

2021

Enabling statistical analysis of the main ionospheric trough with computer vision

<https://hdl.handle.net/2144/43082>

"Downloaded from OpenBU. Boston University's institutional repository."

BOSTON UNIVERSITY
COLLEGE OF ENGINEERING

Thesis

**ENABLING STATISTICAL ANALYSIS OF THE MAIN
IONOSPHERIC TROUGH WITH COMPUTER VISION**

by

GREGORY WALTER SIDOR STARR

B.S., Boston University, 2017

Submitted in partial fulfillment of the
requirements for the degree of
Master of Science

2021

© 2021 by
GREGORY WALTER SIDOR STARR
All rights reserved

Approved by

First Reader

Joshua L. Semeter, PhD
Professor of Electrical and Computer Engineering

Second Reader

Yukitoshi Nishimura, PhD
Research Associate Professor of Electrical and Computer Engineering

Third Reader

Prakash Ishwar, PhD
Professor of Electrical and Computer Engineering
Professor of Systems Engineering

Acknowledgments

I would like to thank Professor Josh Semeter for his guidance and support over the last few years. I have been working in his lab since I was a sophomore and the experience has had a major influence on how I think as an engineer. He has also provided me much needed direction and assistance with this thesis. For providing guidance to me on this and past projects, I would like to thank Professor Toshi Nishimura, Sebastijan Mrak and Professor Michael Hirsch. Toshi and Sebastijan have taught me everything I know about space physics and Michael has been my software development role model for over 5 years. I would also like to thank Shibaji Chakraborty from Virginia Tech for providing SuperDARN data and answering all my questions about how to utilize it. For serving as advisors on my committee, I would like to thank Professors Joshua Semeter, Toshi Nishimura and Prakash Ishwar. Finally I would like to thank my friends and family, for their support over the course of this project. I would especially like to acknowledge my parents and Jess for reading and critiquing my thesis-related writing.

Gregory Starr

ENABLING STATISTICAL ANALYSIS OF THE MAIN IONOSPHERIC TROUGH WITH COMPUTER VISION

GREGORY WALTER SIDOR STARR

ABSTRACT

The main ionospheric trough (MIT) is a key density feature in the mid-latitude ionosphere and characterizing its structure is important for understanding GPS radio signal scintillation and HF wave propagation. While a number of previous studies have statistically investigated the properties of the trough, they have only examined its latitudinal cross sections, and have not considered the instantaneous two-dimensional structure of the trough. In this work, we developed an automatic optimization-based method for identifying the trough in Total Electron Content (TEC) maps and quantified its agreement with the algorithm developed in (Aa et al., 2020). Using the newly developed method, we created a labeled dataset and statistically examined the two-dimensional structure of the trough. Specifically, we investigated how K_p affects the trough’s occurrence probability at different local times. At low K_p , the trough tends to form in the postmidnight sector, and with increasing K_p , the trough occurrence probability increases and shifts premidnight. We explore the possibility that this is due to increased occurrence of troughs formed by subauroral polarization streams (SAPS). Additionally, using SuperDARN convection maps and solar wind data, we characterized the MIT’s dependence on the interplanetary magnetic field (IMF) clock angle.

Contents

1	Introduction	1
1.1	Motivation	1
1.2	Thesis Overview	5
2	Background	7
2.1	Space Environment Near Earth	7
2.1.1	Electron Density and Total Electron Content	7
2.1.2	Ionospheric Continuity Equation	9
2.2	Main Ionospheric Trough	14
2.2.1	MIT Identification	19
2.3	Inverse Problems and Regularization	20
3	Methods	22
3.1	Datasets	23
3.1.1	Madrigal GPS TEC	23
3.1.2	SWARM	25
3.1.3	Defense Meteorological Satellite Program - Special Sensor Ultraviolet Spectrographic Imager (DMSP - SSUSI)	27
3.1.4	Geophysical Indices and Solar Wind	28
3.1.5	SuperDarn	29
3.2	Problem Setup	30
3.3	MIT Labeling Methods	32
3.3.1	General Considerations	33

3.3.2	Trough Labeling Method 1 (TLM1)	39
3.3.3	Trough Labeling Method 2 (TLM2)	44
3.3.4	Performance	47
4	Experiments	56
4.1	Replication	57
4.1.1	MLat vs Kp	57
4.1.2	MLat vs MLT	60
4.1.3	Occurrence Rate	60
4.2	Clock Angle	65
4.2.1	B_z Effect on MIT Position	67
4.2.2	IMF Effect on MIT Occurrence Regions	67
4.3	Kp - MLT	73
5	Conclusion	77
5.1	Summary	77
5.2	Future Work	78
5.2.1	Data Processing	78
5.2.2	Applications	79
	References	81
	Curriculum Vitae	85

List of Tables

3.1	Binary error types	49
3.2	Parameter settings for <i>TLM1</i> , found with random search	50
3.3	Parameter settings for <i>TLM2</i> , found with random search	51
3.4	Performance comparison at all longitudes. <i>B1</i> : 2-D Aa 2020; <i>B2</i> : Deminov 2018	53
3.5	Performance comparison excluding magnetic longitudes [130, 260]. <i>B1</i> : 2-D Aa 2020	55
4.1	Gaussian dithering standard deviations	56

List of Figures

1·1	An example of the MIT in TEC measurements. It appears here as a dark curve spanning the entire night-side ionosphere. Dashed lines show satellite orbits (see text).	2
2·1	Copied figure 1 from (Vierinen et al., 2016) illustrating sTEC (S line) vs vTEC (V dashed line).	9
2·2	Copied figure 5 from (Thomas and Shepherd, 2018) showing average convection patterns for different IMF clock angles.	13
2·3	Diagram of the stagnation mechanism copied from (Spiro et al., 1978)	16
2·4	Map of SAPS occurrence rate copied from (Foster and Vo, 2002) . . .	18
3·1	Madrigal GPS TEC	23
3·2	Conversion and averaging process for Madrigal TEC maps	24
3·3	(Left) TEC data coverage over magnetic longitude. (Right) TEC data coverage over time.	25
3·4	Copied from figure 1 in (Aa et al., 2020). Examples of the trough identification method for SWARM data. The dotted line in the "Ne" plots shows the background electron density.	27

3.5	Overview of the data used in this project. Top left is a plot of solar wind measurements. The solar wind influences the ionospheric convection pattern shown top middle. The convection pattern carries plasma around the ionosphere and is largely responsible for the TEC distribution, shown on the right. Under the TEC plot and to the left are trough labels identified in TEC and SWARM data.	31
3.6	Preprocessing	34
3.7	<i>TLM1</i> performance metrics at different thresholds. Top: Accuracy (Acc.), true positive rate (TPR), true negative rate (TNR), false positive rate (FPR), false negative rate (FNR). Bottom: Poleward ("P") and equatorward ("E") wall latitude error mean and standard deviation ("mean", "std. dev.")	38
3.8	<i>TLM1</i> intermediate vectors	43
3.9	Overview of <i>TLM2</i> scoring step (see " <i>TLM2</i> Overview Diagram Description" in text for details).	45
3.10	(a) SWARM satellite MLAT vs time, showing time alignment of satellite orbital segments with TEC maps. Black vertical lines indicate one TEC map, blue and yellow lines mark selected segments. (b) SWARM trough vs TEC trough comparison diagram.	48
3.11	Poleward and equatorward wall latitude error distributions for the three methods	54
4.1	MIT position and occurrence rate dependence on Kp . (a) Figure 8b Aa 2020: MIT position vs Kp in various studies. (b) Figure 8a Aa 2020: MIT occurrence rate vs Kp . MIT position vs Kp in our dataset, within 2 hours (c) and 5 hours (d) of midnight. (e) (left) Distribution of MIT Kp and MLat, (right) occurrence rate of MIT in Kp - MLat bins.	58

4.2	MIT Position vs MLT as shown in several studies (a) (Yang et al., 2015); (b) (Le et al., 2017); (c) (Aa et al., 2020), bars indicate poleward and equatorward walls; (d) ours, bars indicate standard deviation . . .	61
4.3	MIT Occurrence rates for $Kp \leq 3$. (a) (Aa et al., 2020), (b) ours . . .	62
4.4	MIT occurrence rate in season - MLT bins. (a) (Aa et al., 2020), white lines indicate average solar terminator at 60° MLat. (b) MIT season-MLT distribution and occurrence probability on the left and right respectively.	64
4.5	Trough occurrence rate during east, south and west IMF clock angle. Top: absolute occurrence rate. Bottom: Occurrence rate differences between top row plots.	65
4.6	MIT position and occurrence rate dependence on B_z . $2 \leq B \leq 6$ nanotesla, MIT labels delayed by 1 hour from IMF measurement . . .	67
4.7	Average TEC profile of evening MIT (16 - 20 MLT) corrected for season	68
4.8	Average TEC profile of morning MIT (4 - 8 MLT) corrected for season	69
4.9	Average convection and corrected TEC in the vicinity of the MIT . . .	70
4.10	Motion of MIT towards pre-midnight with increased Kp	72
4.11	SAPS contribution to MIT occurrence rate	74
4.12	Portion of troughs associated with SAPS	76

List of Abbreviations

DMSP		Defense Meteorological Satellite Program
EISCAT		European Incoherent Scatter (radar)
GNSS		Global Navigation Satellite System
GSM		Geocentric Solar Magnetic (coordinates)
IGRF		International Geomagnetic Reference Field
IID		Independent and Identically Distributed
IMF		Interplanetary Magnetic Field
LOS		Line-of-sight
MIT		Main Ionospheric Trough
MLat		Magnetic Latitude
MLon		Magnetic Longitude
MLT		Magnetic Local Time
NaN		Not a Number
RBF		Radial Basis Function
ROC		Receiver Operating Characteristic (curve)
SSUSI		Special Sensor Ultraviolet Spectrographic Imager
TEC		Total Electron Content
TEC _u		Total Electron Content Units
TLM		Trough Labeling Method (1, 2)
TV		Total Variation (regularizer)

Chapter 1

Introduction

1.1 Motivation

The plasma in the ionosphere affects all electromagnetic waves that pass through it. This can have a big impact on trans-ionospheric communication via radio. If the signals pass through a quiet / uniform region of the ionosphere, the effects are not significant or they can be easily corrected for, however, when the ionosphere becomes more turbulent communications can be disrupted. Every year, we rely more and more on being able to send signals through the ionosphere. More satellites are sent into orbit every year and they are becoming a part of every aspect of our lives. Crucial portions of our infrastructure rely on satellites including weather forecasting, positioning and navigation, and recently even internet. In the distant future, communicating with humans on other planets will require sending signals through two ionospheres. Understanding the ionosphere's effects on our communications and being able to predict ionospheric conditions are crucial to humanity's future in space and on other planets.

One particularly disruptive ionospheric phenomenon is the main ionospheric trough (MIT). The MIT is a band of low electron density which forms between the high latitude and mid latitude regions of the ionosphere. It is among the most consistently observed and largest scale features of the ionosphere. It can occur in both hemispheres and during any season. The trough can negatively affect communication signals due to the large electron density gradients as well as smaller scale irregularities which form




Figure 1.1: An example of the MIT in TEC measurements. It appears here as a dark curve spanning the entire night-side ionosphere. Dashed lines show satellite orbits (see text).

in it (Rodger et al., 1992), (Kintner et al., 2007), (Le et al., 2017). Because of the trough’s continental size, it can have a serious impact on our ability to communicate with satellites. Understanding the trough in particular will help us predict where interruptions will occur and how to mitigate them. However, the trough is also a major part of the larger ionospheric and magnetospheric systems of the earth, so advancing our knowledge of the trough will also help us understand those systems on earth and other planets. Finally, because the trough is dynamic and spread across the world, measuring it is a challenge. The only way we can accomplish this is with a large and diverse set of sensors, and analyzing such a dataset requires new techniques. These new techniques for data collection and analysis will be hugely beneficial to science and engineering in general.

Despite the importance of the MIT, it is still not fully understood. Many of its characteristics have been established through measurement and its position is well-modeled empirically. Additionally, the primary mechanisms which create and maintain the trough have been identified. However there is still a lot left to discover. The relative importance of the various mechanisms has not been fully established, and some of the most popular ionospheric models do not properly reproduce its behavior (Yang et al., 2015). Finally, adequate statistics of how the MIT behaves during heightened periods of geomagnetic activity have not been obtained. The full spectrum of MIT behavior including what edge cases exist is not yet known.

Progressing our understanding of the MIT is difficult. Because it can get very large, simultaneously measuring the MIT along its length requires a vast network of sensors. Most existing studies have been conducted with satellite measurements which inherently can only collect data at a single location at a time. Understanding the generation mechanisms requires measurements of many different ionospheric parameters, some of which can only be made with expensive radar facilities, e.g. plasma

flow. In order to specifically study the MIT in a dataset, you first need to identify it. This is also difficult because there is no concrete definition which qualifies an electron density depletion as the MIT. To make things worse, there are also other troughs which occur in the ionosphere like the high latitude trough or the ring ionospheric trough. These troughs can physically overlap with one another which makes it difficult to decide whether they should be classified differently. Such phenomena should be classified according to the underlying physics which produce them, but this is also tricky because the MIT is caused by many different processes. While datasets exist that contain global measurements of total electron content, it is more difficult to label a 2D dataset than 1D, either manually or automatically and no one has undertaken the task of developing and testing a method to accomplish this.

Though two-dimensional data is more difficult to process, we believe that it will ultimately lead to a much better dataset of the MIT. Previous statistical studies of the MIT have mostly only considered latitudinal cross sections rather than global 2D or 3D scalar measurements. This has two drawbacks. The first is that it inherently has less data. For example, in Aa 2020, they 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 would make 350,000 measurements of the trough region. A two-dimensional TEC dataset has measurements at most local times (conservatively 90) and so you could think of it as having almost 8,000,000 measurements of the trough. Looking at figure 1-1 gives an idea of the vastly increased coverage with 2D data over 1D. Using a 2D dataset like Madrigal GPS TEC will allow for statistics to be compiled at a much higher level of detail.

Given that such significant progress has been made in the field of machine learning recently, it is an exciting prospect to utilize some of the new techniques for space physics problems. However, it has been difficult for researchers in the space physics

to find good uses for such techniques. The most popular problem settings in machine learning, for example image classification, do not have an obvious application for space physics. One reason for this is that the ionosphere does not produce discrete categories. Most ionospheric phenomena exist on a continuum and are caused by a variety of interacting processes which makes categorizing them difficult and sometimes futile. Additionally, complex machine learning models which have the best chance of being able to represent ionospheric phenomena are notoriously difficult to interpret. If you want to model a phenomenon then there is a trade off you have to make between expressibility and interpretability and this becomes a difficult choice if your goal is to understand the underlying physics.

1.2 Thesis Overview

The goal of this thesis is to enable statistical analysis of the MIT by developing methods to automatically identify it in global maps of total electron content. Creating an automatic method is crucial because it allows us to apply it to large datasets, which is necessary to properly establish the statistics of the trough. To accomplish this, we cast the problem into the framework of image segmentation, where a label is assigned to each pixel of an input image. In this case, the labels are binary: a pixel can either be part of the trough or not. This approach has the advantage of producing a very descriptive and flexible dataset from which a large variety of measurements can be easily taken. We developed two methods and established their validity by quantifying their agreement with the method used by Aa et. al for satellite measurements of electron density (Aa et al., 2020). Finally, we performed three "replication" experiments to further validate our labeled dataset, and two experiments in which we produced novel results about the mechanisms responsible for MIT formation.

In this chapter we explained the importance of the MIT, the gaps that exist in

our understanding of it and we outlined the goals of this project. Next, in chapter two, we introduce background information about the ionosphere and MIT. In chapter three we describe the datasets we are used, how we processed them and the details of the trough identification methods. In chapter four we discuss the experiments we performed to demonstrate the labeled dataset's value and improve our understanding of the MIT. Lastly in chapter five, we summarize the project and our plans for future work.

Chapter 2

Background

2.1 Space Environment Near Earth

The various "-spheres" of the atmosphere are defined by the characteristics and behaviors of the gas contained in them. Through various processes, the gasses of the atmosphere can become ionized. The ionosphere is the region from about 100km to 1000km in altitude, where the atmospheric gas has been ionized to a significant level such that it behaves as a plasma rather than neutral gas. At higher altitudes, the atmosphere becomes thin and collisions between particles rare. In this region, the motion of particles is determined by the earth's magnetic field and so it is called the magnetosphere.

2.1.1 Electron Density and Total Electron Content

Electron density is defined as the number of free electrons per unit volume. It can be measured by instruments on satellites, rockets, or with radar. The refractive index of the ionospheric plasma depends on the wavelength of the incident radiation and the electron density of the plasma. The plasma imparts a phase shift on all radio waves which pass through it. The additional phase shift of a wave travelling through the ionosphere is proportional to the electron density integrated along the path from transmitter to receiver.

$$\Delta\Phi \propto \text{TEC} = \int_{\mathcal{P}} N(\mathbf{r}) dl \quad (2.1)$$

In the above expression, $\Delta\Phi$ is the ionospheric phase shift of a wave travelling along the path $\mathcal{P}$, N is the electron density at point $\mathbf{r}$, and dl is an infinitesimal distance along $\mathcal{P}$. Integrated electron density is called total electron content (TEC) and it is a density value which varies in two dimensions rather than three. It has units of electrons/m² but it is usually presented in "total electron content units" (TECu) where 1 TECu = 10¹⁶ electrons/m². Since the total phase shift on waves transmitted through the ionosphere depends on their frequency, TEC can be measured by transmitting and receiving on two separate frequencies. Global Navigation Satellite Systems (GNSS) utilize this method to remove phase shifts and improve their positioning accuracy, and as a by-product, they provide useful data for space science research.

There are two ways to define TEC. Slant TEC (sTEC) is the quantity shown in equation 2.1, which is integrated along the satellite-receiver line-of-sight (LOS). While sTEC is useful for correcting GPS errors, it strongly depends on the geometry of the measurement, in particular, the elevation angle of the satellite-receiver LOS. A lower elevation LOS has a longer intersection with the ionosphere, increasing the sTEC. This makes it difficult to compare measurements of TEC from multiple satellite-receiver pairs. To address this shortcoming, researchers use vertical TEC (vTEC) instead, which is the integrated electron density along a vertical line. In the special case where a satellite-receiver LOS is a vertical line, the sTEC and vTEC are equivalent. Since true vTEC measurements are only available for a small subset of LOS's, it is always estimated from sTEC measurements. This is accomplished by assuming some electron density altitude profile for the ionosphere and deriving a "mapping function" which, when multiplied by the sTEC, produces the vTEC. The most common way to perform this general process is to assume the ionosphere is a thin shell at a fixed height. Then the mapping function only depends on the elevation angle. Figure 2.1




Figure 2.1: Copied figure 1 from (Vierinen et al., 2016) illustrating sTEC (S line) vs vTEC (V dashed line).

was copied from figure 1 in (Vierinen et al., 2016) and it illustrates the concept. It shows a slightly more complicated altitude profile (which they call "Shape function"), but it still results in a mapping function v which only depends on the LOS elevation angle α .

2.1.2 Ionospheric Continuity Equation

The electron density in the ionosphere changes according to the continuity equation:

$$\frac{\partial N}{\partial t} = q - \beta N - \nabla \cdot (N\mathbf{V}) \quad (2.2)$$

where N is the electron density, q is the production rate, βN is the loss rate due to chemical recombination and the final term represents change due to the bulk motion of the gas, whose velocity is $\mathbf{V}$. The two main sources which make up the production term q are solar ionization and ionization from energetic particles. The rate of solar ionization depends on the elevation of the sun, the intensity of the radiation, and the availability ionizable gas, which decreases at higher altitudes. At night, the only source of ionization is energetic particles.

The earth's magnetosphere contains trapped electrons in regions known as radiation belts. Electrons gyrate around magnetic field lines due to the Lorentz Force. The angle between an particle's velocity and the field line around which it gyrates is called the pitch angle. As an electron travels from the equatorial plane towards the poles along a field line, the magnetic field strength increases which causes the pitch angle to increase. At some point, called the mirror point, the magnetic field strength will cause the pitch angle to become greater than 90° , at which point the electron will begin moving in the opposite direction along the field line. The altitude of the mirror point is determined by the electron's pitch angle at the equator. For a range of equatorial pitch angles, the mirror point is at a low enough altitude to be inside the atmosphere. This is called the loss cone because most electrons within that pitch angle range will collide with the particles in the atmosphere rather than reversing their direction. This is known as electron precipitation. A precipitating electron can ionize or excite the gasses in the atmosphere, depending on its energy. The high latitude region where most electron precipitation occurs is called the auroral oval. Electrons with pitch angles outside the loss cone continue to bounce between northern and southern mirror points. Various interactions can occur which will change a trapped electron's pitch angle, possibly causing it to precipitate.

Magnetic Coordinate Systems

Within 2 earth radii ($R_E = 6,371\text{km}$) of the surface, the magnetic field is roughly approximated by a dipole. Because so many ionospheric phenomena are related to the magnetic field, it is usually more revealing to study them using magnetic coordinate systems. The simplest of these coordinate systems is called centered dipole, in which the z axis is aligned with the magnetic dipole (positive north), the y axis is defined such that it is perpendicular to the plane containing the magnetic dipole axis and the earth's rotation axis, and x completes the right handed coordinate system (Laundal

and Richmond, 2016). At ionospheric altitudes (100 - 1000km), the dipole approximation is valid but not very accurate, and so if a closer correspondence is required, then researchers will use a coordinate system based on a more detailed model. One such coordinate system is called magnetic apex coordinates (Richmond, 1995), (Emmert et al., 2010), which determines longitude and latitude by tracing the field lines of the International Geomagnetic Reference Field (IGRF). The IGRF is a periodically updated spherical harmonics model of the earth's magnetic field. Magnetic apex coordinates have the property that latitude and longitude are constant along any field line. Since the sun is a major driver of the dynamics of the ionosphere, often times researchers will replace longitude with magnetic local time (MLT). MLT is defined so that the magnetic longitude of the subsolar point is noon and 180° longitude from the subsolar point is midnight. In the below equation for MLT, ϕ_s is the magnetic longitude of the subsolar point and ϕ is the magnetic longitude of the point (Laundal and Richmond, 2016).

$$\text{MLT} = (\phi - \phi_s)/15 + 12 \quad (2.3)$$

We utilize magnetic apex coordinates with magnetic local time throughout this work.

Solar Wind and Interplanetary Magnetic Field

The sun constantly emits a stream of particles, mostly electrons and hydrogen ions, called the solar wind. The wind carries a weak magnetic field with it which is called the interplanetary magnetic field (IMF). The magnetopause is the boundary between the earth's magnetosphere and the solar wind. The magnetosphere is where a particle's motion is controlled by the earth's magnetic field. On the other side of the magnetopause, the motion of the particles is controlled by the sun. As the solar wind and IMF encounter the earth's magnetic field, it is distorted into a tear-drop shape. The coordinate system which is typically used for solar wind measurement is called

geocentric solar magnetic (GSM), in which the x axis points from the earth to the sun, the y axis is perpendicular to both the x axis and the dipole axis, and the z axis completes a right handed coordinate system (Laundal and Richmond, 2016). Key variables of the solar wind plasma include its speed, density and temperature. More important to this study are the IMF parameters including its z and y components, its magnitude, and its angle in the yz plane, clockwise from $+z$, called the IMF clock angle. In the equation below for clock angle, arctan is the full 360-degree version.

$$\theta_{\text{clock}} = \arctan(B_y, B_z) \quad (2.4)$$

Convection

Circulations and plasma flow velocity in the ionosphere can be described with electric fields according to the following equation:

$$\mathbf{V} = \frac{\mathbf{E} \times \mathbf{B}}{|\mathbf{B}|^2} \quad (2.5)$$

This is the drift a charged particle would experience in both an electric and magnetic field (called "E-cross-B drift"). Fields are just a mathematical tool to describe electromagnetic interaction. The laws of electromagnetism can be thought of in multiple ways depending on the reference frame. For example, if a magnet passes through a loop of wire, in the reference frame of the magnet, you would say that the current is due to the electrons accelerating because of the Lorentz force. However, in the reference frame of the wire, you would say that the changing magnetic field created an electric field according to Faraday's law. The two viewpoints are equivalent, which is why we can describe plasma motion in the ionosphere with electric fields.

Part of the IMF merges with the earth's magnetic field, and pulls it back in the antisunward direction. The magnetic field lines moved by the solar wind drag ionospheric plasma with them, creating circulations. In the sun-earth reference frame,




Figure 2.2: Copied figure 5 from (Thomas and Shepherd, 2018) showing average convection patterns for different IMF clock angles.

the ionosphere has two electric fields. One is due to the magnetosphere's interaction with the solar wind and is called the convection electric field, and the other is due to the earth's rotation and is known as the corotation electric field. Close to the equator, the plasma mostly moves according to the corotation electric field and towards the poles, it moves according to convection. Since convection is driven by the solar wind, its shape is related to the IMF. An example of average convection patterns at different IMF clock angles is shown in figure 2-2, which was copied from (Thomas and Shepherd, 2018). The convection electric field and the resulting convection flows are typically visualized with the scalar electric potential. $\mathbf{E} \times \mathbf{B}$ drift implies that plasma motion is perpendicular to the electric field, which means that plasma travels along electric equipotential surfaces. From figure 2-2 we can see that at all IMF clock angles, plasma generally flows over the polar cap from MLT 12 to MLT 0 (magnetic noon to midnight), then returns to the dayside at lower latitudes. The two "cells" of the convection pattern are referred to as the dawn cell (right) and dusk cell (left). The z component of the IMF controls the overall strength of the convection electric field and the y component rotates its orientation about the dipole axis.

2.2 Main Ionospheric Trough

The main ionospheric trough (MIT), sometimes called the mid-latitude ionospheric trough, is a large region of decreased electron density occurring in the sub-auroral ionosphere. It does not have a concrete definition, but it is usually 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). When writing about the MIT, authors will often begin by specifying what their definition is. At this time, it is best defined by its characteristics, though a better definition would be based on its physical mechanisms. Though some of these

mechanisms are known, their interactions and relative contributions are still unclear. A better understanding of the ionosphere in general will be required to properly classify ionospheric troughs.

The main ionospheric trough is often described as having three parts: 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 and Larson, 2010). The electron density gradient is typically much stronger at the poleward wall than at the equatorward wall (Spiro et al., 1978). The MIT is observed most often in darkness and has an average width of about 5° to 10° of latitude (Aa et al., 2020), (Yang et al., 2015), (Collis and Häggström, 1988). It can be observed in wide-area TEC maps (Zou et al., 2011) and sequential radar scans (Nilsson et al., 2005) that the MIT is very elongated longitudinally. Since it mainly occurs in darkness, its length (longitudinally) is strongly correlated with season (Rodger, 2008), though its length has not yet been directly quantified.

The most frequently studied parameter of the trough is the latitude at its minimum. Like many of the MIT's other parameters, this latitude varies across its length. Depending on how the MIT is defined, its highest is 75° to 80° latitude at noon and steadily decreases with MLT until it reaches its minimum of 55° to 60° latitude 3 - 5 hours MLT after midnight (Werner and Prölss, 1997) (Aa et al., 2020) (Yang et al., 2015). This latitude also varies with the level of magnetic disturbance. As magnetic activity level increases, the convection pattern and the auroral precipitation region expand. This causes the MIT to move to lower latitudes. Many studies have estimated linear models relating the latitude of the MIT to various measures of magnetic activity. Less common examples include a time-averaged version of the au-




Figure 2-3: Diagram of the stagnation mechanism copied from (Spiro et al., 1978)

roral electrojet index (Werner and Prölss, 1997) and the Dst index (Karpachev et al., 1995), but many researchers have used the Kp index (Yang et al., 2015), (Collis and Häggström, 1988), (Deminov and Shubin, 2018), (Dudeney et al., 1983), etc. A recent example is from (Aa et al., 2020):

$$T_M = 65.8^\circ - 1.7Kp \quad (2.6)$$

Many mechanisms which could contribute the trough have been proposed and evaluated. The most commonly mentioned is called the stagnation mechanism. In the evening side ionosphere near the equator, the plasma flow is corotational i.e. antisunward, but at higher latitudes, in the return flow of the dusk convection cell, the flow direction is sunward. At some latitude in between, these two flows must cancel each other out and create a stagnation zone. Plasma equatorward of the stagnation zone, having recently corotated from the dayside, would have a relatively high electron density. Similarly, plasma poleward of the zone would have increased electron density from auroral precipitation. However plasma that winds up in the stagnation zone would remain for an extended period, during which time it would recombine. A schematic of the scenario is shown in figure 2-3. Regions B and C are where flow

reversal and stagnation occurs, whereas region A is dominated by corotation and region D by convection (Spiro et al., 1978). This theory was analyzed, compared with measurements and determined to be plausible in (Spiro et al., 1978) and (Nilsson et al., 2005). In (Collis and Häggström, 1988), the authors 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.

The stagnation mechanism can explain evening and premidnight troughs but not postmidnight or morning troughs. A more general convection based theory, of which the stagnation mechanism is a special case was explained in (Quegan et al., 1989). The authors emphasize that the electron density at any location is due to all of the production and loss along the path on which that plasma travelled. The convection pattern is very complex which can result in two paths with very different histories being brought close together. A trough can form when a path with a short transit from the dayside ends up near a path which meandered on the night side. The poleward wall is built up from auroral precipitation and dayside plasma convected across the polar cap.

Another mechanism which can form troughs is sub-auroral polarization stream (SAPS). SAPS is a high speed flow channel in the sub auroral ionosphere. In a nonlinear process, the high wind speed increases frictional heating, and the increased temperature increases the plasma recombination rate. This creates a local plasma depletion. The occurrence of SAPS was quantified in (Foster and Vo, 2002), and a map copied from their paper is shown in figure 2.4. Because the behavior of the MIT at high Kp , namely that it moves equatorward and is observed having a higher occurrence rate, matches that of SAPS, many researchers believe that SAPS troughs could be contributing to the statistics of the MIT (Aa et al., 2020). Whether or not SAPS troughs should be considered along with or separately from the MIT is up for




Figure 2.4: Map of SAPS occurrence rate copied from (Foster and Vo, 2002)

debate, though distinguishing them from electron density measurements alone can be difficult.

2.2.1 MIT Identification

In past statistical studies of the MIT, many different methods were utilized to identify the trough. Of course one option which is frequently used is manual identification, but this becomes infeasible for large datasets. Algorithmic approaches have so far only been developed for one-dimensional data. It seems that existing methods follow the same basic two 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, the first step is to process it into a latitudinal profile. For latitude-altitude measurements from the European Incoherent Scatter radar (EISCAT), the authors of (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 in (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". In (Yang et al., 2015), the authors used the same TEC dataset as us, but they computed latitudinal profiles by averaging TEC over the course of a day in two hour 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. An identical approach was taken in (Pryse et al., 2006) except using TEC computed from tomography data. In (Aa et al., 2020), they computed the background electron density measured in-situ by satellites, then used a threshold of 50% to identify the

trough. This is described in greater detail in section 3.1.2. Finally, one exception is (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 $-0.1 \text{ TECu} / \text{degree}$.

2.3 Inverse Problems and Regularization

Often times in science and engineering we have a set of measurements and we wish to compute a set of physical parameters which produced our measurements. This is known as an inverse problem. The purpose of solving an inverse problem is to obtain an estimate for parameters which are difficult or impossible to measure directly. So we take measurements of related parameters and attempt to estimate the parameters we are actually interested in. The measurements are related to the parameters of interest via the forward model function f , as shown in the following equation, where $\mathbf{x}$ is a vector of measurements and $\mathbf{u}$ is a vector of the parameters we are interested in.

$$\mathbf{x} = f(\mathbf{u}) \tag{2.7}$$

Then the goal of the inverse problem is to estimate $f^{-1}(\mathbf{x})$. The difficulty is that in most practical cases, f^{-1} has no closed form or is not a function because it does not map $\mathbf{x}$ to $\mathbf{u}$ one-to-one. To get around this, the inverse problem is solved by searching for a value of $\mathbf{u}$ which minimizes a norm of the error between the measurements $\mathbf{x}$ and the theoretical measurements resulting from $\mathbf{u}$. This is shown in the following optimization problem where $\mathbf{u}^*$ is the estimated vector of parameters.

$$\mathbf{u}^* = \arg \min_{\mathbf{u}} \|\mathbf{x} - f(\mathbf{u})\| \tag{2.8}$$

However, one difficulty which is often present is that multiple, possibly infinite, values of $\mathbf{u}$ can produce the same output of f . A simple example of this is when f is a linear

function, i.e. $f(\mathbf{u}) = A\mathbf{u}$; $A \in \mathcal{R}^{m \times n}$; $m < n$. Then $f(\mathbf{u}) = f(\mathbf{u} + \mathbf{v})$ for any vector $\mathbf{v}$ in the null space of A . In this case, equation 2.8 will not be sufficient to solve the inverse problem because every time you perform the minimization you could get a different answer. Problems which have this property are known as "ill-posed" problems, because they don't have a unique solution.

One way to deal with this difficulty is to incorporate prior information into the objective function by adding additional terms. This is called regularization. The additional terms are meant to increase the cost for values of $\mathbf{u}$ which don't conform to prior beliefs. The most common regularization terms are the $L2$ and $L1$ norms of $\mathbf{u}$. Both improve the posedness of the inverse problem by favoring low-norm solutions. Using the $L1$ norm tends to produce sparser solutions, but also is more computationally expensive. A more typical example of an inverse problem is given in the following equation, where λ is a positive scalar which controls the regularization strength.

$$\mathbf{u}^* = \arg \min_{\mathbf{u}} \|\mathbf{x} - f(\mathbf{u})\|_2^2 + \lambda \|\mathbf{u}\|_2^2 \quad (2.9)$$


Depending on the form of f , the above equation could even have a closed form solution, which would make the inverse problem very fast to solve.

Chapter 3

Methods

The goal of this project was to create a large, flexible labeled dataset of the trough in TEC images. In order to make the dataset large, we needed to develop methods to automatically estimate the labels. There were many ways we could have parameterized the trough labels. One option was to estimate the latitudes of the MIT's poleward wall, equatorward wall and minimum at each MLT, producing a matrix of shape $(n \times l \times 3)$, where n is the number of timesteps in the dataset and l is the number of longitude / MLT bins. The problem with this type of approach is that you need to specify ahead of time all of the parameters you wish to estimate, i.e. it is not flexible. If we estimate the three latitude parameters and later decide that we also want to measure the TEC values or the TEC gradients at those positions, then we either have to reprocess the dataset, or we have to come up with a somewhat complex procedure to collect those parameters. This became especially inconvenient because we did not know which parameters would be useful to collect when we started.

Instead of that, we decided that the most natural and flexible way to parameterize the trough labels for a two-dimensional dataset would be to assign a binary label to each pixel of the input. An example of this can be seen in figure 3.5 on the bottom left. One of the biggest advantages of this approach is that it makes fewer assumptions about the shape of the trough, for example, whether it can exist at multiple latitudes at the same MLT. This approach also allows us to refine our estimates of specific parameters down the road. For example, if we wanted to measure attributes of the



(a) Madrigal GPS receiver network in 2020 (b) Example of a TEC map from Madrigal

Figure 3-1: Madrigal GPS TEC

poleward wall, we could take the top pixel at each MLT, or take some average of the surrounding area, etc. The important thing is that we did not lock ourselves into a particular way of doing anything by choosing this approach. Another advantage of this approach is that it makes it easy to select other types of data in the vicinity of the trough as long as it is on the TEC grid, for example, ion flow vectors. This flexibility may make the labeled dataset more useful for other researchers in the future.

3.1 Datasets

In this section we describe the various datasets we utilized and how we processed them. An overview of the different data sources we used is presented in figure 3-5.

3.1.1 Madrigal GPS TEC

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 binned and averaged into (1 degree latitude x 1 degree longitude x 5 minutes) bins. The network on 1/1/2020 had over 6000 receivers and is shown in figure 3-1a. The coverage is good over the United States and Europe but sparse everywhere else. An example of a corresponding TEC map is shown in figure 3-1b. The coverage is improved from the binning, but there is still very little over the oceans. Each file in




Figure 3.2: Conversion and averaging process for Madrigal TEC maps

the dataset has a TEC array of size (180 x 360 x 288) with NaNs (Not a Number) where data is missing.

Since we intended to use this dataset to study the MIT, a phenomenon organized by the magnetic field, we converted the coordinates to magnetic apex (Richmond, 1995). In magnetic apex coordinates, all points along the a magnetic field line are mapped to approximately the same coordinates. During this process, we also average several Madrigal TEC maps together to further improve their coverage. To do this, we first convert each latitude - longitude grid point in the Madrigal grid to apex latitude - MLT. The grid points of several consecutive Madrigal TEC maps are all converted, the number depends on how much averaging we want. Finally all these converted measurements are binned and averaged into a regular magnetic apex latitude - MLT grid which we call the TEC Grid. This process is illustrated in figure 3.2, where each dot in the left and middle plots are one grid point from a single Madrigal TEC map, and the right plot shows the resulting average over one hour in a magnetic coordinates grid. The grid cell size we chose is (1 degree latitude x 2 degrees longitude x 1 hour). Larger amounts of time-averaging result in higher coverage in each map, but less time resolution. For our dataset, we chose 1 hour because it seemed like a good balance between coverage and time resolution. Additionally, because the northern hemisphere has better coverage in the Madrigal dataset, we chose to limit our study to only magnetic apex latitudes above 30° North. Thus our resulting TEC images have a shape of (60 x 180).

Figure 3.3 shows the coverage of our TEC dataset from the years 2010 to 2020.




Figure 3-3: (Left) TEC data coverage over magnetic longitude. (Right) TEC data coverage over time.

In the left plot, the coverage of each latitudinal profile is averaged over the dataset. If the “Data coverage Proportion” is 1 for a particular magnetic longitude (MLon), that means that every grid cell along a latitudinal slice had data for every image in the dataset. This plot shows clearly the sparse coverage the data has over the oceans and Asia. This means we should expect any product or analysis based on this data to exhibit a strong 24 hour periodicity. We should treat any frequency analysis based on this data with caution. One way to avoid this problem is to add or average 24 maps together. The right plot shows the average coverage over time in the dataset in 6 month bins. If the “Data Coverage Proportion” is 1, that means that every grid cell had data during that particular 6 month interval. This plot shows that the receiver network is expanding over the course of the dataset. The years prior to 2010 have even less coverage and so the resulting TEC images are of limited value.

3.1.2 SWARM

In this project, we utilize the SWARM satellite in-situ electron density dataset (Lomidze et al., 2019), which we downloaded via FTP from `swarm-diss.eo.esa.int`

(specifically: `/Advanced/Plasma_Data/2_Hz_Langmuir_Probe_Extended_Dataset`). The SWARM satellite constellation consists of three satellites named A, B and C, which have been providing data since the end of 2013. Satellites A and C orbit at an altitude of 450km and are separated by 1.4° longitude. Satellite B orbits at about 510 km. All three satellites take about 1.5 hours to complete one orbit. The SWARM satellites carry a variety of instruments but we only used the electron density measurements which are provided at 2Hz. After downloading the dataset, we converted the satellite positions from geographic coordinates to magnetic apex coordinates, and consolidated the daily files into more convenient monthly files.

MIT Identification From SWARM

To verify and calibrate our TEC image trough identification methods, we used the method developed in (Aa et al., 2020) as a baseline. Their process is as follows. First they convert the electron density measurements to log-electron-density and remove noise with a three-point moving median filter. Then they estimate and remove the background plasma density level using a 480-point moving average filter, which corresponds to a horizontal distance of about 1800km. At this point, they split up the satellite orbits into segments in both hemispheres between 45° and 75° magnetic latitude (MLat). Within these segments, they identify negative peaks which achieve a minimum of -0.3 or lower as the MIT minimum, which they label with T_M . They label the poleward, T_P , and equatorward, T_E , walls of the trough at the latitudes closest to the minimum where the detrended-log-electron-density returns to zero. If more than one negative peak within the segment achieve a minimum which passes the threshold, then the equatorward one is selected. Examples of identified troughs are shown in figure 3-4 which was copied from (Aa et al., 2020), as well as in figure 3-5, where they are shown with green lines and red lines in the bottom and bottom left plots respectively.




Figure 3-4: Copied from figure 1 in (Aa et al., 2020). Examples of the trough identification method for SWARM data. The dotted line in the "Ne" plots shows the background electron density.

With some parameter adjustment, this algorithm could identify the MIT in any satellite in-situ electron density measurements. Because it does not require any hand-labeling, this process can be used for very large datasets. Another advantage is that it is simple and so its failure modes are understandable.

3.1.3 Defense Meteorological Satellite Program - Special Sensor Ultraviolet Spectrographic Imager (DMSP - SSUSI)

The Defense Meteorological Satellite Program (DMSP) has 4 satellites of interest to this project named F16, F17, F18 and F19 which orbit at around 850km (100 minute

period). These satellites were launched between 2003 and 2014 and one of the payloads they carry is an imaging system called Special Sensor Ultraviolet Spectrographic Imager (SSUSI) (Johns Hopkins University Applied Physics Laboratory, 2021). One of the data products that the SSUSI team provides is an estimate of the boundaries of the auroral oval. This estimate is the result of fitting a model to measurements of auroral precipitation. The SSUSI team provides the latitudes of the auroral boundary at regularly spaced geographic longitudes once per orbit. We convert the geographic latitudes to magnetic apex latitudes, then linearly interpolate them onto the MLT coordinates that we use for TEC and the MIT labels. We save all the boundary estimates into monthly files, then when we open the files, we interpolate the boundary in time to match the TEC map times. This product is useful because the poleward wall of the trough is usually found just equatorward of the auroral boundary and so we use the boundary to guide our trough identification algorithm. This will be explained in greater detail in section 3.3.

3.1.4 Geophysical Indices and Solar Wind

The primary science goal of this project was to clarify how the MIT responds to various ionospheric drivers and conditions. One way to do this is to sort the trough into bins of a geophysical index. The purpose of any geophysical index is to quantify some complex aspect of the ionospheric state with a single scalar value. One example is the planetary K index, more commonly called Kp , which is meant to be a measure of the global level of geomagnetic activity. Every three hours, local geomagnetic disturbances are measured at 13 subauroral location, and these measurements are standardized and averaged to obtain the global estimate (GFZ Helmholtz Center Potsdam, 2021).

While Kp provides a representation of the current state of the ionosphere, it does not directly measure ionospheric drivers. For that, we use measurements of the solar

wind. The Goddard Space Flight Center Space Physics Data Facility provides a low resolution version of these measurements which consist of hourly averages. An example of the solar wind magnetic field measurements is shown in the top left plot of figure 3.5. Missing data is common in any space physics data source and the solar wind measurements are no exception. These averages are made up of measurements from several different satellites which have been time-aligned to account for the different satellite locations. The time alignment procedure is explained in detail in (NASA Goddard Space Flight Center Space Physics Data Facility, 2021). When comparing the solar wind data to the trough, we add an additional delay of 1 hour to improve correlation. This is because there is some delay between a change in the solar wind and the corresponding change in the ionospheric convection pattern, and there is a further delay between a change in the convection pattern and the corresponding change to the MIT.

3.1.5 SuperDarn

The Super Dual Auroral Radar Network (SuperDARN) is made up of over 30 high-frequency phased-array radars distributed in both hemispheres. Each radar measures the backscatter from decameter-scale irregularities in the ionosphere. These irregularities drift with the convection pattern, and by estimating the doppler shift of the scattered wave, the radar is able to determine the LOS component of the $\mathbf{E} \times \mathbf{B}$ plasma drift. Radars are operated in pairs so that multiple velocity components are measured. This strategy allows for the drift velocity to be fully determined. A SuperDARN radar only receives a backscattered signal if the irregularities are present and the radar wave encounters them from a direction orthogonal to the magnetic field. Though the radars operate continuously, these conditions are met only a portion of the time (Ruohoniemi and Baker, 1998). Using the individual velocity measurements, the SuperDARN team fits a spherical harmonics model for the electric potential. An

example is shown in the top middle plot of figure 3.5.

We were provided 10 years of 2-minute SuperDARN convection measurements from 2010 to 2020. The maps are organized on an equal area grid based on altitude adjusted corrected geomagnetic coordinates (AACGM). To make the data more easily comparable to our dataset, we binned and averaged the 2-minute values into the 1 hour-1 degree latitude-2 degrees longitude grid that our labels and TEC data is on. Finally, for each pixel of the processed SuperDARN maps, we classified it as SAPS or not using the following procedure. We initially label a pixel as SAPS if the westward component of its velocity is above a threshold. For the threshold, we tried 300 m/s and 400 m/s. Then we expanded the positively labeled pixels using a binary dilation operation. Finally, we removed positive labels from within the auroral oval. This is discussed in greater detail in section 4.3.

3.2 Problem Setup

In the computer vision community, the task of labeling each pixel of an input image is called image segmentation. In this setting, the input is set of images, $X \in \mathcal{R}^{m \times n}$ where m is the size of the image and n is the size of the dataset, and our goal is to estimate a set of binary labels, $\hat{Y} \in \{0, 1\}^{m \times n}$ which has positive values corresponding to the trough. Note that X and $\hat{Y}$ are the same size which means that every pixel in the input is assigned a label. In the supervised learning setting, we would also have a set of ground truth labels $Y_{GT} \in \{0, 1\}^{m \times n}$. Typically, one would choose a class of models controlled by a set of parameters Θ and a loss function $L(X, Y_{GT}, \Theta)$ such that L can be minimized with respect to Θ .

Our setting is slightly different in that we do not have proper ground truth labels. The MIT is not a precisely defined phenomenon and so in one sense, ground truth labels do not exist. Ideally, with a better understanding of the MIT in the future,




Figure 3-5: Overview of the data used in this project. Top left is a plot of solar wind measurements. The solar wind influences the ionospheric convection pattern shown top middle. The convection pattern carries plasma around the ionosphere and is largely responsible for the TEC distribution, shown on the right. Under the TEC plot and to the left are trough labels identified in TEC and SWARM data.

we could come up with a precise definition which distinguishes it from other trough-like phenomena in the ionosphere. Although we could have hand-labeled a significant number of TEC images and used that as training data, our labels would be subjective and may be influenced by our biases. Instead, we decided to replicate the algorithm used in (Aa et al., 2020), which identifies the MIT in measurements of electron density by the SWARM satellite constellation (described in section 3.1.2). However, this approach also introduces complications because the SWARM labels only provide a sample of the MIT at a six MLTs (one sample on each side of the north pole per satellite), whereas our labels sample the MIT at all MLTs. Mathematically, you

could think of this as partially observing a noisy version of Y_{GT} .

$$Y_{SWARM} = h(Y_{GT} + \eta) \quad (3.1)$$

Equation 3.1 models using the SWARM labels as ground truth. Y_{SWARM} are the SWARM labels from Aa 2020, Y_{GT} are the (theoretical) ground truth labels, η is a noise process which corrupts the true labels into the labels determined by the Aa 2020 algorithm, and h is an observation function which samples the noisy labels at particular MLTs. Unfortunately, since the performance of the Aa 2020 algorithm is probably affected by the shape, size and background conditions of the MIT, the noise process η is not likely to be independent from Y_{GT} . The hope in any study like this is that the noise process is almost independent from the ground truth labels and so sample statistics are close to their true values. What this means for our study is that the Aa 2020 labels provide a good check to see whether our algorithm is working properly, but we should not be overly concerned with fitting their results as closely as possible.

3.3 MIT Labeling Methods

During this thesis, we developed two methods for identifying the MIT in TEC images. Our approach is very much inspired by previous work by others: (Yang et al., 2015), (Ishida et al., 2014), (Aa et al., 2020), (Pryse et al., 2006), (Voiculescu et al., 2006). First we (1) perform a preprocessing step to the TEC image, then we (2) assign a score to each pixel, then we (3) threshold the scores and (4) do a postprocessing step. The scoring step is different for the two methods we developed, and each has advantages and disadvantages. We found that for two-dimensional data, the preprocessing by itself was insufficient to separate the trough from the rest of the ionosphere which is why we add additional processing before and after thresholding.

While two dimensional data may be more complicated to process, we believe that it should lead to a more accurate determination of the MIT because of the longitudinal coherence of the trough. 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 gives us higher confidence that it is indeed part of the MIT.


While developing these methods, we made an effort to keep our processing as consistent with previous methods as possible. This is important because the MIT is not a precisely defined phenomenon. In fact, the most precise definition that the MIT has is provided by the algorithms which are used to identify it. By making our methods similar to previously developed ones, we prevent our implicit definition of the trough from being too far away from what is accepted.

3.3.1 General Considerations


Before we explain the details specific to each method, we will describe the processing steps that are utilized in both methods. For the remainder of this section, the input will be called the *TEC vector* ($\mathbf{T}$), the output of the preprocessing step will be called the *preprocessed vector* ($\mathbf{x}$), the output of the scoring step will be called the *score vector* ($\mathbf{u}$), the output of the thresholding step will be called the *thresholded vector* ($\mathbf{q}$) and the final output after postprocessing will be called the *label vector* ($\hat{\mathbf{y}}$). Examples of the input (TEC) and output (labels) can be seen in the top right and bottom left plots of figure 3-5 respectively.

Preprocessing


The input to this stage are the TEC vectors, and the output are the preprocessed vectors, both of which can be thought of as either (60×180) arrays or $(10,800 \times 1)$ vectors. We have shown an example of the input and output from this stage in figure 3-6a. Each grid cell in the input holds the TEC in TECu for that location. We begin




(a) Input (left) and output (right) to the preprocessing step: TEC and pre-processed vector, respectively



(b) Distributions of TEC (left) and Log TEC (right)



(c) Distributions of preprocessed values with different filter sizes



(d) Examples of preprocessed vectors calculated with different filter sizes. Left: (9 x 9), right: (21 x 21)

Figure 3.6: Preprocessing

preprocessing by throwing out any out-of-bounds data which we decided to be values below zero or above 150 TECu.

We then converted TEC to log-TEC. By converting to log-TEC we are placing importance on the relative decrease in TEC in the trough as opposed to absolute decrease. This means a dip from 2 TEC to 1 TEC might be considered a trough but a dip from 15 to 10 might not. This could exaggerate the occurrence rate of the trough during the winter when the TEC values are lower in general, but this isn't necessarily a problem because the trough doesn't have a true definition. Using the log scale more closely aligns our definition of the MIT with previous work. The distribution of log-TEC is more symmetric than the distribution of TEC which is generally beneficial for analysis and machine learning. This can be seen in figure 3-6b where we have plotted histograms of TEC and log-TEC from random sample of the TEC dataset.

Next, we estimate the background using a sliding window average. We experimented with different window sizes, but anything around (17×17) seems to work. The sliding window size used in (Aa et al., 2020) corresponds to about 17° of latitude. By padding the log-TEC at the top and bottom with the edge value and at the left and right with the values at the opposite side of the array, we avoid MLT edge effects. Finally, following (Aa et al., 2020), we subtract the background from the log-TEC image. This step is equivalent to high pass filtering with the cutoff frequency controlled by the size of the sliding window. A larger window sets a lower cutoff frequency. We have shown examples of the preprocessing step performed with different filter sizes in figure 3-6d and the resulting output distributions in figure 3-6c. The larger filter size results in an output distribution with more variance, which means more of the input power has been let through. This can also be seen in the examples in figure 3-6d where the plot on the right appears to have higher amplitude.

By viewing this step as a high pass filter, we see that its purpose is to filter out variations in the data that aren't helpful for identifying the MIT. The highest amplitude of those variations will be low-frequency including seasonal TEC variations and day-night variations. This can be seen in figure 3-6a where in the left TEC plot, there is a large difference between the TEC on the day side versus on the night side, but there is no such difference in the right plot of the preprocessed vector. There is some optimal value for the cutoff frequency which we don't know. If the cutoff frequency is too low, then low frequency variations like day-night are allowed to pass and if the cutoff is too high, then large troughs will be filtered out. We estimated the moving average filter size as part of our random parameter search. 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 degree wide in latitude and so additional low pass filtering is not needed.

At this point in past studies, the authors would set a threshold to identify MIT regions. The reason we could not threshold the preprocessed vectors directly is because there are other trough-like structures at high and low latitudes that would be misclassified as the MIT. Additionally, the preprocessed vectors contain noise which we would prefer to get rid of. Both the noise and the non-MIT trough-like structures can be seen in figures 3-6a and 3-6d.

Thresholding

Many binary labeling methods rely on estimating some continuous score value, then thresholding it to make a final decision. By varying the threshold and measuring true positive rate and false positive rate, you can trace out the receiver operating characteristic (ROC) curve. Therefore to choose our threshold, we wanted to maximize

accuracy and minimize the poleward wall and equatorward wall errors. We could have chosen a cost function which balances these metrics, performed a grid search, and picked the threshold with the lowest cost, but this has a problem. Varying the threshold also affects the size of the trough labels, i.e. if the threshold is lower, more pixels will be considered trough and the resulting output will have larger / more troughs. The continuous errors are only measured where both methods say that a trough is present. For this reason, as the threshold increases, fewer and fewer of the continuous errors are being considered, and the ones that are considered are only for very high confidence troughs, i.e. where the score value is above a high threshold. This causes the poleward and equatorward absolute errors and error variances to decrease as the threshold increases, which means they are not a good indicator of a properly set threshold. An example of this is shown in figure 3-7. Both the top and bottom plots show performance metrics over a range of threshold values. In the bottom plot you can see, after an initial increase, both error standard deviation lines decrease with increasing threshold. Instead, we used accuracy to determine the threshold. In the top plot of figure 3-7, you can see that the red accuracy line exhibits a maximum. During our parameter search, we looked for settings which resulted in low bias and variance in the continuous errors even though the threshold was not explicitly chosen to minimize them.

Prior MIT Model

One of the key ways in which we rejected trough-like structures at high and low latitudes was by performing a sort of weighted regularization based on the expected location of the MIT. We experimented with two models for expected trough position: the empirical model developed in (Deminov and Shubin, 2018), and one based on the auroral boundary measured by SSUSI. SSUSI is described in section 3.1.3. Examples of both prior models can be seen in the top right plot of figure 3-5 as red and black




Figure 3-7: *TLM1* performance metrics at different thresholds. Top: Accuracy (Acc.), true positive rate (TPR), true negative rate (TNR), false positive rate (FPR), false negative rate (FNR). Bottom: Poleward ("P") and equatorward ("E") wall latitude error mean and standard deviation ("mean", "std. dev.")

dashed lines. From these models, we got a vector of MLats for each time step corresponding to our MLT coordinates. We then computed weights for each cell of our TEC grid using the following formula:

$$w(\lambda, \phi) = c|\lambda - m(\phi)|^p + 1 \quad (3.2)$$

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 , c is a scalar which sets the maximum weight and p sets the order of weighting to either linear or quadratic. Using this sort of weighted regularization helped prevent our methods from scoring pixels far from the expected trough location too highly.

3.3.2 Trough Labeling Method 1 (TLM1)

Motivation and Overview

As mentioned earlier, the problem with thresholding the preprocessed vector is that any relative low region will be labeled as trough such as isolated noisy pixels or low regions within the auroral oval or at low latitudes. We wanted to threshold a vector which more closely approximated our level of confidence in each pixel being part of the MIT. We called this vector the score vector. The idea behind *TLM1* is to model the score vector's contribution to the preprocessed vector (forward model), then invert the model to find a score vector given a preprocessed vector. An example of a preprocessed vector is shown in figure 3-8a and its corresponding score vector is shown in 3-8b. In both plots, red indicates a high value and blue indicates a low value so we should expect blue regions in the preprocessed vector to be shaded blue in the score vector. Note how the preprocessed vector has some non-MIT low regions both higher and lower latitude than we would expect the trough, e.g. near MLT 22 and MLT 6. Thresholding the preprocessed vector directly would result in these regions being mislabeled. In the score vector these regions are not scored as highly as the pixels within the MIT.

Trough Image Model

We modeled each preprocessed vector as a linear combination of radial basis functions (RBFs) with the weights given by the score vector values, plus noise. In general an RBF is any function which only depends on the radial distance from a point $\mathbf{c}$ (Bishop, 2006).

$$s(\mathbf{x}) = h(\|\mathbf{x} - \mathbf{c}\|) \quad (3.3)$$

The most common choice for RBFs is the Gaussian RBF:

$$s(\mathbf{x}) = \exp\left(-\frac{\|\mathbf{x} - \mathbf{c}\|^2}{2\sigma^2}\right) \quad (3.4)$$

where $\mathbf{c}$ is the center and σ is a size parameter which is sometimes called the bandwidth. For the remainder of this thesis, "RBFs" will refer to Gaussian RBFs unless otherwise noted. Our forward model for the preprocessed vectors $\mathbf{x}_i$ based on the score vectors $\mathbf{u}_i$ has the following form:

$$\mathbf{x}_i = -A_i \mathbf{u}_i + \epsilon_i \quad (3.5)$$

where ϵ is noise. Each column of the matrix A_i contains an RBF centered on a pixel of the grid. The negative sign means that higher score values will result in lower values in the preprocessed vector. Nominally, all the A_i 's are identical, but because 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 x . If $\mathcal{G}_i$ is the set of indices where x_i has data, then the size of A_i is $(|\mathcal{G}_i| \times 10800)$. By using RBFs, we cause the score vector elements to affect all the elements in the corresponding neighborhood of $\mathbf{x}$. We were motivated to choose a forward model with this property by the fact that the MIT is a large scale structure, i.e. a single low pixel in the preprocessed vector does not constitute the trough, only when many contiguous pixels are low should the region be considered part of the MIT.

This processing step is meant to capture three aspects of / assumptions about the MIT, as it appears in TEC images. The first, which we described above, is that the trough is a large-scale structure which should cause regions of low TEC. The second is that the MIT should appear where we expect it to, i.e. not too far from a prior model like the ones described in section 3.3.1. Lastly, the trough should mainly appear as one large contiguous region of low TEC, not as many smaller regions scattered

around. The first aspect is part of the forward model via the RBF basis, but the second two aspects are implemented in the inversion using regularizers. During the inversion, we search for a score vector which fits the forward model, but constrain the search to score vectors which also satisfy our assumptions about the MIT. We express the trough image model inversion mathematically as:

$$\mathbf{u}_i^* = \arg \min_{\mathbf{u}_i} \mathbf{x}_i^T A_i \mathbf{u}_i + \alpha \|W_i \mathbf{u}_i\|_2^2 + \beta \|C \mathbf{u}_i\|_1 \quad (3.6)$$

where $\mathbf{u}_i^*$ is the fitted score vector, W_i is a diagonal matrix which implements the weighting described in equation 3.2 and C is a matrix which computes the differences between all neighboring pixels of $\mathbf{u}$. α and β are nonnegative coefficients which weight the relative importance of the three components. The cost function is convex because each individual term is convex: the first term is linear which makes it convex and the second two terms are convex given that α and β are nonnegative. In fact, for $\alpha > 0$, the second term makes the cost function strongly convex, guaranteeing that it has a single unique minimum. We performed the minimization using the commercial software *Gurobi* (<https://www.gurobi.com/>) but Python also has open source optimization problem solvers which we could have used, e.g. *CVXPY* (<https://www.cvxpy.org/>).

Typically in a problem like this, the first term would be $\|\mathbf{x}_i + A_i \mathbf{u}_i\|^2$ in order to minimize the euclidean distance between $\mathbf{x}_i$ and $-A_i \mathbf{u}_i$. We instead minimized the negative dot product between the forward model and the preprocessed vector. Minimizing this without the other terms would send $\mathbf{u}_i$ to infinity wherever $\mathbf{x}_i$ is negative. The regularization terms prevent this from happening. We found that setting up the cost function in this way produces sharper score vector boundaries. One intuitive reason for this is that a linear cost term is less restrictive than a euclidean distance cost term, which results in a greater influence from the regularization terms

on $\mathbf{u}_i^*$. For example, the level set of a euclidean distance term is a sphere, which is finite, whereas the level set of a linear cost function is a plane which is infinite. Another reason is that in the transition region between trough minimum and wall, where $\mathbf{x}$ becomes close to zero, the first term has essentially no influence on those pixels. Ultimately, this form of cost function strengthens the regularizers.

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 where we expect the trough to be.


Together with the $L1$ norm, the C matrix implements the total variation (TV) regularizer.

$$(C\mathbf{u}_i)_j = \sum_{k \in \mathcal{N}_j} u_{ij} - u_{ik} \quad (3.7)$$

The j th element of $C\mathbf{u}_i$ is the sum of the differences between the j th element of $\mathbf{u}_i$ and all of its neighbors. 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 this tends to make it sparse, which means the gradients of the score vector will be sparse. This encourages the score vector to have larger contiguous regions of the same value and is minimized when every pixel is the median of its neighbors. The goal of using this is to get score vectors which are less influenced by noise and missing values, instead tending to have one or two larger contiguous trough regions.


Post-processing

The purpose of postprocessing is to clean up specific errors which remain after thresholding the score vector. These errors include small patches classified as MIT due to noise in the preprocessed vector and trough-classified pixels within the auroral oval. Examples of the input and output to the postprocessing step are shown in figures 3-8c and 3-8d respectively. The corresponding auroral oval estimate is shown as



(a) Input to the *TLM1* scoring step: the pre-processed vector. Blue color indicates local plasma depletion

(b) *TLM1* score vector. Red color indicates high score.



(c) Thresholded vector. Still has small patches and patches within the auroral oval (dotted line) which we want to remove.

(d) Final labeled vector

Figure 3-8: *TLM1* intermediate vectors

a dotted black line in figure 3-8c. For removing small patches, we set thresholds for the perimeter and area of each patch, then utilized connected component analysis to determine whether each patch passed (<https://scikit-image.org/docs/stable/api/skimage.measure.html#skimage.measure.label>). We considered the latter scenario an error because the community generally accepts that the auroral precipitation makes up the poleward wall of the trough (Karpachev et al., 1995). We removed positive pixels from within the auroral oval because according to many other papers, electron depletions within the auroral oval are “high latitude troughs” and not the MIT. This error is common because, due to the grid projection, the pixels in the polar cap are very small. This means that a relatively small-scale disturbance in the ionosphere can affect many pixels near the polar cap, which has a disproportionate influence on the cost function. Finally, the output of this step is our label vector $\hat{\mathbf{y}}$.

3.3.3 Trough Labeling Method 2 (TLM2)

While *TLM1* performs well, its two main drawbacks are that it is complicated and that it is slow. It has many parameters which require manual tuning and to run *TLM1* on a 10 year dataset takes a few hours or overnight. Other researchers are more likely to utilize the method if they are able to run it quickly and it is simple enough for them to adapt it to their specific needs. The time it takes to run *TLM1* is dominated by optimizing equation 3.6. The goal of *TLM2* is to approximate the scoring step of *TLM1* without having to perform any optimization.

If you take the preprocessed vector and add the prior model vector (diagonal of W_i from equation 3.6), then when you threshold the result, the pixels far away from the prior model latitudes are less likely to have a positive label. In this way, we have approximated the $L2$ regularization term from *TLM1*. If we could do the same for the other two aspects of *TLM1*, then we would have a simpler algorithm with similar performance and less computational cost. *TLM2* uses the same preprocessing and




Figure 3-9: Overview of *TLM2* scoring step (see "*TLM2* Overview Diagram Description" in text for details).

postprocessing steps as *TLM1*, the only difference is in the scoring step.

Scoring Step

The score vector is a linear combination of three vector components, each meant to mimic a component of equation 3.6. The first component is just the preprocessed vector $\mathbf{x}$, the second component is the prior model vector $\mathbf{w}$, and the final component is a modified version of the preprocessed vector called the *integrated sparse gradients vector* $\mathbf{v}$.

$$\mathbf{u}_i = \mathbf{x}_i + \alpha \mathbf{w}_i + \beta \mathbf{v}_i \quad (3.8)$$

Equation 3.8 shows the computation of the *TLM2* score vector.

The sparse gradients vector is meant to approximate the TV regularizer from *TLM1* and is the most complicated part of *TLM2*. The TV regularizer causes score vectors with more dispersed values to cost more, i.e. it favors score vectors with a single contiguous region. While the TV regularizer "tends to" result in solutions with sparser gradients, to approximate it in *TLM2*, we directly applied operations to the preprocessed vector to sparsify its gradients. We processed the x and y gradients separately and combined them at the end. For each gradient vector, we first applied a sliding median filter to reduce noise, as the initial gradient calculation can result in a noisy vector. Then we applied a "sliding max filter" to create a sparse vector. The max filter, which uses a single-pixel-width window, sets values that are the maximum within the filter window to the sum of the values within the filter window, otherwise it sets them to zero. This has the effect of compressing a neighborhood of gradients into a single pixel. If we were to integrate the gradients, instead of seeing a ramp up, we would just see a single jump. We then spread out the sparsified gradients in their orthogonal direction using a sliding mean filter. Finally, we symmetrically integrated each gradient by adding the forward integration to the negative of the backwards integration. Adding the two integrated gradients together gave us our "integrated sparse gradients vector". The integration process is shown below in equation 3.9, where $v(r, s)$ is the total integrated sparse gradient at the pixel in the r th column and s th row, $v_x(r, s)$ and $v_y(r, s)$ are the separately integrated sparse x and y gradients,

and $\nabla_x v(r, s)$ and $\nabla_y v(r, s)$ are the sparsified gradients in the x and y directions.

$$\begin{aligned}
 v(r, s) &= v_x(r, s) + v_y(r, s) \\
 v_x(r, s) &= \sum_{i=1}^r \nabla_x v(i, j) + \sum_{i=1}^{R-r} -\nabla_x v(R-i, j) \\
 v_y(r, s) &= \sum_{j=1}^s \nabla_y v(i, j) + \sum_{j=1}^{S-s} -\nabla_y v(i, S-j)
 \end{aligned} \tag{3.9}$$

***TLM2* Overview Diagram Description**

Figure 3.9 shows an overview of the *TLM2* scoring step. The bottom left image is the preprocessed image $\mathbf{x}_i$. The sparsified x and y gradients are labeled with Δ_x and Δ_y respectively and the sparsified y gradient image is shown above and to the left of the $\mathbf{x}_i$ image. These two sparse gradient components are integrated and summed to produce the "sparse gradients vector" $\mathbf{v}_i$ which is shown in the top image. Along with the prior model vector $\mathbf{w}_i$, which is shown in the bottom right image, the preprocessed vector and the sparse gradients vector are summed together to produce the score vector $\mathbf{u}_i$. As the final operation of the scoring step, this sum is shown with a bold green arrow and the output score vector is shown in the right image.

3.3.4 Performance

As mentioned earlier, we used several error metrics to measure the performance of our algorithms. Since the Aa 2020 labels are not actually a source of ground truth, we were not expecting or even hoping to match them perfectly. The key characteristics we wanted to observe were low bias and decent binary accuracy.

To test an algorithm, we randomly selected N days from the 7 year span of SWARM data we downloaded and ran the algorithms on all of the data from those days. The Aa 2020 algorithm splits the SWARM orbits into segments between 45° and 75° MLat, through which each satellite passes twice in the northern hemisphere




Figure 3-10: (a) SWARM satellite MLAT vs time, showing time alignment of satellite orbital segments with TEC maps. Black vertical lines indicate one TEC map, blue and yellow lines mark selected segments. (b) SWARM trough vs TEC trough comparison diagram.

per orbit. The segments in which the satellites' latitude is increasing, we called "up" segments and the other segments we called "down". Each of our label vectors correspond to one hour and for comparison, we chose the "up" segment and the "down" segment which were closest in time to each one. This is illustrated in figure 3-10a where the dashed black line indicates a SWARM satellite's MLat over time, the vertical black lines indicate the time span of a TEC map and the blue and yellow lines indicate the selected SWARM orbital segments.

We then extracted a 3-pixel-wide path of our label vectors under the selected SWARM satellite orbital segments and marked the highest and lowest MLats where our labels were positive, corresponding to the poleward and equatorward walls of the MIT, respectively. We compared these latitudes to ones from the Aa 2020 algorithm. In the case of multiple trough candidates within a single orbital segment, the Aa 2020 algorithm, as written in the paper, chooses the one with the lowest MLat. For our comparison, we instead chose the one which best agreed with our labels. This should result in better ground truth labels because rather than somewhat arbitrarily choosing the lower MLat troughs, we chose troughs that are confirmed in two separate datasets. If our algorithm did not detect the trough, then we used the

Error Type	TEC	SWARM
No Error: True Negative	No Detection	No Detection
Error: False Negative	No Detection	Trough Detected
Error: False Positive	Trough Detected	No Detection
No Error: True Positive	Trough Detected	Trough Detected

Table 3.1: Binary error types

original strategy of choosing the lowest latitude trough.

Aligning the Aa 2020 labels with ours nominally resulted in 6 comparisons per TEC map (3 satellites, 2 orbital segments, but often less due to missing SWARM) and 6 variables per comparison. There are two binary variables, one for SWARM and one for TEC, indicating whether any trough was 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. Of course, if no trough is detected in one of the data sources, then there are no values for these latitudes and so we did not perform any comparison. With the two binary variables, there four possible situations: true negative, false negative, false positive and true positive. These are listed in table 3.1. In the true positive case, when a trough is detected by both our algorithm and the Aa 2020 algorithm, then we computed the errors of the wall latitudes as:

$$\begin{aligned}
 E_P &= \lambda_{TP} - \lambda_{SP} \\
 E_E &= \lambda_{TE} - \lambda_{SE}
 \end{aligned}
 \tag{3.10}$$

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. With 6 comparisons per TEC map and 24 TEC maps per day, we got a total of $144N$ comparisons, from which we computed accuracy, rates for the binary error types, and statistics (mean / standard deviation) for the continuous errors.

To find good parameters for our algorithms, we utilized a random search. Because

Parameter	Setting
Background Filter Size	(19, 17)
RBF Bandwidth (distance to half power)	1 pixel
L2 Regularization Strength α	.05
Total Variation Regularization Strength β	.15
Horizontal TV strength	Double
Perimeter Threshold	40 pixels
Area Threshold	40 pixels
Prior model	(Deminov and Shubin, 2018)
L2 weighting order (p in eq. 3.2)	1
Max weight (c in eq. 3.2)	15
Threshold	0.7

Table 3.2: Parameter settings for *TLM1*, found with random search

we do not have actual ground truth labels, we did not simply pick the parameters with best metrics, but rather used the metrics and manual debugging to guide us towards parameters that seemed subjectively reasonable and resulted in good performance. The parameters which we settled on are listed in tables 3.2 and 3.3 for *TLM1* and *TLM2* respectively.

Baselines

It is always a good idea to test out new methods against a simple baseline. Because there is no other method for MIT identification in TEC maps, we used a 2D analog of the Aa 2020 algorithm. We used the same preprocessing step as described in section 3.3.1, which is already a 2D version of the Aa 2020 algorithm. We applied a small (5×5) median filter to smooth out the preprocessed image, then we identified possible troughs by searching for closed < -0.05 contours. We eliminated any contours which did not contain any pixel values < -0.2 . Finally, we labeled all pixels within the remaining contours as trough and applied the postprocessing operations described in

Parameter	Setting
Background Filter Size	(19, 17)
"Prior" weight α	1
Max weight (c in eq. 3.2)	1
"Prior" order (p in eq. 3.2)	1
Prior model	auroral boundary - 1
Integrated sparse gradients weight β	.5
Perimeter Threshold	40 pixels
Area Threshold	40 pixels
x "max filter" size	3
y "max filter" size	5
median filter size	7
mean filter size	15
Threshold	0.0

Table 3.3: Parameter settings for *TLM2*, found with random search

section 3.3.2. Though not identical to the Aa 2020 algorithm, this is fairly close 2D analog with a few modifications to improve performance. Importantly, this baseline algorithm has very few parameters and represents a reasonable first attempt at this task.

To put the following performance results into perspective, the Aa 2020 dataset has positive labels in about 60% of cases, which means an algorithm that guessed positive for every timestep and MLT would have around 60% accuracy. To add additional context to our results, we have also included a baseline which labels every vector with a 6° -wide trough centered on the (Deminov and Shubin, 2018) empirical model. The two dimensional generalization of the Aa 2020 algorithm will be referred to as *B1* and the positive-guesser based on the Deminov 2018 empirical model will be referred to as *B2*.

Results

Table 3.4 summarizes the performance of *TLM1*, *TLM2* and the baselines. For the metrics in table 3.4, we considered all longitudes. Because of the uneven coverage of the TEC maps, we also computed performance metrics ignoring all SWARM - TEC comparisons in the range [130, 260] MLon, which are listed in table 3.5. We roughly estimated this range by looking at figure 3.3. In general, the metrics listed in table 3.5 show improvement over those in table 3.4. The "Time" metric is the time in minutes it took to run the algorithm on 200 days (4800 TEC maps). While we do list "time" as a metric in the tables, our focus is mainly on accuracy, bias and variance because we only need to run the algorithm once through the TEC dataset. In figure 3.11, we show the continuous error distributions for the three algorithms. These distributions are shown for the higher coverage longitude range.

B1 performs well in terms of binary accuracy, agreeing with Aa 2020 in 80% of cases. However, it has a rather high error standard deviation for its estimates of the MIT walls at 4.28 and 5.06 for the poleward and equatorward walls respectively. From figure 3.11c we see that the equatorward error distribution has an especially heavy negative tail. The postprocessing operations most likely prevent the poleward wall estimate from being too far off, but the equatorward wall estimate is less constrained. In general, the equatorward wall of the MIT is less well-defined than the poleward wall, and other authors have found it difficult to parameterize (Prölss, 2007).

TLM1 appears to have the best accuracy. Especially appealing is the fact that it performs the best when we include the poorly covered longitudes. *TLM1* should be able to handle missing data more gracefully because it is simply masked out in the scoring step. In areas with missing data, the cost function is only determined by the L2 and TV regularizers. *TLM1*'s estimate of the poleward wall has lower bias than *TLM2* and somewhat higher standard deviation. *TLM1*'s estimate of

Metric	TLM1	TLM2	B1	B2
Accuracy	0.81	0.75	0.74	0.59
True Positive Rate	0.76	0.78	0.62	1.0
True Negative Rate	0.87	0.67	0.92	0.0
False Negative Rate	0.24	0.22	0.38	0.0
False Positive Rate	0.13	0.33	0.08	1.0
Poleward Error Mean	-0.04	1.02	-1.41	-0.47
Poleward Error Std. Dev.	4.12	3.79	5.80	3.91
Equatorward Error Mean	-1.74	-0.41	-2.85	-0.92
Equatorward Error Std. Dev.	4.08	4.01	6.45	4.69
Time (minutes)	108	4.14	3.50	

Table 3.4: Performance comparison at all longitudes. *B1*: 2-D Aa 2020; *B2*: Deminov 2018

the equatorward wall has higher bias than *TLM2* but they have the same standard deviation. Both overestimate the width of the MIT on average but *TLM1* by slightly less. The lower standard deviation for *TLM2* indicates that it may be more strongly controlled by the prior model. The one place where *TLM2* is a clear winner is in speed, which is what it was designed for. *TLM1* behaves more predictably in the presence of missing data, so it is what we ended up using for this particular study. However, with a larger dataset of higher resolution samples, the speed of *TLM1* would be too low to be feasible. Both algorithms outperform the baselines and have good enough agreement with Aa 2020 to give us confidence in our results, and they would both be useful in different scenarios.




Figure 3.11: Poleward and equatorward wall latitude error distributions for the three methods

Metric	TLM1	TLM2	B1
Accuracy	0.84	0.80	0.80
True Positive Rate	0.83	0.91	0.70
True Negative Rate	0.86	0.63	0.93
False Negative Rate	0.17	0.09	0.30
False Positive Rate	0.14	0.37	0.07
Poleward Error Mean	0.37	1.41	-0.15
Poleward Error Std. Dev.	3.88	3.62	4.28
Equatorward Error Mean	-1.49	-0.63	-1.62
Equatorward Error Std. Dev.	3.90	3.90	5.06

Table 3.5: Performance comparison excluding magnetic longitudes [130, 260]. *B1*: 2-D Aa 2020

Chapter 4

Experiments

In this chapter we demonstrate some of the scientific applications of our labeled dataset. We explain how our flexible labeling scheme allowed us to compute a wide variety of measurements and statistics of the MIT. In order to avoid quantization artifacts in our histograms, we applied a small dither to some variables, e.g. Kp . The dither is sampled independently and identically distributed (IID) from a zero-mean Gaussian distribution and we set the variance as small as possible to make the histograms reasonably smooth. The variables we applied a dither to and the corresponding Gaussian standard deviations are listed in table 4.1. Unless otherwise noted, the dithered versions of the variables in table 4.1 were used throughout this section.

Variable	Dither Std. Dev.
Kp	0.1
MLT	0.02 hours
MLat	0.01 degrees
B_y	0.05 nanotesla
B_z	0.05 nanotesla
Season	0.25 days

Table 4.1: Gaussian dithering standard deviations

4.1 Replication

In order to provide further evidence that our labeled dataset was valid, we first verified that previously known statistical relationships of the MIT were reproduced in our dataset. We have provided an explanation of these effects as well as general background on the MIT in section 2.2.

4.1.1 MLat vs Kp

That the MIT’s latitudinal position decreases with increasing Kp is probably the most well documented aspect of the MIT’s behaviors. Therefore, this was the first result we wished to replicate with our dataset. Because our dataset contains an order of magnitude more measurements than previous studies (section 1.1), we were able to view these statistical relationships with a much greater level of detail.

Figure 4-1a is a copy of figure 8b from (Aa et al., 2020) which shows the linear regression results of Kp versus trough minimum position from eight different papers, including their own (black solid line). In figure 4-1c we show our computation of the same line. Because our labels are low level, i.e. pixel level, it takes a small amount of additional computation to estimate the latitude of the MIT minimum. To do this, we used our labels to mask out non-trough pixels of either the TEC vectors or preprocessed vectors, then we searched for the latitude which achieves the minimum value at each MLT. This gives us an array with shape $(N \times 180)$ where N is the number of TEC maps in our dataset. The array has NaNs at each MLT which has no trough pixels. Since MLT also has a large effect on trough position, in figure 4-1c we limited our regression data to only include MLT values within two hours of midnight. Finally, we binned the MIT positions into 20 Kp bins. Figure 4-1b shows the overall regression line as well as the mean of each bin as a dot and the standard deviation of each bin with error bars. In general, our results agree well with




Figure 4.1: MIT position and occurrence rate dependence on Kp . (a) Figure 8b Aa 2020: MIT position vs Kp in various studies. (b) Figure 8a Aa 2020: MIT occurrence rate vs Kp . MIT position vs Kp in our dataset, within 2 hours (c) and 5 hours (d) of midnight. (e) (left) Distribution of MIT Kp and MLat, (right) occurrence rate of MIT in Kp - MLat bins.

the results from other studies. The Aa 2020 line is slightly higher than most of the others, which they attribute to the fact that they averaged over 19 - 5 MLT whereas other studies only averaged around midnight. The trough before midnight rapidly increases its latitude while after midnight, it decreases its latitude only marginally. We verified the plausibility of this explanation by computing our trend line for the wider range used in (Aa et al., 2020), which indeed caused the line to move poleward. Our midnight line is very close to the line from (Yang et al., 2015) which also was computed from Madrigal TEC maps.

Figure 4.1b, copied from figure 8b of (Aa et al., 2020), shows how the occurrence rate of the MIT varies with Kp . One of the benefits of our large dataset is that we are able to simultaneously show these two relationships with the same histogram, which we have done in figure 4.1e. The left plot of figure 4.1e shows the overall distribution of the MIT's Kp and MLat in our dataset, and the right plot shows the occurrence rate of the MIT in the same bins. To compute these histograms, we checked for any positively labeled pixels within the MLT range $[-1.5, 1.5]$ at each latitude. This resulted in a binary array of shape $(N \times 60)$, of which each column has the same MLat and each row has the same Kp . We then sorted the positive entries of this array into Kp - MLat bins. The left plot of figure 4.1e shows the count of each bin. To compute the occurrence rate shown in figure 4.1e, we divided the count of each bin by the number of times it was observed, i.e. the number of data points, trough or not, we had for that bin. The left side of figure 4.1e shows that the trough moves equatorward with higher Kp and that high Kp occurs less frequently than low Kp . The right side of 4.1e shows the same trough position effect, but also shows that the occurrence rate of the MIT increases with Kp . Finally, it gives a sense of the variability of the trough's position. Overall figure 4.1 confirms the results of previous studies.

4.1.2 MLat vs MLT

Another well-known behavior of the MIT is that its latitude decreases from evening until after midnight. This behavior has been shown in many past studies. Figure 9 from (Yang et al., 2015) is copied in figure 4.2a and it shows the median trough position as well as the upper and lower quartiles of 1 hour MLT bins. They limited their averages to an F10.7 range of [90, 150] and a Kp range of [.3, 2]. A similar plot from (Le et al., 2017) (figure 4) is shown in 4.2b. Theirs was computed from average low- Kp TEC profiles. Both (Yang et al., 2015) and (Le et al., 2017) used the same Madrigal TEC dataset as us. Figure 4.2c shows part of figure 2 from (Aa et al., 2020). All three studies limited their averaging to low- Kp conditions to lessen the variance due to magnetic activity. Our plot, shown in figure 4.2d, was computed in a similar way to figures 4.1c and 4.1d, except we sorted the MIT positions into MLT bins rather than Kp bins.

4.1.3 Occurrence Rate

One of the most interesting contributions from (Aa et al., 2020) was their detailed maps of seasonal MIT occurrence rate, copied in figure 4.3a. 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 their field of view and detail over the earlier maps of (Ishida et al., 2014). The occurrence rates in figures 4.3a and 4.3b are restricted to $Kp \leq 3$. We computed our maps, which are shown in figure 4.3b, by counting the number of times the MIT was observed and dividing by the number of times we had TEC data in each grid cell. We used the same season groups as (Aa et al., 2020) which are November – February for winter; March, April, September, and October for equinoxes; and May – August for summer. No dithering was performed for these plots.




Figure 4.2: MIT Position vs MLT as shown in several studies (a) (Yang et al., 2015); (b) (Le et al., 2017); (c) (Aa et al., 2020), bars indicate poleward and equatorward walls; (d) ours, bars indicate standard deviation

Our plots agree on all of the general patterns. The shapes of the distributions are generally similar. The difference in distribution between the seasons have the same




Figure 4.3: MIT Occurrence rates for $Kp \leq 3$. (a) (Aa et al., 2020), (b) ours

relationship: in the winter the distribution has 2 modes, a small one at around 16-17 MLT and a large one at 3 MLT. In the equinox season, the distribution loses the evening sector mode, and in the summer the trough occurrence rates are the lowest. Because we had more data, we computed the statistics on the same grid that as our labels, which resulted in a more detailed map. The winter bimodal distribution was not explicitly mentioned in (Aa et al., 2020), but both of our plots show it clearly. Because these plots were compiled from low- Kp data, the evening mode is likely caused by an interaction between the sunset terminator and the convection pattern. In all three seasons the trough occurrence rate is diminished in around 21 MLT which is especially interesting because that is where flow stagnation is most likely to take

place. The morning mode of our winter distribution is much wider than in (Aa et al., 2020) and our equinox one seems a bit narrower.

Our trough identification algorithm uses a less strict threshold for classification and so our occurrence rates are considerably higher. The choice of threshold is rather arbitrary, so we do not see this as a problem. It would be more concerning if the general patterns differed significantly. The occurrence rates presented in (Ishida et al., 2014) are even higher than ours.

One of the unique plots of (Aa et al., 2020) was of MIT occurrence rate in season - MLT bins. Figure 4.4a shows the plot copied from their paper (figure 4a) and figure 4.4b shows our version. Both are limited to $Kp \leq 3$. We created our plot (right) by counting the number of times a trough was observed at any MLat in a season - MLT bin and dividing by the number of times there was any data. For the left plot we left the histogram counts unnormalized. The occurrence rates are higher than in figure 4.3 because all MLats were considered for each bin. Our plot agrees well with (Aa et al., 2020). The high occurrence rate region is bounded by the sunset terminator in both plots. There is a subtle decrease in occurrence rate in the premidnight sector around 21 MLT. Again, because we have more observations of the trough, we were able to make our bins smaller and the occurrence rates in our plots appear less noisy.

By reproducing several results of previous statistical studies on the MIT we have accomplished several things. First, we added detail to many of the plots. Additionally, by confirming a relationship with a separate set of measurements from different instruments, we add confidence to the results. Finally, these results serve as an additional piece of evidence that our dataset provides a valid representation of the distribution of the MIT.




Figure 4.4: MIT occurrence rate in season - MLT bins. (a) (Aa et al., 2020), white lines indicate average solar terminator at 60° MLat. (b) MIT season-MLT distribution and occurrence probability on the left and right respectively.




Figure 4-5: Trough occurrence rate during east, south and west IMF clock angle. Top: absolute occurrence rate. Bottom: Occurrence rate differences between top row plots.

4.2 Clock Angle

One aspect of the MIT's behavior that has not been adequately studied is its relationship with the solar wind and interplanetary magnetic field (IMF). As explained in section 2.2, the trough's shape strongly depends on ionospheric convection and precipitation, both of which are influenced by the solar wind. In this section we combine our dataset with measurements of the solar wind and ionospheric convection to investigate how the MIT is affected by the IMF clock angle (equation 2.4).

In figure 4-5 we show the MIT occurrence rate in MLT - MLat bins for different IMF clock angles. For clock angle bins, we used an angular radius of 30 degrees and clock angle centers of 270 for west, 180 for south and 90 for east. We restricted the averaging to time steps where the B field magnitude is between 2 and 6 nanotesla to

isolate the effect of clock angle. B field magnitude has a similar effect on the trough as K_p , i.e. increased B field magnitude expands the auroral oval and convection pattern, pushing the MIT to lower latitude. Another detail is that we delayed the TEC maps by 1 hour from the IMF samples, because this exaggerates the effects. There is some delay between a change in the IMF and the corresponding change in the convection pattern and there is a further unknown delay between a change in the convection pattern and the corresponding change in the trough. A change in the IMF should cause a change in the MIT which develops over a few hours. In the top row of figure 4.5, the difference between the histograms is very subtle. One detectable difference is that for westward IMF conditions, the occurrence rate in the evening sector is higher than for the east IMF condition. Another is that the southward IMF condition produces troughs at lower MLats.

To more clearly show these effects, In the bottom row of figure 4.5 we show the relative change in occurrence rate between the IMF conditions. We created the plots by subtracting the occurrence rate histograms of each pair from the top row. For example, a value of .05 in the left plot means that in a particular MLT - MLat bin, there is a higher occurrence rate during westward B than southward B by 5%, e.g. westward has a 20% occurrence rate and southward has 15%. Both westward and eastward IMF appear to produce the trough at higher MLat. The south IMF produces troughs farther towards noon on the evening side and not as far on the morning side. Additionally, as seen in the top row of figure 4.5, the MIT has a significantly higher occurrence rate in the evening sector during westward IMF than during southward or eastward. One reason that the MIT MLat is higher for eastward and westward IMF could be that MLat decreases with B_z due to expansion of the convection pattern, as seen in figure 2.2.




Figure 4-6: MIT position and occurrence rate dependence on B_z . $2 \leq |B| \leq 6$ nanotesla, MIT labels delayed by 1 hour from IMF measurement

4.2.1 B_z Effect on MIT Position

In figure 4-6, we show the dependence of the MIT minimum on B_z with $|B|$ limited to $[2, 6]$ nanotesla. As before, the trough labels are delayed by one timestep. It was created in the same fashion as 4-1e except with B_z bins. We can see that the MIT's occurrence rate is inversely proportional with B_z and that the MIT's MLat increases with B_z . Figure 4-6b plot shows the MLat of the MIT's minimum as a function of B_z . The dots are averages, the error bars are one standard deviation and the linear regression is plotted.

4.2.2 IMF Effect on MIT Occurrence Regions

To investigate how the IMF clock angle changes the MIT occurrence rate pattern, we looked at the average TEC distribution and average convection pattern during the three IMF configurations. The largest contributors to TEC variance are season and local time. To isolate the IMF effect, we fit a simple sinusoidal model to TEC and subtracted it from the maps. The resulting model is as follows, where T is TEC and t is the day of the year. The subscripts m and c stand for "model" and "corrected"




Figure 4.7: Average TEC profile of evening MIT (16 - 20 MLT) corrected for season

respectively.

$$\begin{aligned}
 T_m(t) &= 0.98 \sin(2\pi t/365) - 2.10 \cos(2\pi t/365) + 9.08 \\
 T_c(t) &= T(t) - T_m(t)
 \end{aligned}
 \tag{4.1}$$

We chose not to remove the MLT variation because we did not want to obfuscate the interaction between solar ionization and the convection pattern.

The left plot of figure 4.7 shows the average evening latitudinal TEC profile of the MIT for the three different IMF conditions. We selected the corrected TEC from within 2 hours MLT of 18, then removed any examples which had no trough pixels, then averaged the profiles for each IMF condition (applying the 1 hour delay as before). In the right plot of figure 4.7, we show the differences between each pair of TEC profiles. Figure 4.8 has the same plots except for the morning sector within 2 hours MLT of 6.

The plot in figure 4.7 show that the increase in evening sector trough occurrence rate during westward IMF conditions is associated with increased plasma density




Figure 4-8: Average TEC profile of morning MIT (4 - 8 MLT) corrected for season

in the poleward wall region. The left plot shows the eastward IMF MIT minimum forming at about $68 - 69^\circ$ MLat and the poleward wall forming at about 75° MLat. For westward IMF, the minimum is around $70 - 71^\circ$ MLat and the poleward wall is closer to 80° MLat. From the right plot we see that the west IMF poleward wall is over 1 TECu higher on average (orange line). Another interesting thing about figure 4-7 is that at magnetic latitudes below 60° , southward IMF is associated higher TEC by 0.5 - 1.5 TECu.

In figure 4-9, we are comparing the average convection pattern in the vicinity of the MIT with TEC. Limiting the B field magnitude to between 2 and 6 nanotesla, for each IMF clock angle bin, we selected the corresponding SuperDarn image and the TEC image and labels from 1 hour later, i.e. if at time t the clock angle is in a particular bin, then we select the SuperDarn image from time t and the TEC image and labels from time $t + 1$ for that bin. To isolate the convection and TEC patterns contributing to the MIT, we enlarged the labels by applying a binary dilation operation with a square element of size (15×15) (<https://scikit-image.org/docs/stable/api/>





Figure 4.9: Average convection and corrected TEC in the vicinity of the MIT

`skimage.morphology.html#skimage.morphology.binary_dilation`), then only selected convection and TEC pixels which coincide with the enlarged labels. Finally, we averaged the convection and TEC in each bin (and grid cell) and discarded any which had less than 100 samples.


In the west (left) plot of figure 4.9, there is strong average flow from the TEC enhancement around 15MLT over the evening side of the polar cap. The flow appears to bend around the north pole towards the evening side, which deposits high-TEC plasma from the dayside in the evening sector around 80 MLat. From figure 4.7, this is where the trough poleward wall should be. Contrasting this with the east (right) plot, we see the high-TEC plasma from the dayside is transported around the north pole on the morning side instead of the evening side. This would deposit the high TEC dayside plasma on the morning side instead of the evening side which explains why the east IMF evening poleward wall is weaker. This also explains why, figure 4.8, there is a higher secondary TEC enhancement at around 85 MLat for eastward IMF. Finally, the south IMF plot (middle) shows the convection passing straight over the north pole. This would cause the high TEC plasma to be distributed on both the evening side and the morning side of the polar cap. The plots of the evening and morning average TEC profiles in figures 4.7 and 4.8, respectively, match this description well.

While not a new finding, this experiment demonstrates and clarifies the major role that convection and the IMF play in the formation of the MIT. Most descriptions of convection and the MIT focus on the stagnation mechanism in the premidnight sector. However, from this analysis, we are reminded that convection also contributes directly to the TEC enhancement making up the poleward wall of the MIT.



(a)

N = 6143485.0



(b)

Figure 4.10: Motion of MIT towards pre-midnight with increased Kp

4.3 Kp - MLT

In section 4.1.1, we demonstrated that known behaviors of the MIT, like Kp -associated latitudinal motion, are readily observable in our dataset. In this section we will show that Kp also appears to have an effect on the MLT of the trough: it forms mostly post-midnight at low Kp and shifts towards pre-midnight at higher Kp . While the MLat of the MIT has been thoroughly researched in the past, less attention has been paid to its MLT. This is partially due to the fact that satellite and radar studies provide very limited sampling of MLT, though even past work which utilized the Madrigal TEC dataset did not focus on it. Our dataset simultaneously has a long time span and good MLT sampling which makes it ideal for this type of study.

In figure 4-10a we plot the regional MIT occurrence rate at different Kp levels. The Kp bins were chosen according to the quantile edges: [0.0, 0.6, 0.8, 0.95, 1.0]. At the lowest Kp , the region of highest MIT occurrence is from midnight to 4 MLT and at each successive Kp level, this region appears to expand towards the evening sector. At the highest Kp level, the distribution rotates towards evening without expanding significantly. We show this relationship more directly in figure 4-10b by plotting the trough occurrence rate in Kp - MLT bins. The left plot shows the total Kp - MLT distribution in the dataset and the right plot shows a normalized version (same procedure as figure 4-1e). Also included in the right plot is the average trough MLT in Kp bins (dots), the standard deviation of each Kp bin (error bars) and a regression line of Kp vs MLT. The line has a small downward slope of -0.24 MLT per Kp . The linear regression only explains about .06 of the variance of the trough's MLT. Of course there are many factors that determine the MLT of the trough and additionally some variance is added by the incomplete coverage of the dataset. For both figures, we dithered Kp as explained in 4. For figure 4-10a, we left the labels on the original TEC grid, but for figure 4-10b we dithered MLT according to table 4.1.




Figure 4.11: SAPS contribution to MIT occurrence rate

A likely cause for the trough's MLT dependence on Kp is subauroral polarization streams (SAPS). To investigate this possibility, we identified SAPS in SuperDarn convection maps with the procedure described in section 3.1.5. Then for each element in our MIT labels we determined whether or not it was associated with SAPS using a procedure similar to the one described in section 4.2.2 (for flow vector averaging). We expanded the SAPS labels using a structuring element which was (5×5) at time t and (3×3) at time $t - 1$. This way, a trough pixel would be associated with SAPS if there was a simultaneous high-speed flow within 5 pixels, or if there was previously a high-speed flow within 3 pixels. The purpose of this was to account for possible delay between the measurement of a high-speed flow and the associated formation of a trough, as well as SAPS-associated troughs which remain for more than one hour. This procedure is a proof-of-concept and can only identify SAPS-associated troughs imprecisely. Finally, for low and high Kp , we determined the occurrence rates of non-SAPS-MIT, SAPS-associated-MIT, and total MIT, which should just be a sum of the previous two. These results are displayed in figure 4-11. To calculate these rates, we could only consider MIT labels where we also had SuperDarn data, otherwise we would mistakenly be counting all instances of missing convection data as "Non SAPS". As in section 4.2.2, we discarded grid cells in which we did not have at least 100 data points.

The bottom row of figure 4-11 shows the overall trend between low and high Kp , which is a shift of the MIT pre-midnight and equatorward. In the top row, we see that the peak occurrence rate of non-SAPS-MIT decreases at high Kp . The middle row shows that the increase in pre-midnight MIT occurrence rate is in large part due to higher SAPS occurrence rate. The peak of high- Kp total MIT occurrence in the bottom row is colocated with the peak of high- Kp SAPS-associated-MIT occurrence in the middle row. Finally, figure 4-12 shows the portion of total MIT that can be




Figure 4.12: Portion of troughs associated with SAPS

associated with SAPS. From these plots, it is apparent that evening-side troughs are often associated with high-speed westward flows, and this increases significantly at high Kp .

Together, these plots provide reasonable evidence that SAPS is a major contributor to the increased pre-midnight MIT occurrence rate at high Kp . Because we have not performed any verification on our SAPS identification method, this evidence is far from conclusive. These statistics depend on choices we made during data processing, and before we can have confidence in our results, our choices need to be interrogated. For example, we performed the same experiment with a flow speed threshold of 400 m/s instead of 300, which reduced the apparent contribution of SAPS. Regardless, this experiment indicates that this is a promising direction of research.

Chapter 5

Conclusion

5.1 Summary

In order to perform a large-scale statistical study of the main ionospheric trough, we developed methods for automatically identifying the MIT in total electron content images. We evaluated these methods by measuring the degree to which they agreed with the algorithm from Aa 2020. Our first method achieves high agreement with Aa 2020 and is well behaved in the presence of missing data. Our second method reduces the computation time from the first method by two orders of magnitude, but does not perform as well. Running our first algorithm on a 10 year dataset of TEC maps from the Madrigal database provided us with the large-scale dataset we wanted. To further validate our labeled dataset, we demonstrated that well known MIT characteristics, such as the MIT positional dependence on Kp and MLT, and the MIT occurrence dependence on season, are present.

Finally, we presented two novel findings about the MIT. First we showed how the IMF and convection pattern contribute to the MIT occurrence rate. Specifically we demonstrated that the evening poleward wall is much larger for westward IMF than for eastward which results in a higher average evening MIT occurrence rate. During this experiment we showed how the high latitude convection pattern intuitively contributes to the poleward wall of the MIT at all IMF clock angles and all MLTs. Our second finding is that the highest MIT occurrence rate occurs postmidnight at low Kp , but high Kp is associated with increased premidnight MIT occurrence rate. We then

provided evidence indicating that SAPS significantly contributes to this shift. At this stage, our findings are only preliminary, but we intend to verify them in the near future.

5.2 Future Work

This project provided a proof-of-concept for our labeling strategies and demonstrated the potential value of a large, annotated TEC image dataset. We are proud of what we have accomplished, but have many ideas for how to continue this research. The future projects we have planned fall into two broad categories: improvements to our data processing and applications for the dataset.

5.2.1 Data Processing

One improvement we wish to make is to improve the performance of *TLM2* or another fast algorithm. While *TLM2* seems to perform adequately, most of the project was spent developing *TLM1*. With a few tweaks, *TLM2* or some other filtering-based algorithm would become a better choice. The largest difficulty we had in working with the Madrigal TEC maps was dealing with missing data. Of course there is no replacement for actual measurements, but a fast algorithm which also gracefully handles missing data would improve on this work.

The optimization framework which we used for *TLM1* is very flexible and so another possibility is to incorporate additional measurements from satellites and radar. To accomplish this task, we would come up with additional forward models for each data source, then we would add the corresponding terms to the inversion problem. This would make our labeling strategy more dependent on measurements, allowing us to reduce our reliance on regularization, or equivalently, reduce the strength of our prior assumptions.

Finally, there are a number of ways in which neural networks could help with this project. Our approach suffered from several drawbacks which could potentially be remedied with a neural network approach. At a high level, when optimizing the parameters of our algorithms, we were trying to balance a few factors: that our labels agreed with the Aa 2020 MIT labels and that they appeared to correspond well to the TEC data. While the first factor we could quantify directly, the second factor was evaluated more subjectively. A potentially more principled way to balance these factors would be to train two neural networks, one to estimate the MIT labels from the TEC image, and a second to reconstruct the TEC from the estimated MIT labels (similar to architecture in (Xia and Kulis, 2017)). The reconstructed TEC data would be reasonable only if the MIT labels correspond well to the input TEC image. This approach provides a way to quantify the estimated labels' agreement with both the SWARM data and the TEC data.

5.2.2 Applications

The first follow-on effort we will make will be to finalize our two novel experiments. To finish the experiment from section 4.3, we will need to verify our SAPS labels. We can accomplish this by comparing the labels to ion drift data from the SWARM satellites. In section 4.2.2, we showed that on average, there is a relationship between IMF and the MIT poleward wall. A quick follow up on this would be to try and show a direct relationship between the clock angle and MIT poleward wall position and size.

Investigating the two dimensional inter-dependencies of different sections of the MIT is an opportunity that is uniquely enabled by this dataset. Unfortunately, we did not have time to pursue this possibility during this thesis, but it is something we would like to try in the future.

Finally, 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 in (Mrak et al., 2020) and determine the relationship between these two phenomena.

References

- Aa, E., Zou, S., Erickson, P. J., Zhang, S., and Liu, S. (2020). Statistical analysis of the main ionospheric trough using swarm in situ measurements. *Journal of geophysical research. Space physics*, 125(3).
- Bishop, C. M. (2006). *Pattern recognition and machine learning*. Information science and statistics. Springer, New York.
- Collis, P. and 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. and Shubin, V. N. (2018). Empirical model of the location of the main ionospheric trough. *Geomagnetism and aeronomy*, 58(3):348–355.
- Dudeney, J. R., Rodger, A. S., and Jarvis, M. J. (1983). Radio studies of the main f region trough in antarctica. *Radio science*, 18(6):927–936.
- Emmert, J. T., Richmond, A. D., and 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>.
- Foster, J. C. and Vo, H. B. (2002). Average characteristics and activity dependence of the subauroral polarization stream. *Journal of Geophysical Research: Space Physics*, 107(A12):SIA 16–1–SIA 16–10.
- GFZ Helmholtz Center Potsdam (2021). Geomagnetic Kp Index. <https://www.gfz-potsdam.de/en/kp-index/>. Accessed: 2021-07-14.
- Ishida, T., Ogawa, Y., Kadokura, A., Hiraki, Y., and 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.
- Johns Hopkins University Applied Physics Laboratory (2021). Special Sensor Ultraviolet Spectrographic Imager. <https://ssusi.jhuapl.edu/>. Accessed: 2021-07-13.

- Karpachev, A., Deminova, G., and Pulinets, S. (1995). Ionospheric changes in response to imf variations. *Journal of atmospheric and terrestrial physics*, 57(12):1415–1432.
- Karpachev, A. T. (2019). Variations in the winter troughs’ position with local time, longitude, and solar activity in the northern and southern hemispheres. *Journal of geophysical research. Space physics*, 124(10):8039–8055.
- Kintner, P. M., Ledvina, B. M., and de Paula, E. R. (2007). Gps and ionospheric scintillations. *Space Weather*, 5(9).
- Laundal, K. M. and Richmond, A. D. (2016). Magnetic Coordinate Systems. *Space Science Reviews*, 206(1-4):27–59.
- Le, H., Yang, N., Liu, L., Chen, Y., and 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.
- Lomidze, L., Burchill, J. K., Knudsen, D. J., Kouznetsov, A., and Weimer, D. R. (2019). Validity study of the swarm horizontal cross-track ion drift velocities in the high-latitude ionosphere. *Earth and space science (Hoboken, N.J.)*, 6(3):411–432.
- Mrak, S., Semeter, J., Nishimura, Y., Rodrigues, F. S., Coster, A. J., and Groves, K. (2020). Leveraging geodetic gps receivers for ionospheric scintillation science. *Radio science*, 55(11).
- NASA Goddard Space Flight Center Space Physics Data Facility (2021). OMNIWeb Data Documentation. https://omniweb.gsfc.nasa.gov/html/ow_data.html. Accessed: 2021-07-14.
- Nilsson, H., Sergienko, T. I., Ebihara, Y., and Yamauchi, M. (2005). Quiet-time mid-latitude trough: influence of convection, field-aligned currents and proton precipitation. *Annales geophysicae (1988)*, 23(10):3277–3288.
- Pedatella, N. M. and Larson, K. M. (2010). Routine determination of the plasma-pause based on COSMIC GPS total electron content observations of the midlatitude trough. *Journal of Geophysical Research: Space Physics*, 115(A9):<https://doi.org/10.1029/2010JA015265>.
- Pryse, S. E., Kersley, L., Malan, D., and 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., and 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.
- Rodger, A. (2008). *The Mid-Latitude Trough—Revisited*, pages 25–33. American Geophysical Union (AGU).
- Rodger, A., Moffett, R., and 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.
- Ruohoniemi, J. M. and Baker, K. B. (1998). Large-scale imaging of high-latitude convection with super dual auroral radar network hf radar observations. *Journal of Geophysical Research: Space Physics*, 103(A9):20797–20811.
- Spiro, R. W., Heelis, R. A., and Hanson, W. B. (1978). Ion convection and the formation of the mid-latitude f region ionization trough. *Journal of Geophysical Research: Space Physics*, 83(A9):4255–4264.
- Thomas, E. G. and Shepherd, S. G. (2018). Statistical patterns of ionospheric convection derived from mid-latitude, high-latitude, and polar superdarn hf radar observations. *Journal of geophysical research. Space physics*, 123(4):3196–3216.
- Vierinen, J., Coster, A. J., Rideout, W. C., Erickson, P. J., and Norberg, J. (2016). Statistical framework for estimating gnss bias. *Atmospheric measurement techniques*, 9(3):1303–1312.
- Voiculescu, M., Virtanen, I., and Nygrén, T. (2006). The f-region trough: seasonal morphology and relation to interplanetary magnetic field. *Annales geophysicae (1988)*, 24(1):173–185.
- Werner, S. and Prölss, G. (1997). The position of the ionospheric trough as a function of local time and magnetic activity. *Advances in Space Research*, 20(9):1717–1722.
- Xia, X. and Kulis, B. (2017). W-net: A deep model for fully unsupervised image segmentation. <https://arxiv.org/abs/1711.08506>.
- Yang, N., Le, H., and 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).

Zou, S., Moldwin, M. B., Coster, A., Lyons, L. R., and Nicolls, M. J. (2011). Gps tec observations of dynamics of the mid-latitude trough during substorms. *Geophysical Research Letters*, 38(14).

CURRICULUM VITAE

