MIT spatial distribution during winter (November–February). Each panel is computed for a different combination of α and threshold. Panel labels are α/threshold.

4 Conclusion

In order to perform a large-scale statistical study of the main ionospheric trough, we develop a method for algorithmically identifying the MIT in TEC images. After preparing the TEC data, our method involves 4 steps: preprocessing, scoring, thresholding and postprocessing. Each of these steps has has several parameters to control their output. Because the MIT has no true definition and we have no ground truth to compare to, verifying our method is very difficult. Comparing our method to a previous one would only be comparing one arbitrary set of decisions with another. To better verify our algorithm, we demonstrate that it has two properties. The first is that the output has no variation or an expected variation to our parameter choices. The second is that we have the ability to calibrate our algorithm to match the results of previous studies. To show these properties in our algorithm, we compare our labels with the labels from Aa et al. (2020) at both the instance level and at the statistical level.

In subsequent work, we are utilizing this dataset to investigate the formation mechanisms and the higher dimensional statistics of the MIT. The most exciting investigation this two-dimensional dataset enables is a statistical study utilizing MIT labels and measurements of ionospheric convection. Additionally, it is well known that plasma irregularities the MIT can have negative effects on satellite communications, but no previous study has directly quantified the frequency and severity of such interruptions. One exciting application of our dataset would be to combine it with the scintillation event dataset developed by Mrak et al. (2020) and characterize the relationship between these two phenomena.

Acknowledgments

The work of YN was supported by NASA grant 80NSSC18K0657, 80NSSC20K0604, 80NSSC20K0725, 80NSSC21K1321 and 80NSSC20K1788, NSF grant AGS-1907698 and AGS-2100975, and AFOSR grant FA9559-16-1-0364. YN thanks ISSI/ISSI-BJ for the “Multi-Scale Magnetosphere-Ionosphere-Thermosphere Interaction” workshop. TEC and Swarm data were obtained through `http://cedar.openmadrigal.org/` and `https://earth.esa.int/`. We acknowledge Goddard Space Flight Center / Space Physics Data Facility and OMNIWeb for the OMNI2 dataset (`https://spdf.gsfc.nasa.gov/`).

References

Aa, E., Zou, S., Erickson, P. J., Zhang, S., & Liu, S. (2020). Statistical analysis of the main ionospheric trough using swarm in situ measurements. *Journal of geophysical research. Space physics*, *125*(3). doi: 10.1029/2019JA027583

Bishop, C. M. (2006). *Pattern recognition and machine learning.* New York: Springer.

Collis, P., & Häggström, I. (1988). Plasma convection and auroral precipitation processes associated with the main ionospheric trough at high latitudes. *Journal of atmospheric and terrestrial physics*, *50*(4), 389–404.

Deminov, M. G., & Shubin, V. N. (2018). Empirical model of the location of the main ionospheric trough. *Geomagnetism and aeronomy*, *58*(3), 348–355.

Emmert, J. T., Richmond, A. D., & Drob, D. P. (2010). A computationally compact representation of magnetic-apex and quasi-dipole coordinates with smooth base vectors. *Journal of Geophysical Research: Space Physics*, *115*(A8), `https://doi.org/10.1029/2010JA015326`.

Ishida, T., Ogawa, Y., Kadokura, A., Hiraki, Y., & Häggström, I. (2014). Seasonal variation and solar activity dependence of the quiet-time ionospheric trough. *Journal of Geophysical Research: Space Physics*, *119*(8), 6774–6783.

Karpachev, A. T. (2019). Variations in the winter troughs' position with local time, longitude, and solar activity in the northern and southern hemispheres. *Jour-*

nal of geophysical research. Space physics*, *124*(10), 8039–8055.

Kintner, P. M., Ledvina, B. M., & de Paula, E. R. (2007). Gps and ionospheric scintillations. *Space Weather*, *5*(9). Retrieved from `https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2006SW000260` doi: https://doi.org/10.1029/2006SW000260

Le, H., Yang, N., Liu, L., Chen, Y., & Zhang, H. (2017). The latitudinal structure of nighttime ionospheric tec and its empirical orthogonal functions model over north american sector. *Journal of Geophysical Research: Space Physics*, *122*(1), 963-977. Retrieved from `https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/2016JA023361` doi: 10.1002/2016JA023361

Mishin, E. V., Burke, W. J., Huang, C. Y., & Rich, F. J. (2003). Electromagnetic wave structures within subauroral polarization streams. *Journal of Geophysical Research: Space Physics*, *108*(A8). Retrieved from `https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2002JA009793` doi: https://doi.org/10.1029/2002JA009793

Mrak, S., Semeter, J., Nishimura, Y., Rodrigues, F. S., Coster, A. J., & Groves, K. (2020). Leveraging geodetic gps receivers for ionospheric scintillation science. *Radio science*, *55*(11).

Nilsson, H., Sergienko, T. I., Ebihara, Y., & Yamauchi, M. (2005). Quiet-time mid-latitude trough: influence of convection, field-aligned currents and proton precipitation. *Annales geophysicae (1988)*, *23*(10), 3277–3288.

Paxton, L. J., Morrison, D., Strickland, D. J., McHarg, M., Zhang, Y., Wolven, B., … Meng, C.-I. (2003). The use of far ultraviolet remote sensing to monitor space weather. *Advances in Space Research*, *31*(4), 813-818. Retrieved from `https://www.sciencedirect.com/science/article/pii/S0273117702008864` doi: https://doi.org/10.1016/S0273-1177(02)00886-4

Pedatella, N. M., & Larson, K. M. (2010). Routine determination of the plasmapause based on COSMIC GPS total electron content observations of the midlatitude trough. *Journal of Geophysical Research: Space Physics*, *115*(A9). doi: 10.1029/2010JA015265

Pryse, S. E., Kersley, L., Malan, D., & Bishop, G. J. (2006). Parameterization of the main ionospheric trough in the european sector. *Radio science*, *41*(5), RS5S14.

Prölss, G. (2007). The equatorward wall of the subauroral trough in the afternoon/evening sector. *Annales Geophysicae*, *25*(3), 645–659.

Quegan, S., Rodger, A. S., Williams, P. J. S., Lockwood, M., & Rishbeth, H. (1989). The influence of convection on the structure of the high-latitude ionosphere [and discussion]. *Philosophical transactions of the Royal Society of London. Series A: Mathematical and physical sciences*, *328*(1598), 119–137.

Richmond, A. D. (1995). Ionospheric electrodynamics using magnetic apex coordinates. *Journal of geomagnetism and geoelectricity*, *47*(2), 191-212. doi: 10.5636/jgg.47.191

Rideout, W., & Coster, A. (2006). Automated gps processing for global total electron content data. *GPS solutions*, *10*(3), 219–228.

Rodger, A. (2008). The mid-latitude trough—revisited. In *Midlatitude ionospheric dynamics and disturbances* (p. 25-33). American Geophysical Union (AGU). Retrieved from `https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/181GM04` doi: 10.1029/181GM04

Rodger, A., Moffett, R., & Quegan, S. (1992). The role of ion drift in the formation of ionisation troughs in the mid- and high-latitude ionosphere—a review. *Journal of Atmospheric and Terrestrial Physics*, *54*(1), 1 - 30. Retrieved from `http://www.sciencedirect.com/science/article/pii/002191699290082V` doi: https://doi.org/10.1016/0021-9169(92)90082-V

Spiro, R. W., Heelis, R. A., & Hanson, W. B. (1978). Ion convection and the formation of the mid-latitude f region ionization trough. *Journal of Geophysical Re-*

search: Space Physics, *83*(A9), 4255–4264.

Vierinen, J., Coster, A. J., Rideout, W. C., Erickson, P. J., & Norberg, J. (2016). Statistical framework for estimating gnss bias. *Atmospheric measurement techniques*, *9*(3), 1303–1312.

Voiculescu, M., Virtanen, I., & Nygrén, T. (2006). The f-region trough: seasonal morphology and relation to interplanetary magnetic field. *Annales geophysicae (1988)*, *24*(1), 173–185. Retrieved from https://doaj.org/article/949d1977dac2439f8d5d0d8aba354d7c

Yang, N., Le, H., & Liu, L. (2015). Statistical analysis of ionospheric mid-latitude trough over the northern hemisphere derived from gps total electron content data. *Earth, Planets and Space*, *67*(196). Retrieved from https://doi.org/10.1186/s40623-015-0365-1

Yizengaw, E., & Moldwin, M. B. (2005). The altitude extension of the mid-latitude trough and its correlation with plasmapause position. *Geophysical Research Letters*, *32*(9).

Zou, S., Moldwin, M. B., Coster, A., Lyons, L. R., & Nicolls, M. J. (2011). Gps tec observations of dynamics of the mid-latitude trough during substorms. *Geophysical Research Letters*, *38*(14). Retrieved from https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2011GL048178 doi: 10.1029/2011GL048178

Figure 1.

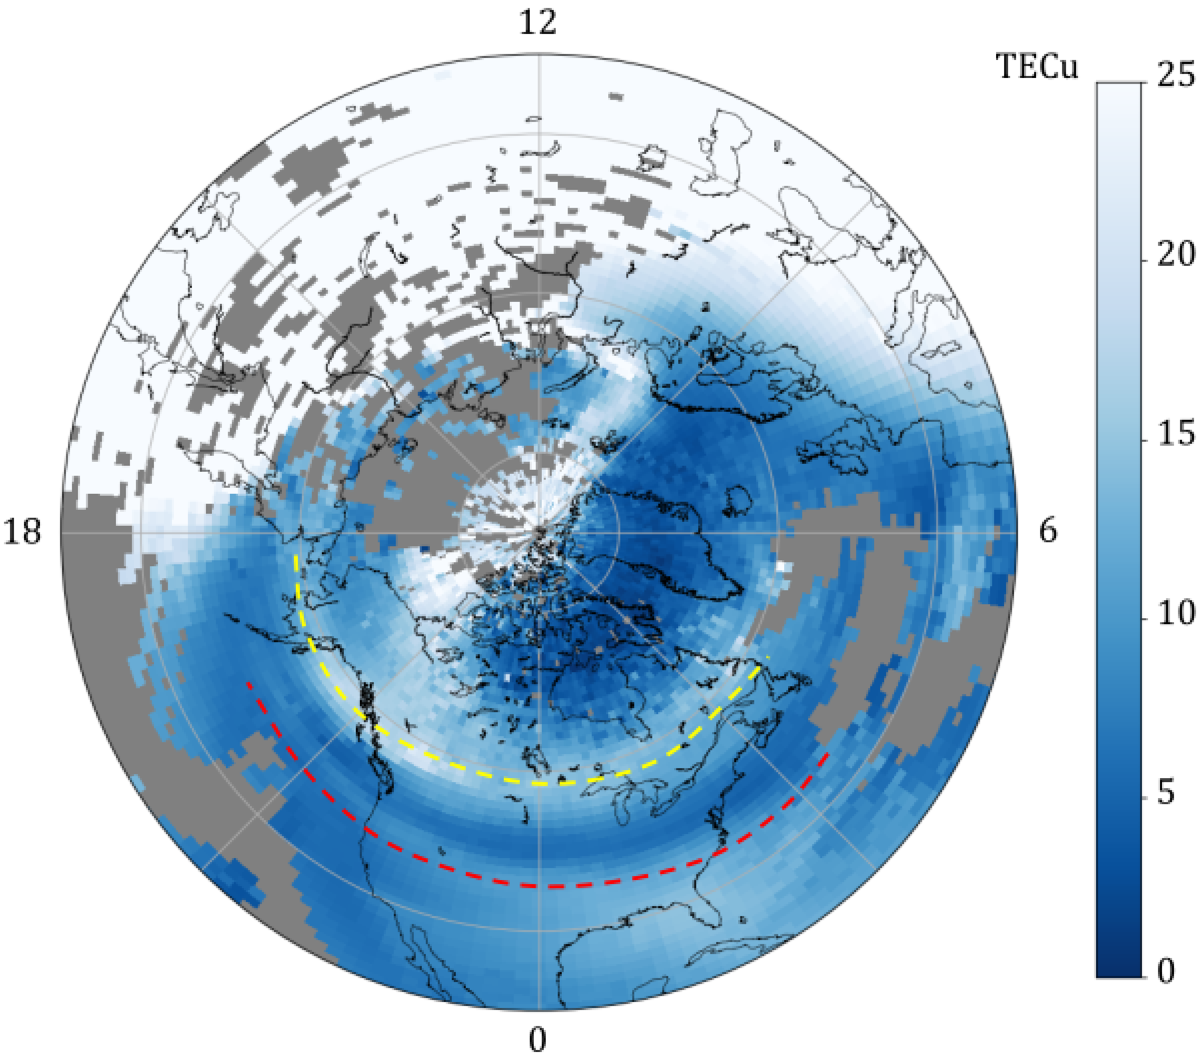

12
18
6
0
TECu
25
20
15
10
5
0

Figure 2.

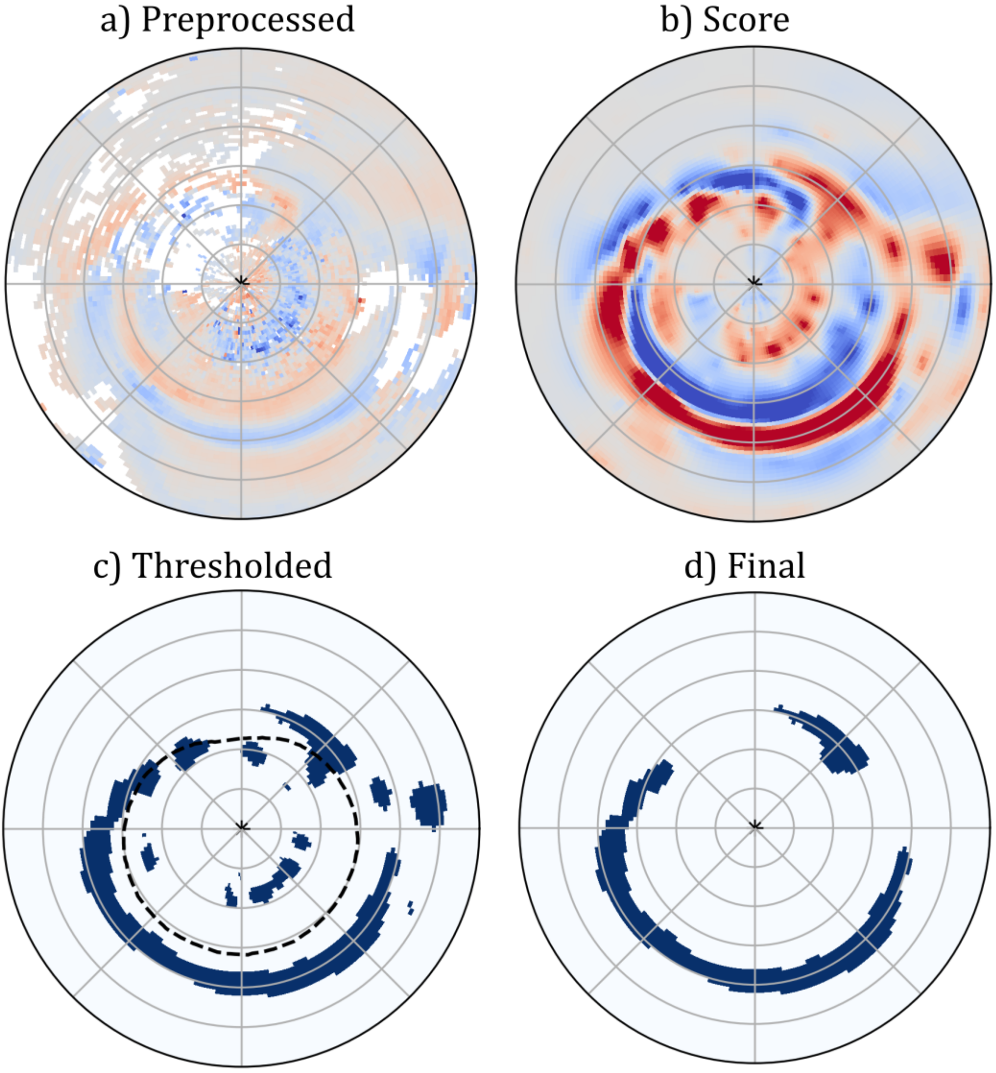
a) Preprocessed
b) Score
c) Thresholded
d) Final

Figure 3.

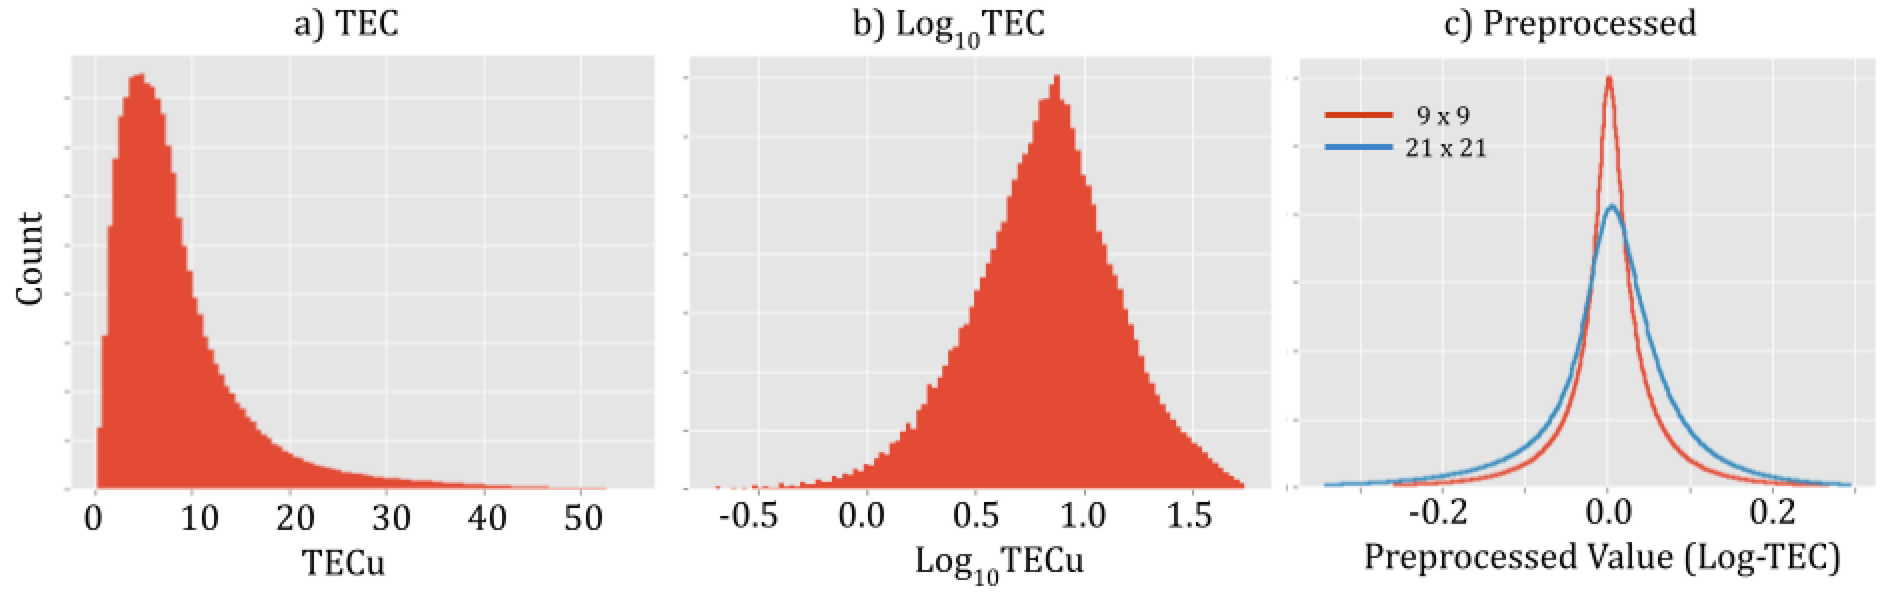

a) TEC
b) Log10TEC
c) Preprocessed
Count
0
10
20
30
40
50
TECu
-0.5
0.0
0.5
1.0
1.5
Log10TECu
9 x 9
21 x 21
-0.2
0.0
0.2
Preprocessed Value (Log-TEC)

Figure 4.

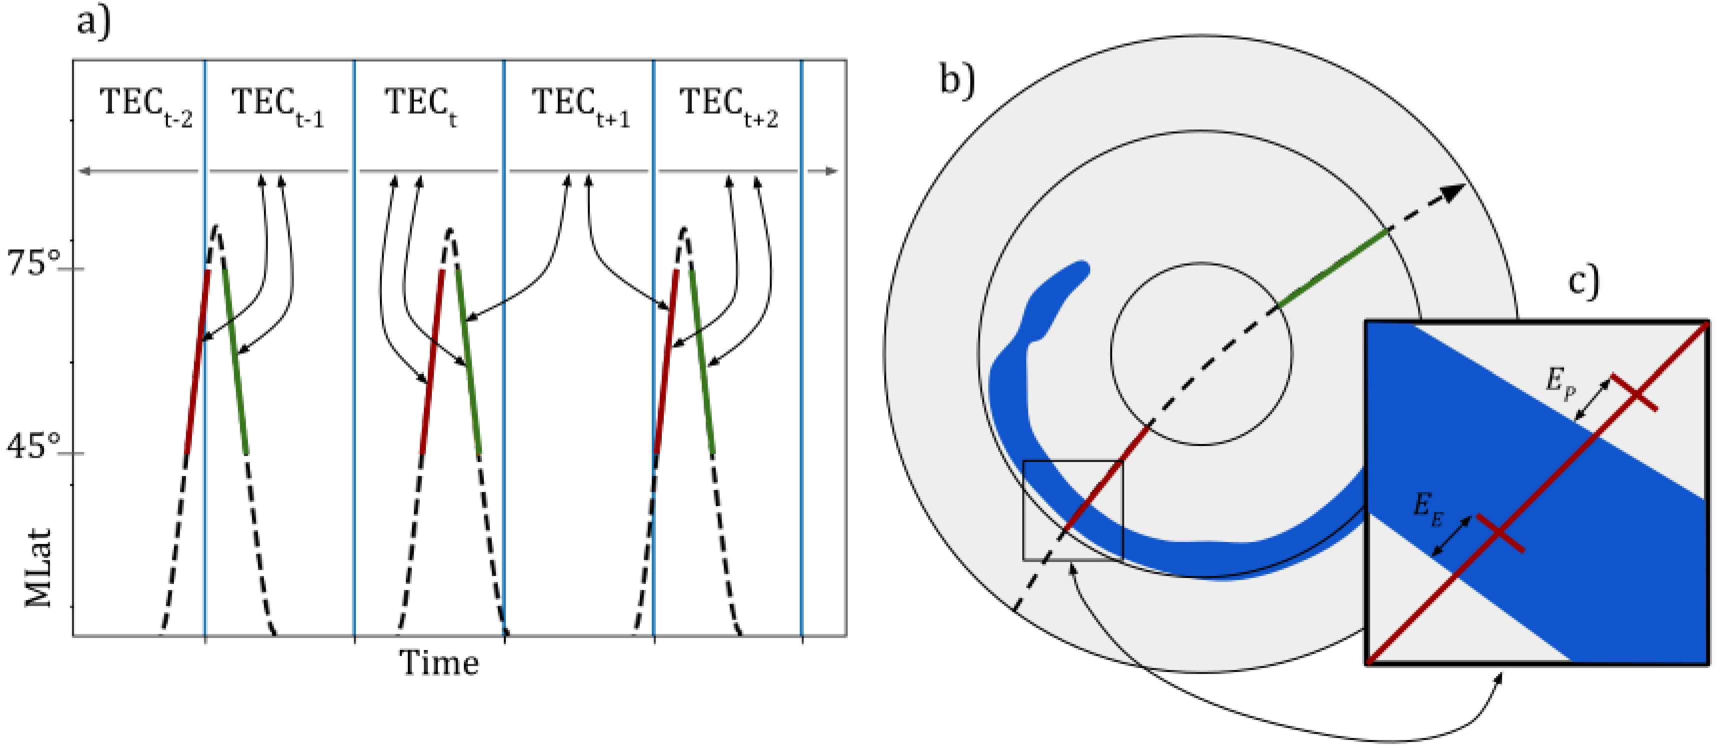

a)
TEC$_{t-2}$
TEC$_{t-1}$
TEC$_t$
TEC$_{t+1}$
TEC$_{t+2}$
75°
45°
MLat
Time
b)
c)
E_P
E_E

Figure 5.

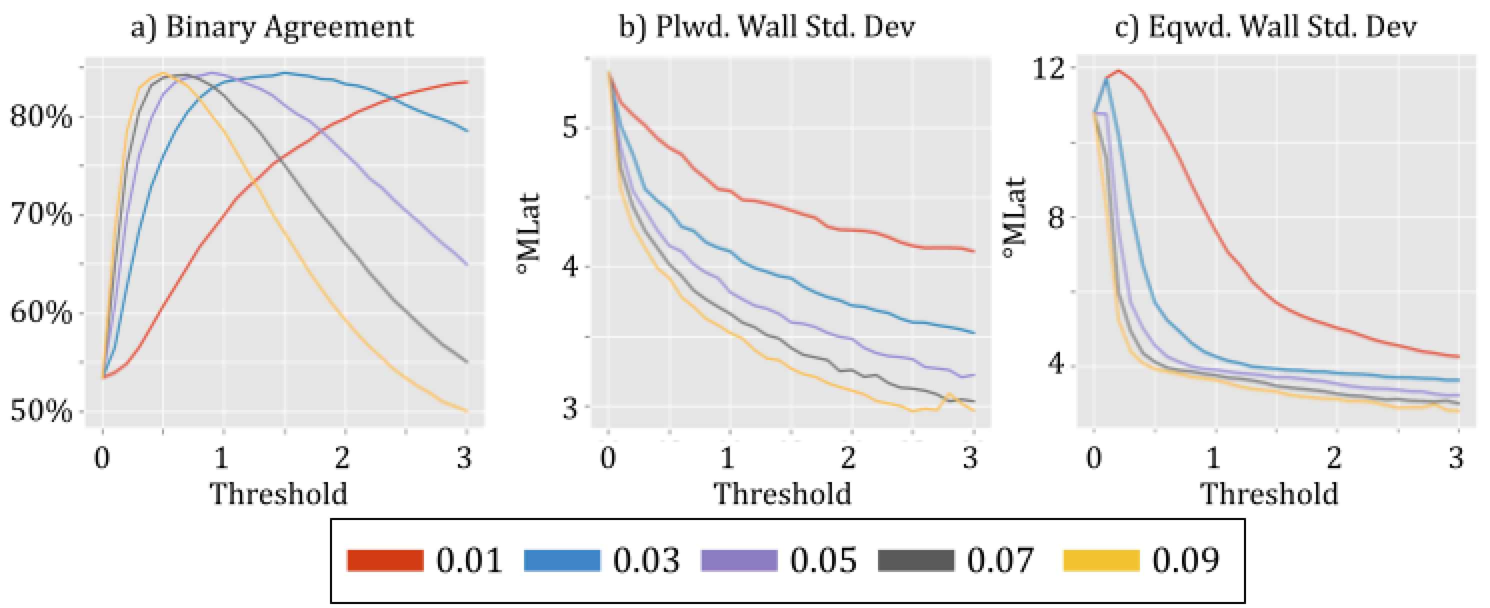
a) Binary Agreement
80%
70%
60%
50%
0
1
2
3
Threshold
b) Plwd. Wall Std. Dev
°MLat
5
4
3
0
1
2
3
Threshold
c) Eqwd. Wall Std. Dev
°MLat
12
8
4
0
1
2
3
Threshold
0.01
0.03
0.05
0.07
0.09

Figure 7.

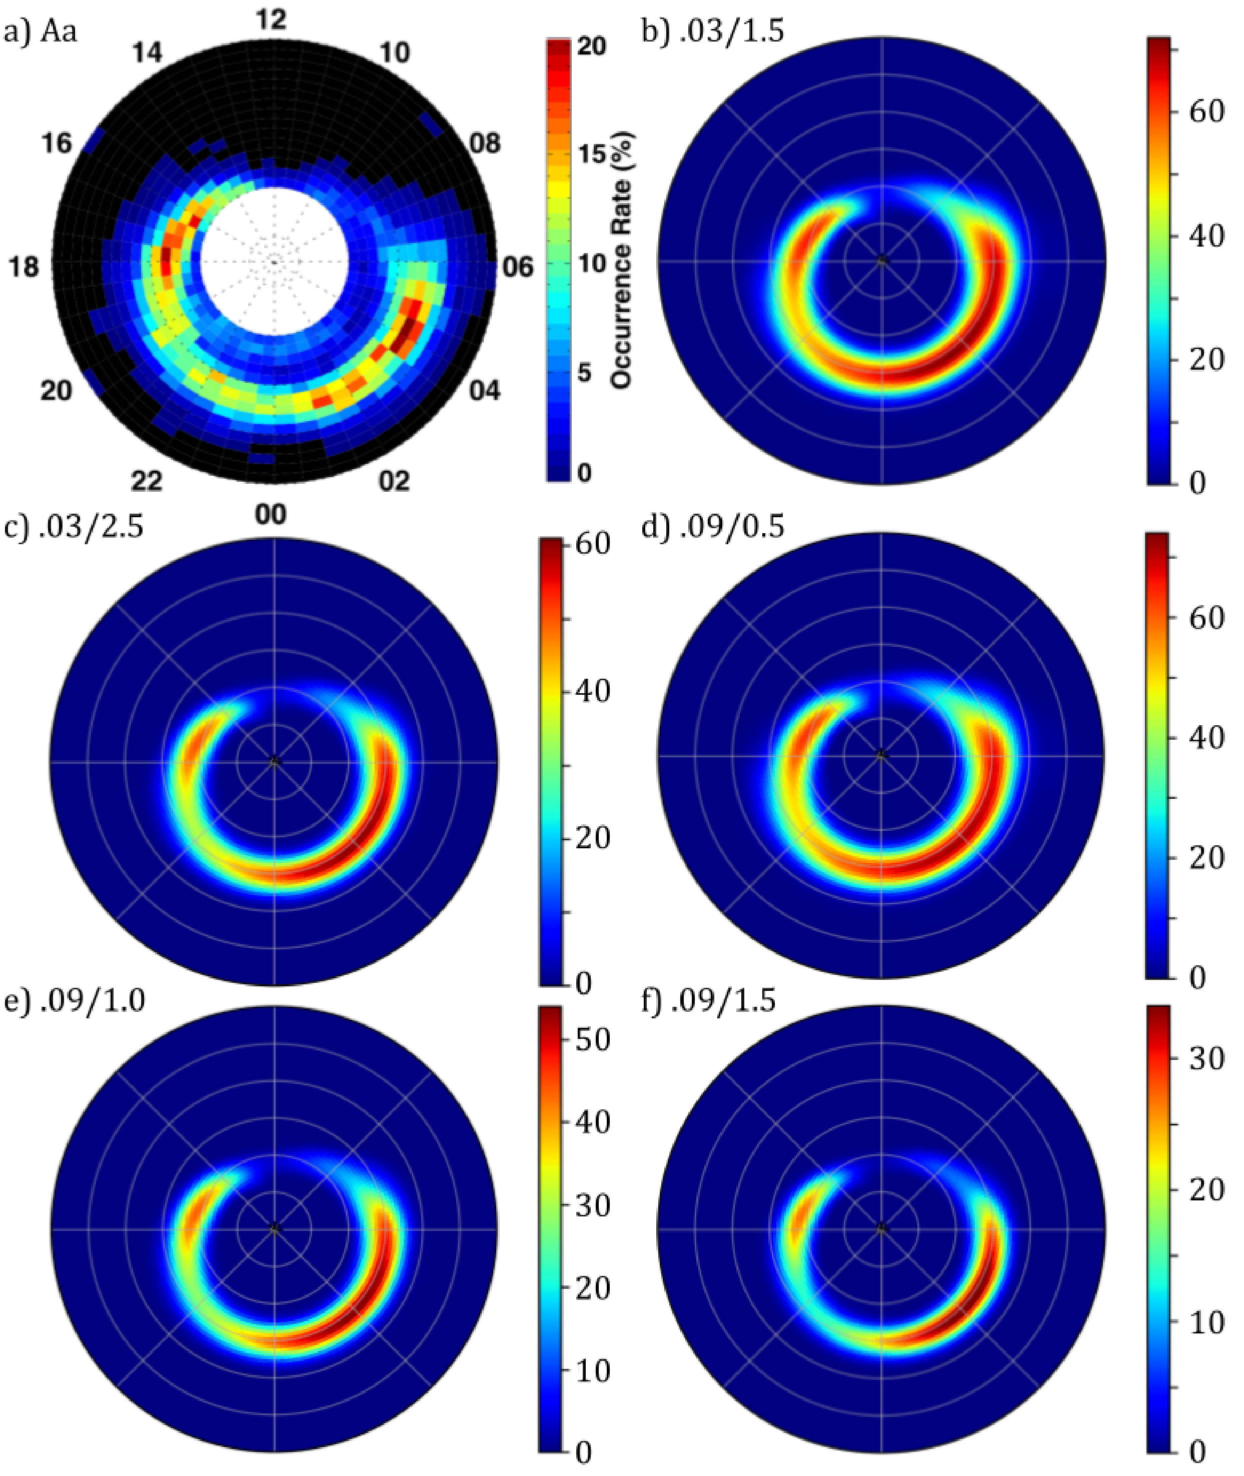

a) Aa
12
14
10
16
08
18
06
20
04
22
02
00
Occurrence Rate (%)
0
5
10
15
20
b) .03/1.5
0
20
40
60
c) .03/2.5
0
20
40
60
d) .09/0.5
0
20
40
60
e) .09/1.0
0
10
20
30
40
50
f) .09/1.5
0
10
20
30

Figure 6.

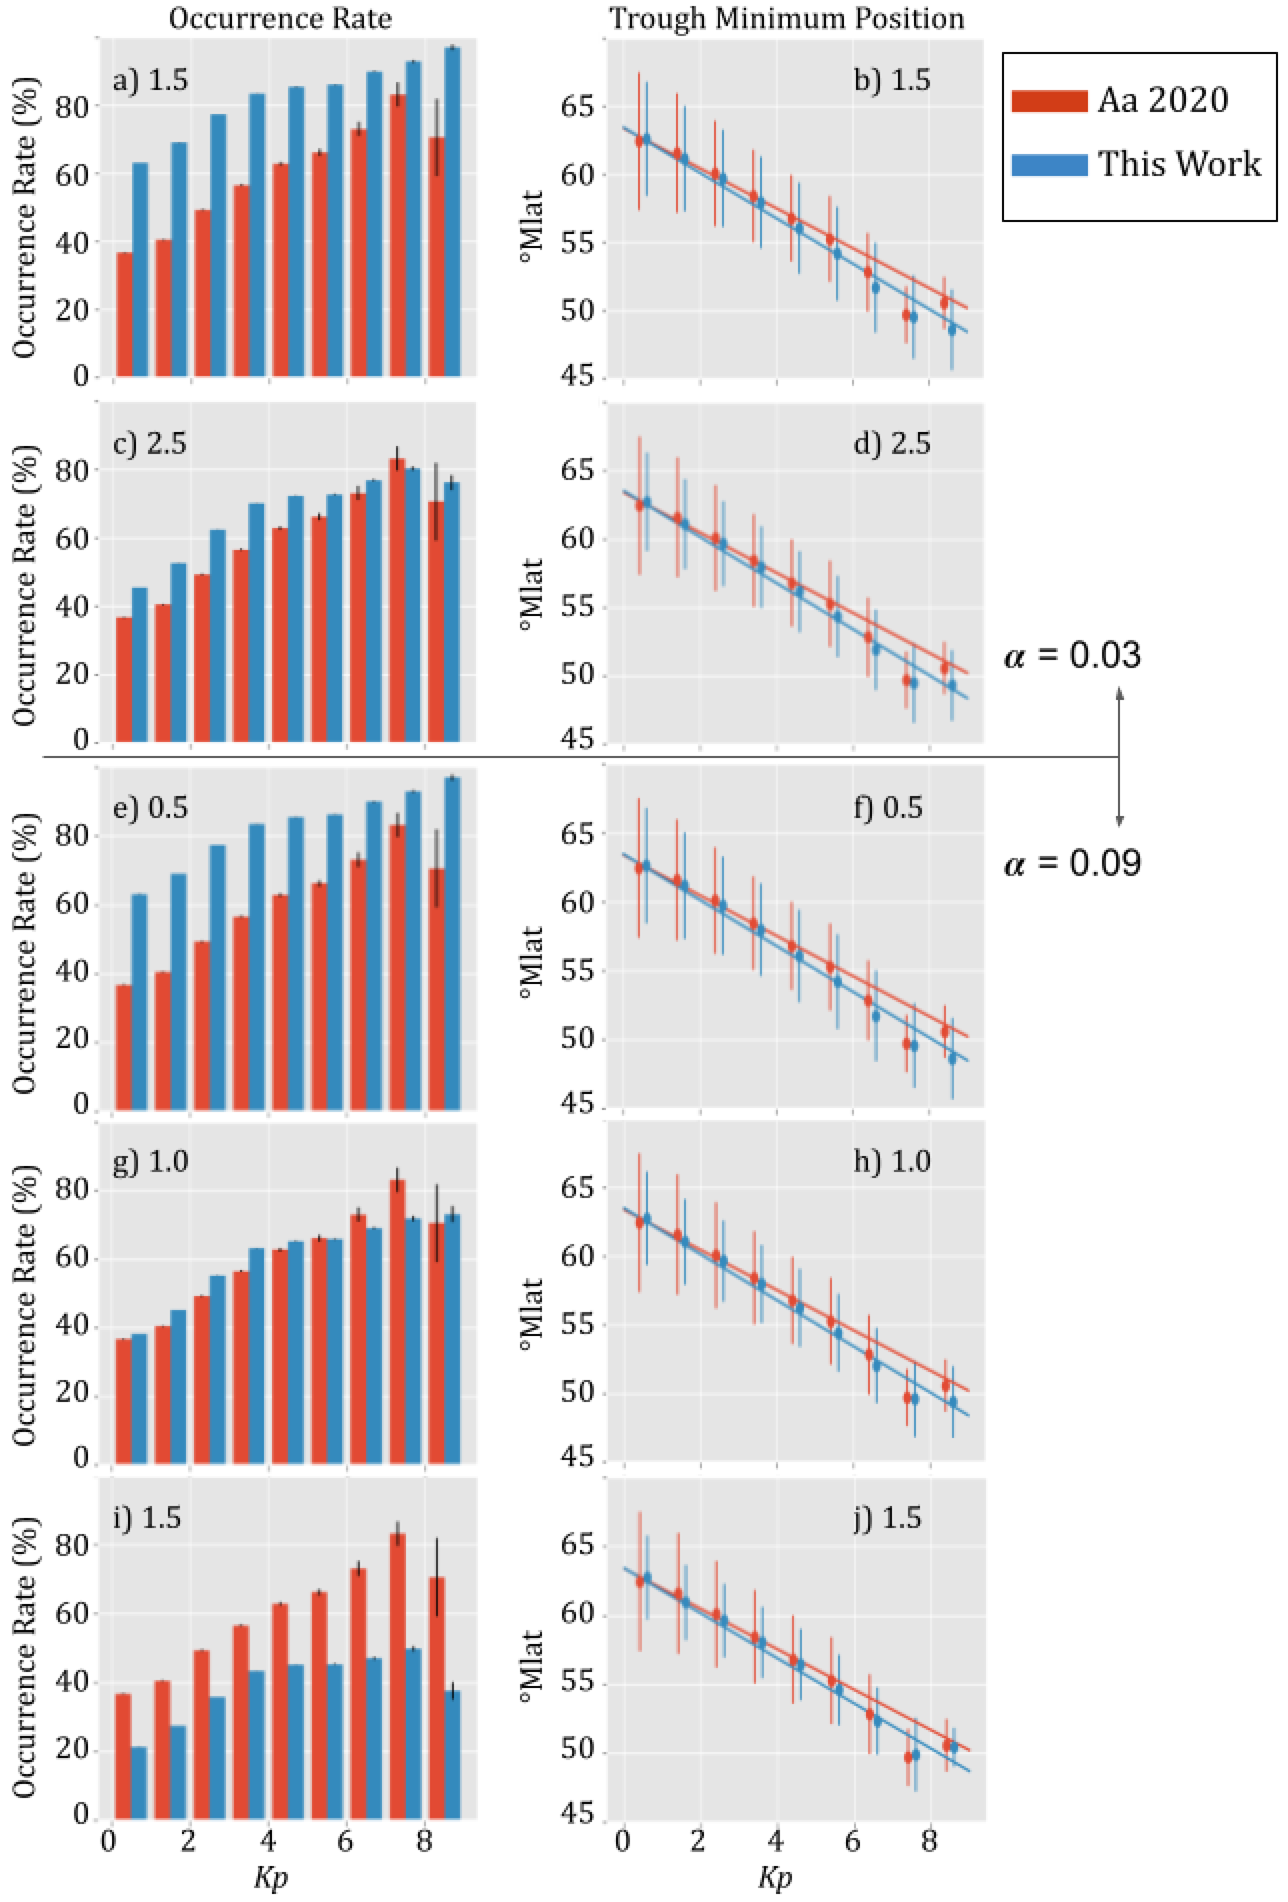

Occurrence Rate
Trough Minimum Position
Aa 2020
This Work
a) 1.5
b) 1.5
c) 2.5
d) 2.5
e) 0.5
f) 0.5
g) 1.0
h) 1.0
i) 1.5
j) 1.5
α = 0.03
α = 0.09
Occurrence Rate (%)
°Mlat
Kp