



A Stacked Ensemble Machine Learning Model for Predicting Ionospheric Total Electron Content Over Uganda

Adriko Norbert Bakole

BU/GS24/MSP/21


A dissertation Submitted to the Directorate of Graduate Studies, Research and
Innovations in Partial Fulfillment of the Requirements for the Award of the Degree
of Master of Science in Physics of Busitema University

August, 2026

Declaration

I Adriko Norbert Bakole declare that the work presented in this dissertation is my original work and has not been presented in any university or institution for any academic award, and where I have used the work of other authors, due acknowledgments and appropriate citations have been made.


Signature: 

Date: 

Approval


The research work for which this dissertation is written was conducted under the guidance and supervision of the following supervisors:

Name: Dr. Geoffrey Andima

Signature: 

Date: 28/08/2026

Name: Assoc. Prof. Emirant Bertillas Amabayo

Signature: 

Date: 27-08-2026

Dedication

To my beloved family, for your endless prayers, unconditional love, and unwavering support throughout this academic journey, your deep sacrifices laid the foundation for my education, and your constant belief in my potential has been my greatest anchor and motivation during the challenging hours of this research.

Acknowledgment

I thank the Almighty God for giving me the strength and guidance to go through with this research work.

My sincere appreciation goes to my supervisors, Dr. Geoffrey Andima and Assoc. Prof. Emirant Bertillas Amabayo, for their constant support towards the progress of my masters work and their willingness to push me to pursue excellence in my work.

This work was funded by the Space Weather Based Early Warning Tool Project under the award BU-STI-OP/2023/AERONAUTICS AND SPACE /06. My sincere thanks goes to the Principal Investigator and Co-principal investigator of the project, Assoc. Prof. Amabayo Emirant Bertillas and Dr. Andima Geoffrey respectively for giving me the opportunity to be part of the beneficiaries of this project.

Special thanks also goes to Busitema University and especially the members of the physics department at the faculty of science and education nagongera campus for their academic prowess and support in helping me build the theoretical concepts in this work and offering me a graduate fellow position at the University.

I also want to extend my heartfelt gratitude to my friends, family and coursemates for their continued teamwork and togetherness during my studies. I particularly appreciate their encouragement and motivation during the challenging hours of my research work.

Table of Contents

Declaration	i
Approval	ii
Dedication	iii
Acknowledgment	iv
Table of Contents	v
List of Figures	ix
List of Tables	ix
Acronyms and Abbreviations	xi
Abstract	xii
1 Introduction	1
1.1 Background	1
1.2 Problem Statement	4
1.3 Aim	4
1.4 Objectives	4
1.5 Scope	5
1.6 Significance	5
2 Theory and Literature Review	6
2.1 The Ionosphere	6

2.1.1	Formation	6
2.1.2	Ionospheric Layers	7
2.1.3	Refraction of Radio Waves in the Ionosphere	8
2.1.4	Total Electron Content and GPS Measurement	10
2.2	TEC Variability at Equatorial and Low Latitudes	14
2.2.1	Solar and Geomagnetic Drivers	14
2.2.2	Equatorial Ionization Anomaly	15
2.2.3	Storm-Time TEC Variations	15
2.3	Ionospheric Models	16
2.3.1	Physics-Based Models	17
2.3.2	The Klobuchar Broadcast Model	17
2.3.3	The International Reference Ionosphere	18
2.3.4	The NeQuick Model	19
2.3.5	Global Ionosphere Maps	19
2.3.6	Regional Models over Africa.	21
2.3.7	Methods of Modelling Ionospheric TEC	21
2.3.7.1	Neural Network Approaches	21
2.3.7.2	Recurrent and Sequence Models	22
2.3.7.3	Gradient Boosting and Tree-Based Methods	23
2.3.7.4	Ensemble and Stacking Methods	23
2.4	Summary	24
3	Data and Methods	25
3.1	Study Area and Regional Context	25

3.1.1	Longitudinal Modulation by Field Geometry	25
3.1.2	Geographic and Ionospheric Context	27
3.1.3	Previous Studies	27
3.2	Data Sources	29
3.2.1	CODE Global Ionosphere Maps	29
3.2.2	Uganda CORS GPS Network	29
3.2.3	Solar and Geomagnetic Indices	33
3.3	Data Preprocessing	33
3.3.1	GIM Spatial and Temporal Interpolation	33
3.3.2	CORS Data Assembly	34
3.3.3	Prevention of Data Leakage	35
3.4	TEC Characterisation Methodology	35
3.5	Feature Engineering	36
3.6	Machine Learning Models	37
3.6.1	Stage 1: XGBoost Gradient Boosted Tree Model	37
3.6.1.1	Model Architecture	37
3.6.1.2	Hyperparameters	38
3.6.2	Stage 2: Artificial Neural Network Correction Model	39
3.6.2.1	Input Features and Architecture	40
3.6.2.2	Training Configuration	40
3.6.2.3	Per-Station Bias Correction	41
3.6.2.4	Implementation and Reproducibility	41
3.7	Two-Stage Stacked Ensemble Architecture	42

3.7.1	Model Formulation	42
3.8	Evaluation Metrics	43
4	Results and Discussion	47
4.1	TEC Characterisation	47
4.1.1	TEC Diurnal Variability	47
4.1.2	Seasonal Variability	48
4.1.3	Storm-Time TEC Variability	49
4.2	Stacked Ensemble TEC Model	51
4.2.1	XGBoost Model	52
4.2.2	Residual Correction	54
4.2.2.1	Final Ensemble	55
4.3	Evaluation of the Two-Stage Ensemble Model Performance	56
4.3.1	Model Comparison Against CORS Ground Truth	56
4.3.2	Diurnal Performance	58
4.3.3	Storm-Time Performance	60
4.3.4	Spatial Performance	61
5	Conclusions and Recommendations	63
5.1	Conclusions	63
5.2	Recommendations	63

List of Figures

2.1	Geometry of the ionosphere	8
2.2	Geometry of the GPS thin-shell ionospheric model	13
3.1	Study area: Uganda and the 40-station CORS GPS network	30
3.2	Temporal division of the datasets	35
3.3	XGBoost gradient-boosted tree ensemble	38
3.4	Two-stage stacked ensemble architecture	44
4.1	Diurnal TEC cycle over Uganda	47
4.2	Semi-annual VTEC anomaly over the Uganda longitude sector	49
4.3	Storm-time TEC variability over the November 2025 case study	50
4.4	SHAP beeswarm plot for Stage 1 XGBoost	51
4.5	XGBoost observed vs predicted scatter	53
4.6	ANN training history	54
4.7	Per-station Stage 1 bias shift from training to test	55
4.8	Three-tier time series at station TIDO	58
4.9	Three-tier time series at ABIM and MBRA	59
4.10	Diurnal RMSE profile for the two-stage ensemble	59
4.11	Spatial distribution of Stage 2 ensemble performance	61
4.12	Leave-One-Station-Out cross-validation of the Stage 2 ANN	62

List of Tables

3.1	Uganda CORS station coordinates	32
3.2	Model input features for the Stage 1 XGBoost model with the target being VTEC measured in TECU	36
3.3	XGBoost Stage 1 hyperparameters	38
4.1	Stage 1 XGBoost performance on the CODE GIM grid	52
4.2	Ensemble model performance metrics	56
4.3	Model comparison against CORS ground truth	57
4.4	Quiet and storm-time stratification of the ensemble	60

Acronyms and Abbreviations

AACGM	Altitude-Adjusted Corrected Geomagnetic coordinates
AIUB	Astronomical Institute of the University of Berne
ANN	Artificial Neural Network
Ap	Geomagnetic Planetary Activity Index
B _z	Southward component of the Interplanetary Magnetic Field
CDDIS	Crustal Dynamics Data Information System
CODE	Center for Orbit Determination in Europe
CORS	Continuously Operating Reference Station
DCB	Differential Code Bias
Dst	Disturbance Storm-Time index
EIA	Equatorial Ionization Anomaly
EPB	Equatorial Plasma Bubble
EUV	Extreme Ultraviolet
F10.7	Solar Radio Flux at 10.7 cm wavelength
GIM	Global Ionosphere Map
GNSS	Global Navigation Satellite System
GPS	Global Positioning System
HF	High Frequency
IGRF	International Geomagnetic Reference Field
IGS	International GNSS Service
IONEX	IONosphere Map EXchange
IPP	Ionospheric Pierce Point
IRI	International Reference Ionosphere
Kp	Geomagnetic Planetary Index
LOSO	Leave-One-Station-Out cross-validation
LSTM	Long Short-Term Memory
LT	Local Time
MAE	Mean Absolute Error
MBE	Mean Bias Error
ML	Machine Learning
MLP	Multi-Layer Perceptron
NASA	National Aeronautics and Space Administration
NmF2	Peak Electron Density of the F2 layer
OOF	Out-Of-Fold predictions
QD	Quasi-Dipole magnetic coordinates
RMSE	Root Mean Square Error
ROTI	Rate of TEC change Index
SC	Solar Cycle
sfu	Solar Flux Unit
SHAP	SHapley Additive exPlanations
SS	Skill Score
TEC	Total Electron Content
TECU	TEC Unit (1 TECU = 10 ¹⁶ electrons m ⁻²)
UT	Universal Time
UV	Ultraviolet
VTEC	Vertical Total Electron Content
XGBoost	eXtreme Gradient Boosting

Abstract

Ionospheric delay, which is directly determined from Total Electron Content (TEC) is the largest source of error in GPS positioning. Because many GPS receivers widely used in Uganda lack the capacity to measure this delay directly, they must rely on models to correct for it. The models currently available are global and therefore capture neither the steep latitudinal gradient of the equatorial ionization anomaly or the day-to-day variability typical of low-latitude regions such as Uganda, which creates a need for a local model. This study first characterised the behaviour of TEC over Uganda and then developed a regional TEC model using a stacked machine learning ensemble trained on 26 years of Global Ionosphere Map data and six months of TEC observations from the Continuously Operating Reference Station (CORS) network in Uganda. The local correction stage relies on observation-derived residual lags, which bounds its evaluation to the station density and time span sampled here. The characterisation revealed a post-noon maximum between 13:00 and 17:00 local time, with night-time values remaining above 5 TECU. A semi-annual anomaly was also evident with the March monthly mean exceeding the June–July mean by 66%. The storm-time response proved to be a latitudinal redistribution rather than a network-wide change during the November 2025 storm: TEC at the northernmost station (ABON) rose markedly above its quiet-day baseline, while the southernmost station (KABA) tracked its own baseline to within 1 TECU. Global Ionospheric Map observations and IRI-2020 carried a mean bias of +18.0 TECU and +27.9 TECU respectively when compared against the ground truth CORS network observations. On the test period, the full ensemble achieved a $\text{RMSE} = 3.970$ TECU, $R^2 = 0.958$ and a Skill Score of +0.724 against the climatological mean. This correlates with an 80% reduction in RMSE relative to the background field at the same locations thereby reducing ranging error of L1 from 3.26 m (background error) to 0.64 m (ensemble prediction). Spatial performance showed a strong trough-to-crest TEC latitudinal transition; leave-one-station-out cross-validation confirmed the model generalises to unseen locations, though with degraded accuracy (RMSE rising from 3.97 to 6.41 TECU) relative to stations seen during training. These observations suggest that the ensemble is a useful tool for ionospheric delay correction over Uganda.

Chapter 1

Introduction

1.1 Background

The ionosphere is the ionized region of Earth’s upper atmosphere, roughly 60–1000 km in altitude, where solar extreme ultraviolet (EUV) and X-ray radiation photoionize neutral atmospheric constituents, producing a plasma of free electrons and ions embedded in a neutral background gas (Kelley, 2009; Rishbeth & Garriott, 1969).

These free-electron population interact with GNSS signals while they move through the ionosphere. Due to the dispersive nature of the ionosphere, the first order effect on a GNSS signal depends on the signal frequency and on the electron density intergrated along the satellite to reciever propagation path. This path intergral is what is refered to as the sland total electron content defined as

$$\text{STEC} = \int_{\text{path}} N_e, dl, \quad (1.1)$$

where N_e is the electron density and dl is an infinitesimal element along the signal propagation path. At first order, the ionosphere introduces a positive group delay in the code (pseudorange) measurement and an equal-magnitude phase advance in the carrier-phase measurement, with the magnitude of both effects proportional to STEC/f^2 . Consequently, for a single-frequency receiver, an uncorrected ionospheric delay appears primarily as an additional range error in the pseudorange observations and can degrade the estimated receiver position. The magnitude of this effect varies with frequency, satellite elevation, geographic location, local time, solar activity, and geomagnetic conditions. The ionosphere is partitioned into the D (60–90 km), E (90–150 km), and F (above 150 km) regions (Kelley, 2009; Schunk & Nagy, 2009). The F layer matters most for GPS positioning since it carries the highest electron densities and persist through out the night which contribute substantially to the total electron content (Davies, 1990).

TEC is measured in TEC units (TECU), where $1 \text{ TECU} = 10^{16} \text{ electrons m}^{-2}$. At the GPS L1 frequency of 1575.42 MHz, one TECU introduces a ranging error of approximately 0.162 m (Mannucci et al., 1999). For single-frequency GPS receivers,

which are the most common type in Uganda and across developing countries, this ionospheric error is the largest single source of position inaccuracy in the absence of an ionospheric correction (Komjathy, 1997).

Uganda spans from about 1.3° S to 4° N in latitude and from 30° E to 35° E in longitude. This places the country directly under the Equatorial Ionization Anomaly (EIA) (Kelley, 2009; Mitra, 1946). The EIA generates two belts of elevated electron density on either side of the dip equator, at roughly $\pm 15^\circ$ magnetic latitude. The mechanism behind it is the fountain effect in that during the day, eastward electric fields lift plasma upward at the magnetic equator, the plasma then drifts outward along the geomagnetic field lines toward higher latitudes and builds the two density peaks (Rishbeth, 1997). As a consequence, TEC over Uganda exhibits strong variability with local time, season, and level of solar activity (Olwendo et al., 2016).

Around the maximum of Solar Cycle 24, daytime peak VTEC at East African EIA crest stations ranged from about 60 TECU to about 80 TECU across the equinoctial months, the spread reflecting differences between stations and between individual days rather than an uncertainty on any single value (Olwendo et al., 2016). Solar flux and geomagnetic storm activity in 2025 exceeded SC24 levels by a clear margin which means that the ionospheric conditions over Uganda during the solar maximum were correspondingly intense. For instance, the strong geomagnetic storm of 1 January 2025 produced marked positive and negative TEC perturbations at Mbarara, Uganda, driven by prompt penetration electric fields and disturbance dynamo processes (Akinyemi et al., 2026).

Several studies have examined TEC over East Africa and Uganda. Olwendo et al. (2016) showed that regional TEC is strongly controlled by local time, season, and geomagnetic activity. Nibigira et al. (2019) quantified the solar-cycle effect on TEC at a Ugandan IGS station throughout Solar Cycle 24. Cele et al. (2023) later mapped spatial TEC gradients across the East African sector. Together, these studies show that the low-latitude ionosphere over Uganda behaves very differently from mid-latitudes, and global models built for mid-latitude conditions do not transfer here without local calibration.

For GNSS ionospheric correction, the most widely used global product is the Global Ionosphere Map (GIM), distributed by the International GNSS Service (IGS) (Hernández-Pajares et al., 2009; Schaer et al., 1998). Several IGS analysis centres produce these maps, including CODE at the University of Berne, JPL at NASA, and ESA, each of these apply independent algorithms to a global network of GNSS receivers to derive VTEC on a common grid (Hernández-Pajares et al.,

2009). Among these, the CODE GIM is adopted in this study because it provides continuous VTEC records on a 2.5° latitude by 5° longitude grid at one-hour intervals reaching back to 1996 (Schaer et al., 1998), giving a 26-year training archive that no other publicly available product matches. However, GIMs are constructed predominantly from mid-latitude GNSS stations where data density is highest, and their interpolation schemes tend to underperform at low latitudes, where the ionosphere is most structured (Hernández-Pajares et al., 2009). The EIA introduces steep horizontal TEC gradients that a $2.5^\circ \times 5^\circ$ latitude and longitude global grid cannot adequately resolve, producing systematic positive biases at low latitudes. This limitation motivates the development of regional ionospheric models tailored to the low-latitude environment.

One regional effort is AfriTEC (Okoh et al., 2019), a feedforward neural network trained on solar flux and geomagnetic indices to predict TEC across the African continent. AfriTEC outperformed IRI-2016 at most test stations, demonstrating that data-driven approaches can capture variability that empirical climatological models miss. However, AfriTEC was trained on a limited set of African GNSS stations, and East Africa where observational density is comparatively low and under-represented in its training data. As a result, its predictions over Uganda carry the same spatial smoothing problem as global GIMs i.e. they reflect a continental average rather than the sharp local gradients of the EIA crest.

Uganda presents a particularly demanding test case. The magnetic dip equator crosses this longitude sector near 8° N, so the whole country lies in the southern magnetic hemisphere, spanning roughly 5° to 11° south magnetic latitude. That places it across the transition from the equatorial trough to the southern EIA crest, where TEC variability is highest, solar cycle modulation is strongest, and global model biases are largest. Studies over the broader East African region have employed a range of methods such as Empirical models like the International Reference Ionosphere (IRI) that provide a useful climatological background but could not adapt to storm-time dynamics or inter-solar-cycle changes. Single-station machine learning approaches like neural networks, gradient boosting, and LSTMs have shown skill improvements over IRI (Nigusie et al., 2024) but are limited in spatial coverage by design. Stacked ensemble methods that combine multiple learners have consistently reduced prediction error relative to any single model (Mao, Pan, et al., 2025; Nigusie et al., 2024; Zhou et al., 2024), yet none has been applied to Uganda using a dense ground-based network.

This study addresses that gap by characterizing the multi-scale variability of VTEC

over Uganda and developing a dedicated regional ionospheric model. Through leveraging long-term global ionospheric maps alongside dense local observations from the Uganda CORS network, this research captures both large-scale drivers and fine-grained regional features.

1.2 Problem Statement

Many GPS receivers widely used in Uganda cannot directly measure ionospheric delays and therefore must rely on model predictions to correct it. Total Electron Content is the largest source of the GPS errors. Global models are limited in that they lack the high spatial resolution needed to capture strong TEC gradients caused by the Equatorial Ionization Anomaly and also the models struggle to represent short-term variations driven by rapid solar and geomagnetic changes. Therefore, there is a need for regional models to be applied in the delay correction. Such a model is developed in this study using a regional network of 40 Continuous Operating Reference Stations (CORS) in Uganda with solar and geomagnetic activity within a machine-learning framework to develop a locally calibrated VTEC prediction model for Uganda.

1.3 Aim

The aim of this study was to develop a TEC model over Uganda using a two-stage stacked ensemble machine learning method.

1.4 Objectives

The following objectives guided the study:

1. To characterise the variability of TEC over Uganda by using Uganda CORS data to investigate diurnal and storm-time variations, and Global Ionospheric Map (GIM) data to investigate seasonal trend.
2. To model TEC using an XGBoost gradient-boosted tree as a background predictor and a feedforward Artificial Neural Network to correct its systematic bias against ground-based CORS GPS measurements.

3. To evaluate the diurnal, storm-time, and spatial performance of the two-stage ensemble model against IRI-2020, the Stage 1 background, and simple reference forecasts, using standard statistical metrics.

1.5 Scope

The study focuses on characterization of TEC over Uganda using CORS data from the period August 2025 to April 2026 and GIM data over Uganda for the period 2000 to 2026. The background model was trained on GIM data and residue correction model used the CORS data of Uganda. The study is limited to modelling TEC over Uganda.

1.6 Significance

Earlier studies relied on single IGS stations or on regional datasets not specific to Uganda (Nibigira et al., [2019](#); Olwendo et al., [2016](#)). This work therefore adds to that existing body of knowledge especially for the case of spatial TEC modelling over Uganda.

The model establishes a basis for regional ionospheric correction in surveying, precision agriculture and unmanned aerial vehicle applications that depend on low-cost single-frequency receivers.

Chapter 2

Theory and Literature Review

This chapter reviews the background on which the study rests. Sections 2.1 and 2.2 cover the physics of the ionosphere and the drivers of TEC variability at equatorial and low latitudes. Section 3.1 turns to East African and Ugandan TEC studies. Section 2.3 reviews the ionospheric models in common use for TEC specification and correction, and Section 2.3.7 reviews the methods by which such models are constructed, including the machine learning approaches adopted here. Section ?? states the resulting research gap.

2.1 The Ionosphere

2.1.1 Formation

The ionosphere is a region of the upper atmosphere of Earth having significant number of charged particles. It has no fixed boundaries but is typically considered to exist within 60 to 1000 km altitude (Goodman, 2005). The ionosphere forms through photoionisation where Solar extreme ultraviolet (EUV) radiation in the 10–121 nm range, together with soft X-rays, carries sufficient energy to eject an electron from a neutral oxygen or nitrogen atom or molecule, leaving a positive ion in its place (Rishbeth & Garriott, 1969; Schunk & Nagy, 2009). The production rate at any altitude depends on the solar zenith angle and the local density of absorbing neutral particles, following the Chapman production function .i.e it peaks at local noon when the Sun is directly overhead and falls to zero on the night side as the ionising source is extinguished.

Free electrons are lost through recombination, and two distinct mechanisms operate. In radiative recombination, an electron combines directly with a positive ion and the excess energy is released as a photon. This process is extremely slow and is negligible at F-layer altitudes. Dissociative recombination is far more efficient where an electron attaches to a molecular ion such as NO^+ or O_2^+ , and the resulting neutral molecule immediately dissociates into two atoms.. The rate coefficient for

dissociative recombination ($\sim 10^{-7} \text{ cm}^3 \text{ s}^{-1}$) is several orders of magnitude larger than for radiative recombination ($\sim 10^{-12} \text{ cm}^3 \text{ s}^{-1}$), making it the dominant electron-loss mechanism wherever molecular ions are abundant (Rishbeth & Garriott, 1969; Schunk & Nagy, 2009). In the upper F layer, where atomic O^+ is the dominant ion, dissociative recombination requires a prior charge-exchange step with N_2 or O_2 before a molecular ion becomes available. This two-step pathway is slow, which is why F-layer electrons survive overnight while the lower layers do not.

A third process, plasma transport by electric fields and neutral-wind-driven diffusion, redistributes electrons along and across magnetic field lines and shapes the large-scale structures described in the sections that follow (Kelley, 2009). TEC by definition is the total integral of electron density over all layers and therefore tracks the overall ionisation state i.e. it peaks at solar maximum, when EUV output is highest, and falls at solar minimum, with low-latitude TEC changing by a comparable factor to the roughly twofold variation in EUV flux across the solar cycle (Hathaway, 2015).

2.1.2 Ionospheric Layers

Ionospheric structure is conventionally split by altitude and ion chemistry into three regions.

The D layer (60–90 km) is daytime-only. Solar Lyman- α radiation and cosmic rays drive its ionisation, and the layer collapses quickly after sunset as electrons recombine (Laštovička, 2009). It absorbs high-frequency radio waves but does little to GPS signals.

The E layer (90–150 km) is also primarily a daytime feature. NO^+ and O_2^+ dominate its ion chemistry, and daytime electron density tracks the solar zenith angle closely enough to be predictable on a routine basis (Rishbeth & Garriott, 1969). At night, the abundant molecular ions recombine rapidly through dissociative recombination, reducing E-layer electron density by two to three orders of magnitude within hours of sunset. A thin, patchy residual known as sporadic E occasionally persists at night through wind-shear concentration of metallic ions, but its occurrence is irregular and unpredictable (Kelley, 2009).

The F layer (above 150 km) is the layer that matters most for GPS and satellite navigation. During the day it splits into an F1 sublayer (150–200 km) and an F2 sublayer (above 200 km), with O^+ ions dominating the upper region. F2 carries the highest electron density of all three layers and survives the night. At F2 altitudes,

the neutral gas is thin enough that recombination is slow, so daytime electrons persist until the next sunrise (Davies, 1990; Kelley, 2009).

Figure 2.1 shows the ion density distribution in the various ionospheric layers which are described above. it also details the ion density in these various layers.




Figure 2.1: Representation of ion density in the ionosphere with structured D, E F1 and F2 layers. (Banks & Kockarts, 1973)

2.1.3 Refraction of Radio Waves in the Ionosphere

The ionosphere is a dispersive medium which means its refractive index depends on the local electron density and on the frequency of the wave traversing it. For a magnetised, collisional plasma the phase refractive index n is given by the Appleton–Hartree formulation (Davies, 1990):

$$n^2 = 1 - \frac{X}{1 - jZ - \frac{Y_T^2}{2(1 - X - jZ)} \pm \sqrt{\frac{Y_T^4}{4(1 - X - jZ)^2} + Y_L^2}}, \quad (2.1)$$

where the dimensionless quantities relate the wave frequency f to the plasma frequency, the collision frequency ν and the electron gyrofrequency:

$$X = \frac{N_e e^2}{4\pi^2 \varepsilon_0 m f^2}, \quad Z = \frac{\nu}{2\pi f}, \quad Y_L = \frac{e B_L}{2\pi m f}, \quad Y_T = \frac{e B_T}{2\pi m f}. \quad (2.2)$$

Here N_e is the electron density, e and m are the electron charge and mass, ε_0 is the permittivity of free space, and B_L and B_T are the longitudinal and transverse components of the geomagnetic field relative to the direction of phase propagation.

At F-region altitudes the collision frequency is small enough that $Z \approx 0$ at GPS frequencies, and for a first-order treatment the geomagnetic terms may be dropped ($Y_L = Y_T = 0$), reducing Equation 2.1 to

$$n^2 = 1 - X = 1 - \frac{N_e e^2}{4\pi^2 \varepsilon_0 m f^2}. \quad (2.3)$$

The explicit f^2 dependence in Equation 2.3 is the origin of the dispersive behaviour exploited throughout this research. Retaining only first-order terms, the phase and group velocities of a wave in the ionosphere are

$$v_\phi = c \left(1 + \frac{40.3 N_e}{f^2} \right), \quad v_g = c \left(1 - \frac{40.3 N_e}{f^2} \right), \quad (2.4)$$

so the group velocity always lags the phase velocity. GNSS code measurements propagate at the group velocity and carrier phase measurements at the phase velocity (Hofmann-Wellenhof et al., 2007), which is why code pseudoranges are lengthened and carrier phase ranges shortened by the same ionospheric term.

The signal delay is the difference between the true geometric range and the optical path actually travelled, $I = \int_R^T (1 - n) ds$, where ds is an element of path length between receiver R and transmitter T . Substituting Equation 2.4 gives the group and phase delays

$$I_g = \frac{40.3}{f^2} \int_R^T N_e ds, \quad I_\phi = -\frac{40.3}{f^2} \int_R^T N_e ds. \quad (2.5)$$

The two are equal in magnitude and opposite in sign. The integral $\int_R^T N_e ds$ is the number of electrons in a column of 1 m^2 cross-section along the ray path, that is, the slant total electron content of Section 2.1.4.

2.1.4 Total Electron Content and GPS Measurement

The ionosphere delays radio signals which is inverse proportional to the square of the carrier frequency. For a GPS code signal at carrier f , the ionospheric range error follows directly from the group delay of Equation 2.5:

$$\Delta R = \frac{40.3}{f^2} \text{STEC}, \quad (2.6)$$

with ΔR in metres when STEC is in el/m^2 (Klobuchar, 1996). At the GPS L1 carrier ($f_1 = 1575.42 \text{ MHz}$) one TECU ($10^{16} \text{ el}/\text{m}^2$) costs about 0.163m of range error. Biases of this magnitude therefore translate directly into metre-level single-frequency positioning errors.

GPS data also yield TEC itself by exploiting this same frequency dependence. A dual-frequency receiver observes the same satellite signal at two carriers, f_1 and f_2 , and logs both a code pseudorange P_i and a carrier phase L_i at each. In range units these observables are (Hofmann-Wellenhof et al., 2007)

$$P_i = \rho + \frac{40.3}{f_i^2} \text{STEC} + b_i^r + b_i^s + \varepsilon_{P_i}, \quad (2.7)$$

$$L_i = \rho - \frac{40.3}{f_i^2} \text{STEC} + \lambda_i N_i + \varepsilon_{L_i}, \quad (2.8)$$

where ρ collects the geometric range, the receiver and satellite clock offsets and the tropospheric delay, every term that is independent of frequency – b_i^r and b_i^s are the receiver and satellite hardware biases, $L_i = \lambda_i \phi_i$, N_i is an integer ambiguity constant over an unbroken (cycle-slip-free) arc, and ε denotes measurement noise. The ionospheric term enters Equations 2.7 and 2.8 with opposite sign, as required by Equation 2.5, the ionosphere retards the code modulation but advances the carrier phase by the same first-order amount.

The two carriers are not arbitrary choices. Every GPS satellite derives its transmissions from a single onboard atomic standard running at a fundamental frequency of 10.23 MHz, and the two legacy civil carriers are coherent integer multiples of it, $f_1 = 154 \times 10.23 = 1575.42 \text{ MHz}$ and $f_2 = 120 \times 10.23 = 1227.60 \text{ MHz}$, with wavelengths $\lambda_1 \approx 19.0 \text{ cm}$ and $\lambda_2 \approx 24.4 \text{ cm}$ (Hofmann-Wellenhof et al., 2007). That coherence is what makes dual-frequency ionospheric sensing possible. Because both carriers leave the spacecraft locked to the same oscillator, a measured difference between them cannot be blamed on the transmitter, and the dispersive ionosphere becomes the dominant remaining explanation. The second frequency exists in the

system design for precisely this purpose.

A carrier phase observable is not a range in the sense that a pseudorange is. The receiver generates a local replica of the incoming carrier and tracks the phase difference between replica and signal in a phase-locked loop. What it records at each epoch is the fractional part of the current cycle, resolved to roughly 0.01 of a cycle — about 2 mm at λ_1 together with a running count of whole cycles accumulated since the loop first achieved lock. What it cannot observe is how many whole wavelengths separated satellite and receiver at the instant of that first lock. This unknown is the integer ambiguity N_i of Equation 2.8: an integer, unobservable in a single epoch, and constant for as long as lock is maintained (Hofmann-Wellenhof et al., 2007). The precision of the phase observable and the indeterminacy of its absolute level are therefore two faces of the same measurement.

The sign difference between Equations 2.7 and 2.8 is worth dwelling on, since it is easily misread as a contradiction. Equation 2.3 gives a phase refractive index below unity, so the phase velocity of the carrier exceeds c and the carrier arrives phase-advanced rather than delayed. Nothing travels faster than light in this description. The modulation that carries the navigation message propagates at the group velocity of Equation 2.4, which falls below c by the same first-order amount, and it is the modulation that conveys information (Davies, 1990). For TEC estimation the consequence is useful rather than paradoxical: code and phase respond to the same electron content with equal magnitude and opposite sign, which is what allows the two to be combined into a single calibrated series once the geometry-free combinations below have been formed.

The constancy of N_i holds only while lock holds. If the line of sight is obstructed, if the tracking loop is stressed at low elevation, or if the received amplitude fades sharply, the receiver may briefly lose and then reacquire the carrier, and N_i jumps by an integer amount. A cycle slip of this kind leaves a discontinuity in the phase series that is indistinguishable, in a single epoch, from a genuine ionospheric change. Slips must therefore be detected and either repaired or used to terminate the arc, and automatic editing schemes based on the geometry-free and wide-lane combinations are the standard means of doing so (Blewitt, 1990). The problem is particularly acute in the region studied here. Equatorial plasma irregularities produce deep post-sunset amplitude fades over East Africa (Amabayo et al., 2014; Spogli et al., 2023), so cycle slips cluster in the evening hours, which is also when equatorial TEC is changing most rapidly and when observations are most wanted.

L1 and L2 are also not equally easy to track. Under Anti-Spoofing the precise code on

L2 is encrypted, and receivers without access to it must recover the L2 carrier by semi-codeless techniques that cross-correlate or square the signal. These techniques carry a signal-to-noise penalty, with the result that the L2 phase observable is noisier than L1 and loses lock more readily, especially at low elevation angles and under scintillation (Hofmann-Wellenhof et al., 2007). Since every ionospheric quantity derived below depends on differencing L1 against L2, the weaker of the two observables sets the quality of the result. This is part of the justification for the 30° elevation mask adopted in Chapter 3, which discards precisely the geometry in which L2 tracking is least reliable.

Because ρ is common to both frequencies, differencing across f_1 and f_2 eliminates it and leaves the geometry-free combinations

$$P_2 - P_1 = 40.3 \text{ STEC} \left(\frac{1}{f_2^2} - \frac{1}{f_1^2} \right) + B_P, \quad (2.9)$$

$$L_1 - L_2 = 40.3 \text{ STEC} \left(\frac{1}{f_2^2} - \frac{1}{f_1^2} \right) + (\lambda_1 N_1 - \lambda_2 N_2) + B_L, \quad (2.10)$$

in which B_P and B_L are the combined receiver and satellite differential biases on the code and phase observables respectively. The two combinations carry an identical ionospheric coefficient and differ only in what contaminates it. Solving Equation 2.9 gives the code-derived slant TEC (Mannucci et al., 1999)

$$\text{STEC}_P = \frac{f_1^2 f_2^2}{40.3 (f_1^2 - f_2^2)} (P_2 - P_1 - B_P), \quad (2.11)$$

and Equation 2.10 the phase-derived estimate

$$\text{STEC}_L = \frac{f_1^2 f_2^2}{40.3 (f_1^2 - f_2^2)} (L_1 - L_2 - \lambda_1 N_1 + \lambda_2 N_2 - B_L). \quad (2.12)$$

The two estimates fail in complementary ways. Equation 2.11 is absolutely calibrated once B_P is known, but inherits decimetre-to-metre code noise. Equation 2.12 is an order of magnitude more precise – carrier tracking noise is millimetre-level – but the arc-constant ambiguity term $(\lambda_1 N_1 - \lambda_2 N_2)$ leaves it precise in shape and unknown in absolute level. The standard resolution is carrier-to-code leveling: for each continuous arc, the constant offset between the noisy code-derived STEC and the precise phase-derived STEC is estimated and added back onto the phase series, so the leveled TEC keeps the absolute calibration of Equation 2.11 and the low noise of Equation 2.12 (Ciraolo et al., 2007). This is the technique implemented in GPS-TEC processing software such as Seemala (2023), and is the method used in

Chapter 3 to derive station-level TEC from the CORS RINEX archive.

The slant TEC of Equations 2.11 and 2.12 depends on the viewing geometry, so observations from different satellites are not directly comparable until they are mapped to a common vertical reference. A thin-shell mapping function converts STEC to VTEC at the Ionospheric Pierce Point (IPP), the point where the satellite-receiver line of sight crosses a fixed altitude h_I , as illustrated in Figure 2.2 (Komjathy, 1997). Writing z' for the zenith angle of the ray at the IPP,

$$\sin z' = \frac{R_E}{R_E + h_I} \cos \varepsilon, \quad \text{VTEC} = \text{STEC} \cos z', \quad (2.13)$$

where ε is the satellite elevation angle and R_E is Earth's mean radius (6371 km). A shell height of $h_I = 350$ km is the standard choice for equatorial work; observations below 30° elevation are typically discarded to keep mapping errors small (Mannucci et al., 1999).




Figure 2.2: Geometry for the conversion of slant TEC to vertical TEC

2.2 TEC Variability at Equatorial and Low Latitudes

2.2.1 Solar and Geomagnetic Drivers

Electron density at any point in the ionosphere is set by a balance of production, loss and transport. The continuity equation for electron density states that balance compactly,

$$\frac{\partial N_e}{\partial t} = q - L(N_e) - \nabla \cdot (N_e \mathbf{v}), \quad (2.14)$$

where q is the photoionisation rate, L the recombination rate and $\mathbf{v}$ the plasma drift velocity (Rishbeth & Garriott, 1969; Schunk & Nagy, 2009). Every solar and geomagnetic driver enters through one of these three terms. Solar extreme ultraviolet (EUV) output controls q , geomagnetic activity acts mainly on the transport term and secondarily on L by changing the neutral composition. At low latitudes it is transport that dominates.

The geomagnetic field geometry is what makes low latitudes distinctive in that at the dip equator the field is horizontal, and because conductivity along a field line vastly exceeds that across it, field lines act as equipotentials and E-region dynamo electric fields map into the F region almost unchanged (Richmond, 1995; Rishbeth, 1997). Plasma then drifts at

$$\mathbf{v} = \frac{\mathbf{E} \times \mathbf{B}}{B^2}, \quad (2.15)$$

which for an eastward $\mathbf{E}$ field and a northward $\mathbf{B}$ field is vertically upward. In the E region the vertical Hall current cannot escape the layer, so a polarisation field builds up, raises the effective conductivity to the Cowling value and drives the equatorial electrojet (Kelley, 2009). The same zonal field produces the plasma fountain and the Equatorial Ionization Anomaly of Section 2.2.2. Crucially, lifting the F layer moves plasma into thinner molecular gas where recombination is slower, so a modest change in the zonal electric field produces a large change in TEC without any change in production (Fejer & Kelley, 1980).

Geomagnetic activity is simply solar wind energy transferred into the magnetosphere when the interplanetary field turns southward. Three indices are used throughout this work, each sampling a different current system and timescale. Kp is a three-hourly quasi-logarithmic index built from 13 subauroral observatories, ap is its linearised equivalent in nanotesla, and Dst is an hourly index from four low-latitude

magnetometers that measures the ring-current depression (Kelley, 2009). The disturbance itself reaches the equator by two competing routes. Prompt penetration electric fields arrive within minutes; under southward IMF they add an eastward field that strengthens the fountain and raises TEC. The disturbance dynamo acts hours later, when auroral Joule heating drives equatorward winds that generate a westward field and suppress the fountain (Bolaji et al., 2022; Fejer & Kelley, 1980). Storm heating also lowers the O/N₂ ratio at high latitudes, and where that air reaches lower latitudes recombination accelerates and TEC falls below the quiet-day level (Schunk & Nagy, 2009). These three responses differ in sign and in timescale, so equatorial TEC is not a single-valued function of K_p; it depends on storm phase and on recent history (Olowendo et al., 2016). That is the mechanism. What it produces in TEC measured from the ground over East Africa, and what that behaviour demands of a model, is taken up in Section 2.2.3.

2.2.2 Equatorial Ionization Anomaly

The Equatorial Ionization Anomaly (EIA) is the dominant structural feature of the low-latitude ionosphere. It forms from the daytime upward $\mathbf{E} \times \mathbf{B}$ drift at the magnetic equator, driven in turn by the eastward electric field of the equatorial electrojet (Kelley, 2009; Mitra, 1946). Plasma lifted at the equator then slides along magnetic field lines toward higher latitudes on both sides. Two TEC peaks form at roughly $\pm 15^\circ$ magnetic latitude, one on either side of the magnetic equator. VTEC in these crests typically exceeds the equatorial trough between them by 20–40 TECU (Kelley, 2009; Olowendo et al., 2016).

Uganda lies in the southern magnetic hemisphere, running from the trough side of the anomaly at its northern border to the southern crest at its southern border. The country therefore sits along the steepest part of the meridional TEC gradient, which is what makes the coarse cells of a global product a poor fit here: a $2.5^\circ \times 5^\circ$ CODE GIM cell straddles that gradient and returns an area-weighted mean that overshoots the column overhead of any single receiver inside it.

2.2.3 Storm-Time TEC Variations

The mechanisms are already discussed as Prompt penetration and the disturbance dynamo compete for control of the zonal electric field, and storm heating works separately on recombination through the O/N₂ ratio (Section 2.2.1). This section is

about what those mechanisms leave behind in the measurements, because it is the measured behaviour, not the mechanism, that a data-driven model has to reproduce.

The regional record is ordered in two ways at once. The first ordering is in time. Olwendo et al. (2016) followed East African TEC through several disturbed periods and found enhancement during the main phase giving way to depletion during recovery, so the sign of the departure from the quiet-day level depends on when in the storm the observation is taken as much as on how severe the storm is. Bolaji et al. (2022) reported the same ordering across the African sector for the St. Patrick's Day storms of March 2013 and 2015.

The second ordering is in latitude. Because the storm-time electric field acts on the fountain, the response at a station is set by its position relative to the EIA crest (Section 2.2.2) rather than by its distance from the geographic equator. Stations on the equatorward slope of the crest and stations sitting in the trough can therefore record departures of opposite sign during the same hours of the same storm. Uganda spans barely four degrees of latitude, yet that is enough to put stations on either side of the divide, which is exactly what makes a dense national network more informative during a storm than any single site can be (Section 4.1.3).

Taken together, the two orderings are what make storm-time TEC hard to model. An instantaneous Kp or Dst value says nothing about which phase of a storm produced it, so a model given only the current indices is asked to return two different answers to the same input. The usual remedy is to supply time history alongside the indices and to report storm and quiet performance separately rather than as a single average (Section 2.3.7.3).

2.3 Ionospheric Models

An ionospheric model is a computational description that returns electron density, or an integrated quantity such as TEC, for a stated position, time, and level of solar and geomagnetic activity. The need for model is due to receivers not being able to fully sample and represent the ionosphere and ionospheric delay corrections which are to be applied in real time which would require to you waiting for observations to occur, Models provide a better alternative by supplying values that measurements cannot.

Ionospheric models are conventionally grouped into three families (Bilitza et al.,

2022; Schunk & Nagy, 2009). *Physics-based* models integrate the coupled continuity, momentum and energy equations for the ionospheric plasma, driven by specified neutral winds, electric fields and solar irradiance. *Empirical* models fit analytical functions to archives of past observations and reproduce the average behaviour those archives contain, without solving any transport equation. *Semi-empirical* models sit between the two: the functional form of the electron density profile is taken from ionospheric theory, while the coefficients that anchor it are determined from measurements. The subsections that follow describe the models most widely used for TEC specification and correction, in increasing order of the amount of observational data each one consumes: the Klobuchar broadcast algorithm, IRI, NeQuick, the Global Ionosphere Maps, and the regional AfriTEC model. The mathematical machinery these models are built from is treated separately in Section 2.3.7, since the same technique often appears in several different models.

2.3.1 Physics-Based Models

Physics-based models such as SAMI2/SAMI3, the Thermosphere–Ionosphere Electrodynamics General Circulation Model and the IRAP Plasmasphere Ionosphere Model solve the plasma transport equations on a field-line or global grid (Schunk & Nagy, 2009). Because they represent the processes themselves, they can in principle reproduce conditions never previously observed, which no empirical fit can do. Their accuracy, however, is set by the accuracy of the drivers supplied to them. Over the equatorial and low-latitude sector the two drivers that matter most, the equatorial electric field and the thermospheric neutral wind, are the least well observed of all (Fejer & Kelley, 1980; Richmond, 1995), so the modelled EIA crests are frequently displaced in latitude or misestimated in strength. A single global run is also computationally expensive. For these reasons physics-based models are used mainly as research tools rather than as the correction layer inside a GNSS receiver, and they are not considered further in this study.

2.3.2 The Klobuchar Broadcast Model

The Klobuchar model (Klobuchar, 1987) is the ionospheric correction that every GPS satellite broadcasts and that essentially every single-frequency receiver applies by default. It represents the diurnal variation of vertical delay as a constant night-time

value plus a half-cosine centred on 14:00 local time,

$$I_v = \begin{cases} 5 \times 10^{-9} + A \cos\left(\frac{2\pi(t - 50400)}{P}\right), & |t - 50400| < P/4, \\ 5 \times 10^{-9}, & \text{otherwise,} \end{cases} \quad (2.16)$$

where t is local time at the ionospheric pierce point in seconds, and the amplitude A and period P are third-order polynomials in geomagnetic latitude whose eight coefficients α_n and β_n are uploaded to the satellites and transmitted in the navigation message. The vertical delay is then mapped to the slant direction by an obliquity factor. The model is constrained by design, utilizing only eight broadcast coefficients and operating completely independent of receiver observations removing about half of the root-mean-square ionospheric range error on a global average (Klobuchar, 1996). That target is met at mid-latitudes. At equatorial and low latitudes, where the actual delay is largest and where the EIA introduces spatial structure that a latitude polynomial cannot express, the residual error after applying Klobuchar remains substantial. This is precisely the population of users, and the region, that motivates the present work.

2.3.3 The International Reference Ionosphere

The International Reference Ionosphere (IRI) is a semi-empirical model maintained jointly by the Committee on Space Research (COSPAR) and the International Union of Radio Science (URSI). Established in the late 1970s and updated on a continuing cycle, it describes the ionosphere between roughly 50 km and 2000 km altitude (Bilitza et al., 2022). The functional forms of its electron density profile are physically motivated, while the coefficients are derived from a large archive of ground- and space-based measurements: ionosondes and digisondes for the bottomside profile, incoherent scatter radar for high-resolution electron and ion profiles, topside sounders, and in-situ satellite measurements of the plasma parameters. From these data the model generates monthly median values of electron density, ion composition, and electron and ion temperature, expressed as functions of geographic and geomagnetic position, altitude, local time, day of year, and solar and geomagnetic indices. Vertical TEC is obtained by integrating the modelled profile between two specified heights.

The current operational release, IRI-2020, carries an improved representation of the D and E regions and updated boundary conditions (Bilitza et al., 2022). IRI is the ISO standard for ionospheric background specification (ISO 16457) and is

used routinely in satellite orbit determination, high-frequency propagation planning, and the specification of the ionospheric background against which GNSS corrections are assessed. Its status as the accepted reference is the reason it is adopted as the benchmark for the model developed in this study.

2.3.4 The NeQuick Model

NeQuick (Nava et al., 2008) is a semi-empirical quick-run electron density model developed at the Abdus Salam International Centre for Theoretical Physics and the University of Graz. The bottomside profile is assembled from three semi-Epstein layers anchored to the E, F1 and F2 peaks, whose critical frequencies and peak heights come from the CCIR/ITU-R coefficient maps; the topside is a further semi-Epstein layer with an empirically determined, height-dependent thickness parameter. TEC along any ray path is obtained by numerical integration through the resulting three-dimensional density field, which is what makes the model fast enough for operational use. NeQuick is driven either by the solar radio flux F10.7 and sunspot number or, in the NeQuick-G variant adopted by Galileo for single-frequency correction, by an effective ionisation level computed from three coefficients broadcast in the navigation message.

Because the CCIR coefficient maps are themselves built from a mid-latitude-dominated ionosonde archive, NeQuick shares the weakness of IRI over the African equatorial sector. Twinomugisha et al. (2017) compared GPS-derived TEC from East African low-latitude stations, including Ugandan receivers, against NeQuick2 and found that the model reproduced the diurnal and seasonal shape of the observations but differed from them systematically in magnitude.

2.3.5 Global Ionosphere Maps

Global Ionosphere Maps (GIMs) are two-dimensional vertical TEC maps generated from the pseudorange and carrier-phase observations of the worldwide IGS receiver network (Hernández-Pajares et al., 2009; Schaer et al., 1998). The construction begins by treating the ionosphere as a thin shell at a fixed height, so that each observation constrains the vertical TEC at the ionospheric pierce point where the ray crosses the shell. The pierce point values from all stations are then fitted with a

spherical harmonic expansion in a Sun-fixed frame (Hernández-Pajares et al., 2009):

$$E(\beta, s) = \sum_{n=0}^{n_{max}} \sum_{m=0}^n \tilde{P}_{nm}(\sin \beta) (a_{nm} \cos ms + b_{nm} \sin ms), \quad (2.17)$$

where E is the vertical TEC, β is the geocentric latitude of the pierce point, s is its Sun-fixed longitude, $\tilde{P}_{nm}$ are the normalised associated Legendre polynomials of degree n and order m , and a_{nm} , b_{nm} are the coefficients estimated from the network. CODE uses $n_{max} = 15$. The satellite and receiver differential code biases are estimated simultaneously with the coefficients, since the two are inseparable in the pseudorange observable. The fitted coefficients are evaluated on a $2.5^\circ \times 5^\circ$ latitude–longitude grid at one-hour resolution and distributed in the IONEX format, which stores the gridded VTEC values alongside the bias estimates (Schaer et al., 1998).

Several IGS analysis centres publish their own GIM products. The maps produced by the Centre for Orbit Determination in Europe (CODE) at the Astronomical Institute of the University of Bern are among the most widely used, largely because they are updated daily and the archive reaches back to 1996 (Schaer et al., 1998). That archive spans more than two complete solar cycles, a record length no regional GNSS dataset over the African sector can match.

The accuracy of GIM products has been assessed against independent TEC measurements. Hernández-Pajares et al. (2009) reported a global RMSE of 2–4 TECU, rising to 10–20 TECU in the equatorial region under high solar activity. The degradation is a consequence of the grid spacing rather than of the estimation procedure. A single $2.5^\circ \times 5^\circ$ cell spans roughly 275 km in latitude and 550 km in longitude at the equator, so one cell can contain both the equatorial trough and part of the adjacent EIA crest. The value reported for that cell is an area average, not the TEC above any particular point on the ground, and over a country that straddles a steep latitudinal gradient the difference between the two is systematic rather than random. This is a limitation of spatial resolution, not of the underlying data, and it is the specific limitation that a regional correction stage is well placed to remove. The distinction matters for how the product is best used. The long-term solar-cycle and seasonal climatology carried in the archive is not in question; it is the regional scale that requires calibration against observations made on the ground.

2.3.6 Regional Models over Africa.

Global products are calibrated for global performance and their representation in the African sector is constrained due to under-representation of the receiver network in the continent. AfriTEC (Okoh et al., 2019) was built to address exactly that imbalance. It is a continental TEC model for Africa constructed from GPS-derived TEC recorded by GNSS receivers distributed across the continent, with solar activity, geomagnetic indices, and time variables as inputs; the fitting technique is a feedforward neural network, described in Section 2.3.7.1. Evaluated across the African regions tested, AfriTEC outperformed IRI, cutting VTEC prediction RMSE by roughly 40–60% relative to the IRI baseline.

Two features of that result are worth drawing out. The first is that a regionally trained model can outperform a global climatology over Africa using nothing more than the same solar, geomagnetic and temporal drivers, so the improvement comes from the training data rather than from any additional physical input. The second is that the largest residual errors remained in the equatorial sector, where the EIA gradient is steepest. The residual pattern is therefore a consequence of receiver density rather than of the fitting technique.

2.3.7 Methods of Modelling Ionospheric TEC

2.3.7.1 Neural Network Approaches

Feedforward artificial neural networks (ANNs), or multilayer perceptrons, were among the first machine learning methods applied to TEC prediction. Their tabular input format pairs naturally with the solar, geomagnetic and temporal variables that drive ionospheric variability, and the hidden layers can represent the non-linear interactions between those variables that a polynomial fit cannot.

Work in Africa began with Habarulema et al. (2009), who showed that ANNs trained on South African GPS data could outperform IRI climatology at low and mid-latitudes. Habarulema and McKinnell (2012) extended this by comparing backpropagation training algorithms over the same South African network and identifying which optimisation strategies produce the most accurate TEC estimates. Uwamahoro et al. (2018) subsequently compared neural network reconstruction against ionospheric tomography for storm-time TEC over Africa, and Ssessanga et al. (2024) applied a neural network to the prediction of low-latitude scintillation indices over the African

region. The continental-scale application of the same method is AfriTEC (Okoh et al., 2019), discussed in Section 2.3.6.

More recent deep learning work has targeted gridded map forecasting directly. Liu et al. (2022) forecast the IGS global VTEC maps 24 hours ahead using convolutional and recurrent layers, reporting RMSE of 1.5–3.0 TECU against the maps themselves. Luo et al. (2023) added a self-attention module to a ConvLSTM operating on the same class of global product and reported an RMSE improvement over Liu et al. (2022). Ren et al. (2022) compared encoder–decoder convolutional networks against graph neural networks for global TEC forecasting and found that architectures which represent spatial structure explicitly consistently outperformed flat multilayer perceptrons. In a review of machine-learning approaches to ionospheric modelling, Mao, Hernández-Pajares, and Soja (2025) noted a convergence towards hybrid designs in which a global learner supplies the large-scale background and a second stage corrects it locally. Two conclusions recur across this body of work. Spatial position is better supplied to a network as continuous latitude and longitude than as a categorical station index, since only the continuous form permits evaluation at locations absent from the training set; and a local correction applied on top of a global background generalises better than a single monolithic fit covering both scales at once.

2.3.7.2 Recurrent and Sequence Models

Long Short-Term Memory (LSTM) networks (Hochreiter & Schmidhuber, 1997) use gated memory cells to carry information across a sequence, which makes them a natural fit for TEC time series where the present state depends on the preceding hours. Jeong et al. (2024) compared LSTM and convolutional LSTM architectures for regional TEC prediction over Korea and reported RMSE of 2–4 TECU for one-hour-ahead forecasts, and Ren et al. (2022) found comparable behaviour for recurrent architectures at global scale, with both classes of model outperforming climatological baselines on geomagnetically quiet days.

The known weakness of sequence models in this application is generalisation beyond the range of conditions represented in training. A network that learns its hidden-state dynamics from one interval of the solar cycle has no mechanism to recognise that it is being evaluated under a different level of activity, and the error grows with the size of that mismatch; the difficulty is not specific to any one cycle but arises whenever the evaluation period lies outside the training distribution (Mao, Hernández-Pajares,

& Soja, 2025). Tree ensembles behave differently. A gradient boosted tree partitions the feature space into discrete regions, and explicit physical inputs such as F10.7, Kp and Dst place each prediction in the correct region directly, without relying on a learned sequence memory to infer the activity level (Chen & Guestrin, 2016). The trade-off is that the tree ensemble extrapolates poorly outside the range of the training targets. Either family therefore imposes the same practical requirement, which is a training record long enough to span the range of solar activity over which the model will be used.

2.3.7.3 Gradient Boosting and Tree-Based Methods

Gradient boosted decision trees grow an ensemble of regression trees in sequence, each new tree fitted to the residual errors of those already in the ensemble (Friedman, 2001). XGBoost (Chen & Guestrin, 2016) is a fast, regularised implementation that tolerates outliers, handles missing values without imputation, and requires no feature scaling. It also supports SHAP-based attribution (Lundberg & Lee, 2017), which decomposes an individual prediction into the contribution of each input feature. For a study in which the learned behaviour has to be checked against known ionospheric physics rather than accepted on the strength of its error statistics, that property is a requirement rather than a convenience.

For the East African sector specifically, Nigusie et al. (2024) applied gradient boosting to TEC over Ethiopia, a region whose ionospheric conditions resemble Uganda's, and reached $R^2 > 0.97$ on the test set, with a stacked Ridge regression meta-learner improving accuracy further. Sisay et al. (2024) evaluated comparable machine-learning ensembles over Ethiopia during geomagnetic storms, reporting $R^2 > 0.92$ through the main and recovery phases while IRI-2020 degraded substantially under the same conditions. These two studies are the nearest published precedent for gradient-boosting TEC modelling in the East African sector, and both validate against GIM-derived TEC rather than against a dense local GPS network.

2.3.7.4 Ensemble and Stacking Methods

Stacked generalisation (Wolpert, 1992) feeds the predictions of one or more base models into a second-level model that learns how best to combine or correct them. The practical requirement (Breiman, 1996) is that the base-model predictions used to train the second level be generated out of fold, so that the second-level model

never sees a prediction made on data the base model was trained on. Without that precaution the second level learns the base model’s memorised training error rather than its genuine residual structure.

Applied to TEC, ensemble averaging of several machine learning models trained on global GPS data has produced RMSE improvements of 10–20% over the individual members (Zhou et al., 2024), and stacking XGBoost with a linear meta-learner delivered a further gain of about 5% over single-model XGBoost in the Ethiopian study of Nigusie et al. (2024). Regional real-time schemes have taken a similar route: Sun et al. (2025) combined a learned residual network with a regional GNSS network over China to refine a global background in near real time.

What these published schemes have in common is that both levels of the ensemble are trained on the same class of data, so the second level refines the first without adding information of a different kind. The alternative arrangement, in which the two levels draw on records that differ in length and in spatial density, so that each level is confined to what its own data can support, has not been reported for the African sector.

2.4 Summary

Ionospheric Total Electron Content (TEC) modeling over Uganda and the wider East African region faces distinct spatial and validation constraints across current literature. The Standard climatological models, including Klobuchar, IRI, and NeQuick, map broad seasonal trends over Uganda but systematically depart from local amplitude observations, while global products like CODE GIM lack the spatial grid resolution required to resolve the fine structure of the Equatorial Ionization Anomaly (EIA) trough and crest. Regional observational studies have historically been with sparse localized data, past TEC and scintillation research (Amabayo et al., 2014; Andima et al., 2019; Cele et al., 2023; D’ujanga et al., 2012; Nibigira et al., 2019; Twinomugisha et al., 2017) relied on single receivers or widely spaced regional networks where station gaps of several hundred kilometers are too coarse to sample localized latitudinal gradients. Finally, while recent regional machine learning models (Nigusie et al., 2024; Sisay et al., 2024) demonstrate the utility of ensemble techniques, their validation against GIM-derived VTEC limits their capacity to reflect actual ground-based anomalies. Consequently, the high-density spatial observations available from dense local networks, such as the 40-station Ugandan CORS network (Ogwang, 2020), remain unexploited within existing predictive modeling frameworks.

Chapter 3

Data and Methods

This chapter covers the data sources, data preprocessing steps, model architectures, and evaluation metrics used in this study. Data sources are described first from Sections 3.2 to 3.3, followed by the input features from Section 3.5, the model architectures from Sections 3.6 to 3.7, and the evaluation metrics used in Section 3.8.

3.1 Study Area and Regional Context

The processes set out in Section 2.2 are not African processes. The equatorial fountain, the anomaly it builds, the pre-reversal enhancement at sunset and the storm-time competition between penetration and disturbance dynamo fields operate in the same way in every longitude sector, and any study of the ionosphere over Uganda inherits that physics whole. What changes from one sector to the next is not the mechanism but the geometry of the magnetic field through which the mechanism acts. Two geometric quantities carry almost all of that difference: the magnetic declination, and the offset between the geographic and the dip equator. This section isolates them before turning to what has been measured over the region.

3.1.1 Longitudinal Modulation by Field Geometry

Thermospheric winds move F-region plasma only along the magnetic field. A horizontal neutral wind with geographic northward and eastward components U_n and U_e therefore acts on the plasma through its component in the magnetic meridian plane,

$$U_m = U_n \cos D + U_e \sin D, \quad (3.1)$$

where D is the magnetic declination, reckoned positive eastward. That component is projected onto the vertical by the dip angle I , giving a wind-driven vertical plasma drift of magnitude

$$w = U_m \sin I \cos I, \quad (3.2)$$

an equatorward U_m raising the layer into a region of slower recombination and a poleward U_m lowering it (Kelley, 2009; Rishbeth, 1997).

Equation 3.1 is where declination earns its importance. The zonal thermospheric wind at low latitudes is strong, reaching well over 100 m s^{-1} eastward in the evening sector (Kelley, 2009), but it enters the plasma transport only through $\sin D$. Where declination is large, the zonal wind becomes an effective driver of field-aligned transport; where it is near zero the zonal wind is almost invisible to the plasma and the meridional wind does the work alone. Declination also fixes the season at which the sunset terminator aligns with the magnetic meridian, and the pre-reversal enhancement is strongest when that alignment occurs, because the E-region conductivity then collapses simultaneously along the whole flux tube (Abdu, 2016). The result is a longitudinal and seasonal modulation of anomaly strength and of post-sunset irregularity occurrence that is produced entirely by field geometry, with no difference in the underlying physics.

The African sector sits at one extreme of that modulation. Evaluated with IGRF-14 at 350 km for epoch 2025.5, declination across the 40 Uganda CORS stations runs only from $+0.4^\circ$ to $+1.4^\circ$ east, so $|\sin D| \leq 0.024$ and under three per cent of the zonal wind projects into the magnetic meridian. The same calculation in the American sector gives $D \approx -19^\circ$ near 50° W , where $|\sin D| \approx 0.33$ and roughly a third of the zonal wind contributes. Uganda therefore lies close to the idealised case in which the meridional wind alone drives the wind-induced vertical drift, and in which the terminator aligns with the magnetic meridian near the equinoxes rather than at a solstice. That is consistent with the pronounced equinoctial maxima in TEC and in post-sunset irregularity occurrence reported for East Africa (Amabayo et al., 2014; Olwendo et al., 2016), as against the December-solstice maximum characteristic of the American sector (Abdu, 2016).

The second geometric quantity varies steeply inside the country itself. The dip angle at the CORS stations runs from -11.0° at ABON in the north to -23.6° at KABA in the south, so the projection factor $\sin I \cos I$ in Equation 3.2 nearly doubles across the network, from 0.187 to 0.367. The same meridional wind produces almost twice the vertical drift at the southern stations as at the northern ones. Taken with the displacement of the dip equator described in Section 3.1.2, this is the sense in which the East African sector is distinctive. The physics is the common low-latitude physics; what changes is the geometry that converts it into the TEC a receiver on the ground records, and over Uganda that geometry changes appreciably from one side of the country to the other.

3.1.2 Geographic and Ionospheric Context

The offset between the two equators is what makes that geometry unusual. Across the Ugandan longitude band (29.5° – 35° E) the dip equator is displaced well into the Northern Hemisphere, running near 8° – 9° N geographic latitude through South Sudan and southern Ethiopia. East Africa therefore straddles both magnetic hemispheres even though most of it lies south of the geographic equator.

Uganda extends from about 1.5° S to 4.2° N geographic, which places the whole country in the *southern* magnetic hemisphere. In the quasi-dipole coordinates used throughout this work (Section 3.5), the national territory spans roughly 5° S at its northern border to 11° S at its southern border, and the 40 CORS stations cover -5.1° (ABON) to -10.9° (KABA). The anomaly crest relevant to Uganda is thus the southern one. Taking the crest at 10° – 15° south magnetic latitude, it maps in this longitude sector to roughly 1° S– 5.5° S geographic, that is, southern Uganda through northern Tanzania. The northern crest lies near 17° – 22° N geographic, over Sudan and Egypt, far outside the study area.

The consequence is a steep gradient inside a small country. Uganda spans barely 5.7° of geographic latitude but close to 6° of magnetic latitude, running from the trough side of the anomaly in the north to the crest side in the south. Northern and southern stations therefore sample different parts of the EIA structure at the same instant, and the ordering between them is not fixed: it follows the development of the fountain through the day (Section 4.1.1). Geographic latitude on its own is a poor proxy for this behaviour in the East African sector; it is magnetic, not geographic, latitude that orders the observations.

3.1.3 Previous Studies

Ugandan TEC research has been built on a small number of receivers. D’ujanga et al. (2012) used GPS observations from the equatorial anomaly region to characterise TEC variability and to compare it against IRI, reporting the large positive IRI biases that are characteristic of African low-latitude stations. Amabayo et al. (2014) worked with the two Ugandan stations that have the longest records, Mbarara (MBAR) and Entebbe (EBBE), and produced a climatology of ionospheric scintillation occurrence together with the TEC trend over the region, establishing the local pattern of post-sunset irregularity activity. Twinomugisha et al. (2017) used GPS observations from Ugandan and neighbouring East African low-latitude stations to derive TEC and

compared it against the NeQuick2 model, finding that the model reproduced the shape of the variation but not its amplitude. Nibigira et al. (2019) analysed diurnal, seasonal and annual TEC variation at MBAR (0.60° S, 30.74° E) during 2018, in the descending phase of Solar Cycle 24, recovering the equinoctial maximum and June-solstice minimum. These studies establish the temporal climatology of the Ugandan ionosphere. What one or two receivers cannot do is resolve space: none of them samples the trough-to-crest gradient that Section 3.1.2 shows Uganda spans.

Regional studies widened the coverage at the cost of density. Olwendo et al. (2016) combined GPS data from Kenya, Uganda, Tanzania and Ethiopia to establish the East African TEC climatology, documenting the semi-annual equinoctial peaks, the EIA latitudinal gradient, and the storm-time response of the region. The same study reported strong TEC gradients within the latitude range Uganda spans. Mungufeni et al. (2016) built a statistical TEC model for the equatorial African region from GPS observations during a year of low solar activity, and Tariku (2015) documented the seasonal TEC patterns of the neighbouring Ethiopian sector. Andima et al. (2019) went further and produced a model rather than a description, decomposing GPS TEC over the African low-latitude region into empirical orthogonal functions and expressing the mode amplitudes as functions of solar activity, season and local time; the receiver set included the Ugandan stations. Each of these studies is limited by receiver spacing of several hundred kilometres, which is coarse relative to the gradient being measured.

Cele et al. (2023) addressed that limitation directly. Noting that sparse GNSS coverage leaves spatio-temporal gaps across East Africa, they constructed regional TEC maps by kriging interpolation and characterised the low-latitude ionosphere over the region for quiet and disturbed conditions during 2014 and 2018, recovering the 10:00–13:00 UT daily maximum and the March-equinox maximum with a June-solstice minimum. Kriging, however, is a geostatistical interpolator: it fills gaps between receivers but carries no ionospheric physics, cannot extrapolate to conditions outside the observations, and its resolution remains capped by station density.

Storm-time work over Africa has generally relied on either tomography or neural networks applied to sparse regional networks. Uwamahoro et al. (2018) compared ionospheric tomography against artificial neural networks for reconstructing storm-time TEC over the African region, and Habarulema et al. (2013) assembled multi-event composites of the African storm response. Further south, Ssessanga et al. (2019) applied a 4D variational scheme over South Africa and obtained accuracy comparable to global GIMs from regional GPS observations alone, which establishes the principle

that a dense, well-processed regional network can match a coarse global product.

What has not been available until now is the density. The Uganda CORS network, whose data quality was assessed by Ogwang (2020), provides 40 continuously operating stations across the country, an order of magnitude more than any dataset used in the studies above.

3.2 Data Sources

3.2.1 CODE Global Ionosphere Maps

The data for Stage 1 training was got from the CODE Global Ionosphere Maps (CODE GIM) in IONEX format which is produced by the Astronomical Institute of the University of Berne (AIUB) on a daily basis and distributes freely from the CDDIS archive at NASA and from the AIUB FTP server (<ftp.aiub.unibe.ch>). Each daily file has a gridded hourly VTEC resolution of $2.5^\circ \times 5^\circ$ global on a global scale.

These files were downloaded for January 2000 to February 2026 which gave 26 years of continuous data at 1-hour temporal interval spanning the tail of Solar Cycle 23, all of Solar Cycle 24, and the ascending-to-maximum phase of Solar Cycle 25.

3.2.2 Uganda CORS GPS Network

The distribution of the CORS GPS Network is shown in figure 3.1 which clearly shows a dense network of local receivers over Uganda.

The GPS rinex observation files from all the 40 Uganda CORS stations as represented in table 3.1 were processed using GPS_Gopi v3.0 (Seemala, 2023), which is an open-source tool that extracts differential pseudorange TEC from each satellite-receiver pair and converts it to VTEC using a thin-shell mapping function at 350 km. The daily files are in a .cmn format which is the main output per station file and contained 30-second VTEC observations at the IPP of each tracked satellite together with elevation angle, azimuth, satellite PRN, and local time.

The slant TEC recovered from the phase-levelled pseudorange combination that carry three instrumental biases namely the satellite differential code bias, the receiver




Figure 3.1: Map showing the 40-station Continuously Operating Reference Station (CORS) GPS network used in this Study. The network spans from 1.27° S (KABA, south-west) to 3.51° N (ABON, north-east), covering the full north-to-south extent of the country.

differential code bias, and the receiver inter-channel bias (Ciraolo et al., 2007; Seemala, 2023). All the three were removed before any VTEC value entered the dataset used here. Satellite biases were applied from the monthly P1C1 and P1P2 DCB files published by CODE at the Astronomical Institute of the University of Bern, the same source as the IONEX products of Section 3.2.1, the receiver bias and inter-channel bias were estimated by the processing tool for each station and day from the observations themselves (Seemala, 2023). Bias removal was therefore a preprocessing step which was completed before the hourly aggregation described below, and not something left to the model. The per-station term introduced in Section 3.6.2.3 accounts for that residual, not for the DCB itself.

CORS data processing began from 11,570 daily station files, which yielded 369,452,695 individual satellite-epoch observations before quality control. Three quality control filters were then applied in sequence. The elevation angle cut ($\geq 30^\circ$), to limit multipath errors at low elevations (Mannucci et al., 1999), removed 213,859,434 epochs (57.9% of the raw total). The VTEC range filter, $(0, 120]$ TECU, to exclude negative fill values and unphysical single-receiver spikes, removed a further 18,927,207 epochs (12.2% of the epochs remaining after the elevation cut). This 120 TECU ceiling is specific to the raw 30-second CORS epochs and is not applied to the CODE GIM data used to train Stage 1 (Section 3.2.1), where VTEC legitimately reaches 146.06 TECU under solar-maximum conditions in fewer than 0.1% of grid-hours. A total of 136,666,054 satellite epochs (37.0% of the raw total) survived both filters. Each hourly aggregation bin was then required to contain at least 10 valid 30-second epochs for the hourly median to be considered reliable; 1,085 of the resulting 265,349 station-hours (0.4%) failed this criterion and were excluded from the ML-ready record.

After quality control, the 30-second observations were aggregated into hourly median VTEC per station and the final dataset has 233,698 hourly VTEC observations across the 40 stations (Table 3.1) from August 2025 to April 2026. Of these, 159,787 were used for training (August 2025 to January 2026), 25,793 (all of February 2026) were set aside for validation, and 48,118 as the test period (March to April 2026). Of the 25,793 raw validation hours, 24,946 survive into the Stage 2 ANN’s training dataframe once the 24-hour lag-feature exclusion at the split boundary (Section 3.3.3) and the requirement for a valid Stage 1 background at each row are applied. The other 847 rows are excluded from model fitting and evaluation for that reason, though they remain part of the raw 233,698-hour CORS record.

Table 3.1: Geographic coordinates of the 40 Uganda CORS stations used in this study. Station codes follow the Uganda National Geodetic Control Network naming convention.

Code	Name	Lat (°)	Lon (°)	Code	Name	Lat (°)	Lon (°)
ABIM	Abim	2.59	33.67	MBAL	Mbale	1.07	34.17
ABON	Aboni	3.51	34.14	MBRA	Mbarara	-0.61	30.64
ADJM	Adjumani	3.37	31.79	MRT0	Moroto	2.53	34.66
APAC	Apac	1.97	32.53	MSKA	Masaka	-0.34	31.72
ARUA	Arua	3.02	30.91	MSND	Masindi	1.69	31.72
BOMB	Bombo	0.59	32.53	MTYN	Mutyana	0.41	32.06
BUGI	Bugembe	0.57	33.74	NAKA	Nakaseke	1.86	34.73
ENTB	Entebbe	0.06	32.48	NAPA	Napak	2.12	34.22
FPRT	Fort Portal	0.65	30.30	OCHO	Ocholero	1.64	33.04
GOLA	Gola	1.31	32.46	OROM	Orom	3.39	33.51
GULU	Gulu	2.78	32.30	RAKA	Rakai	-0.71	31.40
IBAN	Ibanda	-0.13	30.49	RUKU	Rukungiri	-0.78	29.93
JING	Jinja	0.42	33.21	SEMB	Sembabule	-0.08	31.46
KABA	Kabale	-1.27	29.98	SRTI	Soroti	1.72	33.62
KALA	Kalangala	-0.32	32.29	TIDO	Tidonoi	3.04	34.13
KAMU	Kamuli	0.96	33.12	WACH	Wachewa	2.47	31.48
KASE	Kasese	0.18	30.08	YUMB	Yumbe	3.45	31.25
KBGA	Kiboga	0.91	31.78	LIRA	Lira	2.25	32.90
KBLE	Kibaale	0.80	31.08	KTGM	Kitgum	3.31	32.89
KIKU	Kikuube	1.44	31.04	KYEG	Kyegegwa	0.49	31.06

3.2.3 Solar and Geomagnetic Indices

Solar and geomagnetic indices were obtained from NASA’s OMNIWeb OMNI2 data service (Space Physics Data Facility) as an hourly time series for the full 2000–2026 study period: F10.7, Kp and Dst. Their native cadences differ. Dst is a genuinely hourly index. Kp is a 3-hourly index, held constant across each 3-hour window to populate the hourly series. F10.7 is the OMNI2 adjusted-to-1-AU daily solar radio flux at 10.7 cm, held constant across all 24 hours of its day.

3.3 Data Preprocessing

3.3.1 GIM Spatial and Temporal Interpolation

The processing of GIM dataset was done in six chronological steps.

The daily IONEX files were first decompressed and parsed.

The native CODE GIM grid ($2.5^\circ \times 5^\circ$ globally) was then spatially interpolated to a $1^\circ \times 1^\circ$ Uganda grid with cell centres at latitudes $[-0.5^\circ, 0.5^\circ, 1.5^\circ, 2.5^\circ, 3.5^\circ, 4.5^\circ]$ and longitudes $[29.5^\circ, 30.5^\circ, 31.5^\circ, 32.5^\circ, 33.5^\circ, 34.5^\circ]$ using a bicubic spline method. This finer grid does not add independent spatial information beyond the native $2.5^\circ \times 5^\circ$ product. It resamples the field onto a denser grid so that the bilinear interpolation to each CORS station’s location (Section 3.6.1) has more nearby grid points to interpolate between.

For IONEX files from before 27 November 2022, the temporal distribution for the CODE GIM was at 2-hourly cadence (13 map epochs per day), which were temporally interpolated to 1-hourly resolution (25 epochs per day) using cubic spline interpolation along the time axis at each grid point (falling back to linear interpolation where fewer than four valid epochs remained after NaN masking). Files from 27 November 2022 onward were already natively 1-hourly and required no temporal interpolation.

The VTEC grid point in the GIM data was assigned a 1 hour UT timestamp with geomagnetic and geographic coordinates (quasi-dipole latitude and longitude) computed at each grid point (Richmond, 1995).

For effective coordinate transformation to geomagnetic coordinates, the International Geomagnetic Reference Field (IGRF) was used. For instance, data extended up

to December 2025 the IGRF-13 coefficients were used (valid through 2025) and for the 2026 data these were replaced by IGRF-14 (released December 2024, valid from 2025). The geomagnetic field difference between IGRF-13 and IGRF-14 at low latitudes is below 0.1° in QD coordinates, which is negligible relative to the TEC signal. Within each calendar year, the IGRF coefficient set is evaluated once, at a single representative epoch for that year, rather than being re-evaluated at every individual timestamp; the sub-annual secular drift of the geomagnetic field is negligible at the resolution used here.

Cyclic encodings and lag features (Section 3.5) were added last. The pipeline was then closed with a strict chronological split with the train windows from 2000 to 2022, validation from 2023 to 2024, and test from (1 January 2025 to 23 April 2026. Any other GIM archival gaps were excluded from the training windows for consistency and reliability of training data.

3.3.2 CORS Data Assembly

The processing steps for the CORS data is as it appears in Section 3.2.2. The hourly CORS data was merged with OMNIWeb indices and Stage 1 XGBoost predictions then interpolated to the stations latitude and longitude. The data also follows a strict chronological split with no overlaps as follow: training window is from 7 August 2025 to 31 January 2026 which is approximately 6 months, validation windows was from 1 February 2026 to 28 February 2026, and lastly test windows was from 1 March 2026 to 23 April 2026. The training window covers the August to September active season, the December to January transition and the start of the equinoctial enhancement so the ANN sees a broad range of seasonal and geomagnetic conditions.

To identify the geomagnetically disturbed periods, a threshold was used as a standard threshold of $K_p \geq 4$. Therefore the disturbed hours made up 15.2% of the March to April 2026 test set and this same threshold drove the storm/quiet stratification in Section 4.3.3.

Figure 3.2 clearly consolidates both splits on a single timeline against the daily F10.7 record over 2000–2026 where the GIM training window spans Solar Cycles 23 and 24 plus the ascending phase of Solar Cycle 25 while the GIM and CORS test windows fall shortly on the Solar Cycle 25 maximum, including the November 2025 storm ($D_{\text{st}_{\text{min}}} = -217 \text{ nT}$). The above visual gap between the long GIM splits and the compact six-month CORS training window was why the two-stage architecture was

particularity needed.




Figure 3.2: Temporal division of the GIM and CORS datasets into training, validation, and test sets, overlaid on the daily F10.7 solar flux record from 2000 to 2026.

3.3.3 Prevention of Data Leakage

Leakage was controlled by adding the lag features. These lag features were exclusively controlled by assigning the first 24 hours of data from each split a NaN value which avoids future information from contaminating the data and this was also excluded from the training and validation data.

The deliberate chronological split (Aug 2025–Jan 2026 / Feb 2026 / Mar–Apr 2026) does not produce identical distributions of solar zenith angle, season, or geomagnetic activity across the three sets which would mean that there will not be a data mixing like in random splitting.

3.4 TEC Characterisation Methodology

The diurnal variability of TEC was computed by obtaining the median VTEC across all days in the entire CORS data, calculated for each hour. The network median is the median of all the 40 stations for each hour. Three stations, ABON (3.51°N, northernmost), APAC (1.97°N, central), and KABA (−1.27°N, southernmost) were selected to represent the network’s latitudinal extent alongside the median.

Seasonal variability was computed as the monthly mean VTEC from the CODE

GIM data, averaged over the 1° latitude by 1° longitude Uganda grid, across the full 2000–2026 study period.

The November 2025 storm was selected as a case study because it was the most geomagnetically intense event within the CORS observation window, with its main phase defined as 12 November 2025. Geomagnetically disturbed hours elsewhere in this study use the Kp greater than 4 threshold.

3.5 Feature Engineering

An initial candidate set of seventeen engineered features was derived from the raw data. SHAP feature-importance analysis on the Stage 1 XGBoost model (Section 3.6.1) ranked `ap`, geographic `lat`, and `lon` last of the seventeen – redundant with `dst/kp` and `geomag_lat/geomag_lon` respectively – and the daily `f107_81d` running mean was found redundant with the daily `f107` value itself. All four were dropped with no meaningful change in test RMSE. The resulting thirteen features (Table 3.2) formed the final input vector for the Stage 1 XGBoost model. These selected features contained most of the physical drivers for TEC variability which can be captured from the data with no ad-hoc predictors.

Table 3.2: Model input features for the Stage 1 XGBoost model with the target being VTEC measured in TECU

Group	Feature	Description
Solar	f107	F10.7 solar flux (sfu)
Geomagnetic	kp	Kp planetary index
	dst	Dst storm index (nT)
Time	sin_hour	$\sin(2\pi \times \text{UT}/24)$
	cos_hour	$\cos(2\pi \times \text{UT}/24)$
	sin_doy	$\sin(2\pi \times \text{DOY}/365.25)$
	cos_doy	$\cos(2\pi \times \text{DOY}/365.25)$
Geometry	sza_deg	Solar zenith angle (degrees)
	geomag_lat	Quasi-dipole latitude (degrees)
	geomag_lon	Quasi-dipole longitude (degrees)
TEC lags	vtec_lag1h	VTEC at UT – 1 hour (TECU)
	vtec_lag2h	VTEC at UT – 2 hours (TECU)
	vtec_lag24h	VTEC at same hour previous day (TECU)

The (sin, cos) pairs seen are cyclic encodings representing the day of the year and hours in such a way to preserve a continuous flow of data to the night and also through out the year. this is a common practise in designing regression models

(Nigusie et al., 2024). The solar zenith angle that is a function of time, latitude and longitude was added to the input features instead of internally computing it. The geomagnetic latitude and longitudes included here were calculated using quasi-dipole(QD) coordinates with IGRF coefficient at an ionospheric thin-shell height of 350 km at the apex of magnetic field lines as used by (Richmond, 1995). In order to capture short term TEC changes and daily periodicity, the 1h, 2h and 24h auto-regressive features were added to the input features. These will help in reducing the prediction error thereby recording the state of the ionosphere just immediately before each prediction time.

3.6 Machine Learning Models

3.6.1 Stage 1: XGBoost Gradient Boosted Tree Model

The first model is the XGBoost model which is a gradient boosting tree based model that was trained on hourly VTEC values from the CODE GIM dataset with the 17 input features already discussed in the previous Section 3.5 and mainly bounded to Uganda gridded dimensions.

3.6.1.1 Model Architecture

The XGBoost model is a collection of decision tree which are referred to as an ensemble and as the trees grow, the regularization function is to keep decreasing until the desired outcome is reached. this function can be represented as:

$$\mathcal{L} = \sum_{i=1}^N l(y_i, \hat{y}_i) + \sum_{k=1}^K \Omega(f_k), \quad (3.3)$$

where l is the squared-error loss, $\hat{y}_i = \sum_k \eta f_k(\mathbf{x}_i)$ is the total ensemble prediction with learning rate η applied to each tree's contribution (matching Eq. 3.6 and Table 3.3), and $\Omega(f_k) = \gamma T + \frac{1}{2} \lambda \|\mathbf{w}\|^2$ is the term which reduces how complex a tree can grow into by changing the number of leaves T and the leaf weight norms with strength λ . In this plan, a penalty with strength α was also applied. A second order Taylor approximation was used to periodically fit each successive tree which helped in correctly determining the weights of the leaf which are close to the exact values to be predicted.

Figure 3.3 shows the overall structure of the XGBoost model with the additive ensemble parts. Each tree is to be split into the root and an internal node which is per different combination of features. This figure shows how the XGBoost model is structured.




Figure 3.3: Stage 1 XGBoost gradient-boosted tree ensemble.

3.6.1.2 Hyperparameters

Rather than an exhaustive grid, three candidate configurations were evaluated by 5-fold time-series cross-validation (24-hour gap between folds) on the 2020–2022 subset of the training data: (i) $\{n_{\text{estimators}} = 500, \text{depth} = 6, \eta = 0.05, \text{subsample} = 0.8, \text{colsample} = 0.8, \text{min_child_weight} = 5\}$; (ii) $\{800, 6, 0.05, 0.8, 0.7, 5\}$; and (iii) $\{500, 6, 0.1, 0.9, 0.8, 3\}$. The configuration with the lowest cross-validation RMSE was retained (Table 3.3), and early stopping (patience 50 rounds) on the 2023–2024 validation split then set the effective number of trees used at inference.

Table 3.3: The final hyperparameter configuration for the Stage 1 XGBoost model.

Hyperparameter	Value
Number of estimators	500
Maximum tree depth	6
Learning rate (η)	0.10
Row subsampling	0.90
Column subsampling (per tree)	0.80
Minimum child weight	3
Objective	reg:squarederror
Early stopping rounds	50

On performing the five fold cross validation, the out of fold (OOF) was determined for the full training windows. This implied that each fold trained on 80% of the chronologically organized training data and validated on the remaining 20%, which shows no overlaps between folds. The predictions in this case act as the Stage 1 background for CORS station locations during the model's training period and finally the Stage 1 test-set predictions (2025-01-01 – 2026-04-23) come from an ensemble average of the five fold models.

The instant predictions from the stage 1 XGBoost were therefore interpolated using bilinear interpolation from a $1^\circ \times 1^\circ$ Uganda GIM grid. The interpolated values form the Stage 1 background vector $\hat{y}_{\text{GIM}}$ for each station-hour, which then feeds the Stage 2 ANN as an input feature.

3.6.2 Stage 2: Artificial Neural Network Correction Model

Stage 2 in the modelling process is a feedforward Artificial Neural Network that is the local correction model since it was not directly predicting the TEC in Uganda. The stipulated learned corrections are the difference between the Stage 1 background predictions and the ground true value GPS VTEC, with the per-station bias term b_s (Section 3.6.2.3) removed before training, and this particular mapping is represented by:

$$\Delta(s, t) = [y_{\text{GPS}}(s, t) - \hat{y}_{\text{GIM}}(s, t)] - b_s = f_{\text{ANN}}(\mathbf{x}(s, t)) + \varepsilon, \quad (3.4)$$

where s are the station indexes, t as time indexes, and $\mathbf{x}$ is the representation of the extended feature vector described below. Since b_s was removed from the target before training, it is added back at prediction time. The final ensemble prediction is thus represented by the equation:

$$\hat{y}_{\text{final}}(s, t) = \hat{y}_{\text{GIM}}(s, t) + f_{\text{ANN}}(\mathbf{x}(s, t)) + b_s, \quad (3.5)$$

where b_s represent all the bias which have been learnt per the ground truth stations and this is estimated from the training set residues.

The Artificial Neural Network continuously receives latitude and longitude which helps in generalization and the correction being deliberately split into learned term and a fixed term could as well also contribute to the good performance of the model and thus generalizability.

The correction is deliberately split into a learned term and a fixed term. The

network receives continuous latitude and longitude, so in principle it could represent a per-station offset itself; the reason for not asking it to is generalisation. A network that absorbed each station’s offset into its weights would be fitting 40 site-specific constants through a spatial function, and those constants would then be interpolated to any new location as though they were part of the ionospheric field. Holding them in a separate additive term keeps f_{ANN} a function of the geophysical inputs alone, so that it can be evaluated at locations with no station, while b_s carries the site-specific part that has no meaning away from the receiver. For prediction at an arbitrary point the global mean bias replaces b_s , which is only defensible because the two are separated in the first place. The cost of split is that b_s is estimated on the training window and held fixed thereafter, so a change in a station’s residual offset between training and test appears directly in the final error.

3.6.2.1 Input Features and Architecture

For the Stage 2 ANN model, the input features used were f107, kp, dst, d_dst, sin_hour/cos_hour, sin_doy/cos_doy, sin_month/cos_month, ionospheric pierce-point latitude (lat_ipp) and longitude (lon_ipp), Stage 1 background(the interpolated XGBoost prediction $\hat{y}_{\text{GIM}}$) and residual lags of 1 h, 2 h, 3 h, 6 h and 24 h. The Stage 1 solar zenith angle, quasi-dipole coordinates, and GIM VTEC lags are replaced in Stage 2 by the monthly cyclic encodings, the Dst rate of change, and lags of the CORS GPS residual.

The ANN structure consists of three hidden layers starting with the first layer that applied Linear(18 $\rightarrow$ 64) followed by Layer Normalisation and ReLU. The second and third layers apply Linear(64 $\rightarrow$ 64) and Linear(64 $\rightarrow$ 32) respectively with each followed by a normalization layer, ReLU activation, and Dropout(0.3) for regularisation. The final result or output layer is a single Linear(32 $\rightarrow$ 1) neuron with no activation, that predicts the residual Δ . The reason for including a Layer Normalisation is due to unreliability of batch normalization since the observations in questions were from different stations and times. this was effectively observed by (Ba et al., 2016).

3.6.2.2 Training Configuration

The huber loss function was used with $\delta = 3.0$ TECU considered due to its ability to capture outliers which would have escaped the quality control set earlier. These

outliers can include bubbles and cycle slips.

Models were trained with Adam optimizer (Kingma & Ba, 2015) with an initial learning rate of 10^{-3} . To improve convergence, the learning rate was decayed following a cosine annealing schedule (Loshchilov & Hutter, 2017) with a over 100 epoch and with a minimum value of 10^{-5} . The batch size used was 256.

The highest training epoch used was 100 and an early stopping was applied with patience at 30 epoch on the validation set to get the best performance expected from the model.

Due to the 6 months of CORS training data as compared to the vast 26 year GIM data, the architecture and the optimizers were designed prior which is what most established regressional models have been implementing as well documented in (Goodfellow et al., 2016; Kingma & Ba, 2015). The only parameter therefore which was chosen by the model was the early stopping which is done as the model is training and also validating with the February 2026 validation month.

3.6.2.3 Per-Station Bias Correction

Before training the Stage 2 network, the per-station systematic bias term was calculated which in principle is different from the DCB errors which were already handled in the data preprocessing steps before the data was used for modelling. This per-station systematic bias, denoted by b_s , is computed as the mean of the Stage 1 residual, $y_{\text{GPS}}(s, t) - \hat{y}_{\text{GIM}}(s, t)$, over the training window only, per station, then subtracted from the target before the network is trained (Eq. 3.4) and added back to the network's output when reconstructing the final prediction (Eq. 3.5). For the locations which are outside the station location, a global bias term which was calculated as the mean of all the per-station biases is used. This ensures that the biases obtained are handled and the final ensemble performance is free from residual calibration errors, local multi-path at the antenna or any site specific errors at the true pierce-point thereby making it a zero mean output.

3.6.2.4 Implementation and Reproducibility

Both stages were implemented in Python. Stage 1 used XGBoost 3.2.0; Stage 2 used PyTorch, with weights initialised by Kaiming (He) normal initialisation matched to the ReLU nonlinearity and biases initialised to zero. Feature scaling for Stage 2

used a `StandardScaler` (zero mean, unit variance) fitted on the training split only and applied unchanged to validation and test data; Stage 1, as a tree-based model, requires no feature scaling. A fixed random seed (42) was used for all data splitting, model initialisation and training in both stages. Both stages auto-detect and use a GPU when one is available (checked via `nvidia-smi` for Stage 1 and `torch.cuda.is_available()` for Stage 2) and fall back to CPU otherwise. The model-selection criterion was the lowest 5-fold time-series cross-validation RMSE among the three candidate configurations for Stage 1 (Section 3.6.1), and the lowest Huber validation loss under early stopping for Stage 2 (this section).

3.7 Two-Stage Stacked Ensemble Architecture

3.7.1 Model Formulation

The complete two-stage prediction at station location (s) and time t is the following system of equations.

Let $\mathbf{x}(g, t) \in \mathbb{R}^{13}$ denote the Stage 1 feature vector at GIM grid point g and time t , comprising solar indices, geomagnetic indices, cyclic time encodings, geometry, and VTEC lags. The Stage 1 XGBoost ensemble maps $\mathbf{x}$ to a gridded VTEC prediction:

$$\hat{y}_{\text{GIM}}(g, t) = \sum_{k=1}^K \eta f_k(\mathbf{x}(g, t)), \quad (3.6)$$

where f_k is the k -th regression tree, $K = 500$, and $\eta = 0.10$ is the learning rate (Table 3.3).

Bilinear interpolation from the four neighbouring GIM grid corners lifts the gridded prediction to the station’s geographic location:

$$\hat{y}_{\text{GIM}}(s, t) = \mathcal{I}_{\text{bilin}}(\hat{y}_{\text{GIM}}(g, t); s). \quad (3.7)$$

The Stage 2 ANN takes an 18-dimensional feature vector $\mathbf{x}'(s, t)$ that combines the shared solar and geomagnetic drivers with continuous spatial inputs (lat_{IPP} , lon_{IPP}), the Stage 1 background $\hat{y}_{\text{GIM}}$, monthly cyclic encodings, the Dst rate of change, and residual lags – replacing the Stage 1 solar zenith angle, quasi-dipole coordinates, and GIM VTEC lags, as described in Section 3.6.2 – and predicts the debiased residual

defined in Eq. 3.4:

$$\hat{\Delta}(s, t) = f_{\text{ANN}}(\mathbf{x}'(s, t); \boldsymbol{\theta}), \quad (3.8)$$

with $\boldsymbol{\theta}$ the trainable weights of the 4-layer MLP that has been well described in Section 3.6.2.

The ANN is trained by minimising the Huber loss, which is quadratic for small residuals and linear for large ones. The linear tail limits the influence of the rare large VTEC excursions that arrive at storm onset, so they do not dominate the gradient:

$$L_{\delta}(r) = \begin{cases} \frac{1}{2} r^2, & |r| \leq \delta, \\ \delta |r| - \frac{1}{2} \delta^2, & |r| > \delta, \end{cases} \quad (3.9)$$

where $r = \Delta(s, t) - \hat{\Delta}(s, t)$ and $\delta = 3.0$ TECU matches the training configuration (Section 3.6.2).

The final ensemble prediction at any station-hour combines the Stage 1 background, the Stage 2 residual correction, and the station-specific debiasing term added back:

$$\boxed{\hat{y}_{\text{final}}(s, t) = \hat{y}_{\text{GIM}}(s, t) + \hat{\Delta}(s, t) + b_s} \quad (3.10)$$

where b_s is the empirical per-station bias estimated on the CORS training set (Section 3.6.2), subtracted from the target before the network is trained (Eq. 3.4) and added back here, and for locations not present in the training network, b_s was replaced by the global mean bias $\bar{b}$.

Figure 3.4 below shows the complete two-stage pipeline and the internal architecture of the Stage 2 ANN.

From the figure 3.4, its clearly seen that the Stage 1 supplied a global scale information for the 26 year data and the Stage 2 ANN acted as a local correction model. This architecture agrees with the stacked generalization principle of Wolpert (1992).

3.8 Evaluation Metrics

To evaluate the performance of the models designed, the following standard metric were used:




Figure 3.4: Two-stage stacked ensemble: (a) prediction pipeline, (b) Stage 2 ANN architecture.

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^N (y_i - \hat{y}_i)^2} \quad (3.11)$$

$$\text{MAE} = \frac{1}{N} \sum_{i=1}^N |y_i - \hat{y}_i| \quad (3.12)$$

$$\text{MBE} = \frac{1}{N} \sum_{i=1}^N (\hat{y}_i - y_i) \quad (3.13)$$

$$R^2 = 1 - \frac{\sum_i (y_i - \hat{y}_i)^2}{\sum_i (y_i - \bar{y})^2} \quad (3.14)$$

The Willmott index of agreement d (Willmott, 1981) which has a bounded effect especially for large dataset and in this case, was used in the Stage 1 due to the 26 years of data used. This index normally varies from between 0 (no agreement) to 1 (perfect agreement). R^2 was used for the CORS data in stage 2.

$$d = 1 - \frac{\sum_i (y_i - \hat{y}_i)^2}{\sum_i (|\hat{y}_i - \bar{y}| + |y_i - \bar{y}|)^2}. \quad (3.15)$$

Skill Score (Murphy, 1988) was also extensively used here. This represents how much better a model is as compared to a simple baseline set. In this case, the baselines are the persistence which is a 1 hour lag, the climatological mean and the IRI for external comparison.

$$\text{SS}_{\text{pers}} = 1 - \frac{\text{RMSE}_{\text{model}}^2}{\text{RMSE}_{\text{persistence}}^2}, \quad (3.16)$$

where the persistence baseline at station s predicts the VTEC observed one hour earlier at the same station. If $\text{SS}_{\text{pers}} > 0$ that would mean the model outperforms a simple 1-hour lag forecast and the reverse would also be true for underperformance.

$$\text{SS}_{\text{clim}} = 1 - \frac{\text{RMSE}_{\text{model}}^2}{\text{RMSE}_{\text{clim}}^2}, \quad (3.17)$$

where the climatological mean baseline predicts the per-(station, hour-of-day) median VTEC from the Stage 2 training window. SS_{clim} is the primary Stage 2 performance metric because it measures improvement over the best available static predictor.

$$SS_{\text{IRI}} = 1 - \frac{\text{RMSE}_{\text{model}}^2}{\text{RMSE}_{\text{IRI}}^2}, \quad (3.18)$$

In this case, the baseline is IRI-2020 evaluated at the CORS station locations. SS_{IRI} measures improvement over the operational empirical model(IRI2020). For storm time performance of the model, all the conditions for quiet ($Kp < 4$), and disturbed ($Kp \geq 4$) were obtained and followed.

IRI-2020 was implemented through a python package which compiles the official IRI-2020 fortran code with the standard files being driven by *iri-driver.f90* wrapper. The altitude chosen was a bottom altitude of 60 km, a top altitude of 2000 km in a 10 km step. the coefficients used were left at the default. and geographic latitude and longitude were input too for the CORS station locations. This was followed by latitudinal, diurnal and Storm-time assessment of model performance.

For spatial generalisability, a Leave-One-Station-Out (LOSO) cross-validation was carried out, where eight stations spanning the four corners and interior of the Uganda network (ARUA, KTGM, MRTO, MBAL, BUGI, MBRA, KABA, FPRT) were withheld in turn. The ANN was re-trained on the remaining stations and evaluated on the held-out station’s March–April 2026 test rows, with early stopping on the global validation set in each LOSO run. Two evaluation variants were computed. In the first (*lenient*), the withheld station’s own per-station systematic bias b_s estimated from that station’s own August–January residuals and its own residual-lag features (res_lag1h–res_lag24h, drawn from its own recent local GPS history) were still supplied at prediction time. Because both quantities are derived from the withheld station’s own observations, this variant leaks target-adjacent information about the “unseen” site and is reported only for continuity with earlier iterations of this analysis. The second, *strict*, variant removes this leakage: b_s for the withheld station is replaced by the leave-one-out mean bias of the remaining 39 stations, and its residual-lag inputs are set to their training-set mean (zero, after centring) rather than its own history – matching what is actually available at a genuinely new site with no local receiver. All spatial-generalisation results reported in Chapter 4 and Chapter 5 use the strict variant.

Chapter 4

Results and Discussion

4.1 TEC Characterisation

4.1.1 TEC Diurnal Variability

The diurnal TEC variability was studied with data from the 40 Uganda CORS station from the period of August 2025 to April 2026. Figure 4.1 shows the diurnal cycle of VTEC over Uganda from August 2025 to April 2026.




Figure 4.1: Diurnal cycle of VTEC over the Uganda CORS network, August 2025 to April 2026. .

The thick blue line represents the network median across the 40 stations, the dashed lines (ABON, APAC, KABA) are three stations spanning the network's latitudinal extent, chosen to illustrate the north-to-south gradient examined in this study. Longitudinal variation is not examined separately here, since Uganda's longitude extent (29.5° – 35° E) is narrow relative to its magnetic-latitude range (Section 3.1.2). The grayish shaded region represents the night (00:00–03:00 UT and 18:00 UT)

The main feature of the median diurnal cycle is a latitudinal north-to-south VTEC gradient that reverses direction over the course of the day. In the morning (07:00 UT

/ 10:00 LT), ABON (3.51° N, the network's northernmost station) leads KABA (−1.27° N, the southernmost) by up to 4.8 TECU. Around local noon, the two stations converge to within 1 TECU of each other. However, after 15:00 UT (18:00 LT), the ordering inverts completely: KABA exhibits a pronounced lead, passing ABON by as much as 9.1 TECU between 18:00 and 19:00 UT (21:00–22:00 LT). There is also an observed increase in TEC across all the station from 18 – 21 UT. The diurnal pattern observed agrees with observation from other low-latitude research on diurnal TEC variability and especially the evening enhancement. The observed gradient can be explained in terms of plasma drift (Abdu, 2016; Kelley, 2009). During sunrise, solar EUV radiation causes rapid photoionization, and since the magnetic field lines are nearly horizontal at the dip equator, this enhances the $\mathbf{E} \times \mathbf{B}$ drift, moving plasma to a region of lower electron loss and producing the higher values observed near the equator. ABON, which sits closer to the dip equator than the other stations, therefore experiences enhanced ionisation from 6:00–12:00 LT.

The post-sunset enhancement in TEC observed in this study may be explained in terms of the pre-reversal enhancement of the eastward electric field, which drastically accelerates the vertical $\mathbf{E} \times \mathbf{B}$ lift and causes a massive secondary dumping of plasma down the field lines. Because KABA (−10.9° magnetic) sits deeper toward the subtropical anomaly crest zone, it acts as the primary catchment basin for this transported plasma. This is why KABA experiences a prominent evening enhancement, driving the 9.1 TECU surplus over ABON long after local sunset.

4.1.2 Seasonal Variability

For seasonal variation, instead of the limited CORS data, the CODE Global Ionospheric Map data was used, as it has been shown to provide reliable VTEC data (Hernández-Pajares et al., 2009). Figure 4.2 shows the monthly mean VTEC over the Uganda 1° × 1° grid (−0.5°–4.5° N, 29.5°–34.5° E; Section 3.3.1) from 2000 to 2026.

From the figure, the minimum monthly mean VTEC is observed during the June–July solstice (21.6 TECU in June, 20.7 TECU in July) while high values were generally observed during the equinox months of March–April and November to December. March reaches the annual maximum at 35.0 TECU, exceeding the June–July mean (21.1 TECU) by 66%. There was an equinoctial asymmetry observed where March–April VTEC values were generally higher than the September–October VTEC values, and the September equinox showed a delayed response with a monthly mean peak in October (34.0 TECU) rather than September (29.1 TECU). There was also an




Figure 4.2: Monthly mean VTEC over the Uganda longitude sector from the CODE GIM record, 2000–2026. Vertical dashed lines mark the astronomical equinoxes.

asymmetry in the solstice months of June-July with generally low VTEC values as compared to the December-January months (25.8–27.4 TECU).

The observations of the equinox months having higher TEC than solstice is because during equinoxes, the sun is at zenith and therefore there is increased absorption of EUV flux which leads to increased ionization thereby causing higher TEC values as compared to solstice periods.

The unique one month lag of the equinox month of September is attributed to thermal variations in thermospheric winds $[O]/[N_2]$ that suppresses the september month while elevating the october month. This observation was also noted by (Kelley, 2009). This lag was also reported in an East African longitude sector research by Olwendo et al. (2016) that make it a regional feature.

The december solstice is the summer solstice in relation to the southern hemisphere. During summer equinox, the sun is closer to the southern stations and this means more ionisation for stations in the south as compared to the june solstctice.

4.1.3 Storm-Time TEC Variability

Storm time variability study of TEC was achieved by picking a representative storm event which was the November 2025 storm from 10th to 17th. Figure 4.3 aboves

shows a stacked-panel case study of the November 2025 event ($Dst_{\min} = -217$ nT). Generally the main phase of the storm produced pronounced TEC variability which




Figure 4.3: November 2025 geomagnetic storm ($Dst_{\min} = -217$ nT) over Uganda: VTEC at two CORS stations (a, b) and geomagnetic indices (c, d).

is evident in both stations. During the day time, TEC at the ABON station reached above its daily peak while RUKU station TEC tracked its own baseline by 1 TECU. The separation widened sharply after sunset.

The general high TEC variability during the main phase of the storm is due to prompt penetrating electric fields. During the main phase, the latitude electric fields get intensified which increases penetration to low and middle latitudes creating strong vertical EXB drift which lifts or depresses to the ionospheric F layer. TEC arose where the array is closest to the dip equator and fell where it is closest to the southern crest. this shows that an equatorial fountain was suppressed leading to these observations because plasma that was not lifted or carried to the poles stayed near the magnetic equator instead of feeding the crest and therefore the disturbance dynam drives it westwards. This effect intensifies while heading to the evening hours. The November 2025 storm case study shows the latitudinal asymmetry that Habarulema et al. (2013) documented across multi-event composites for the African sector.

4.2 Stacked Ensemble TEC Model

The TEC over Uganda was modeled in two stages using XGBoost model as the background model and Artificial Neural Network as the residual correction model. Before developing the model, feature importance was first assessed to see their contribution to the modelling process. This was done by a Shapley Additive EXplanation (SHAP) which works by assigning each feature contribution a numerical value to every individual prediction depending on importance. Figure 4.4 shows the direction of each input feature ranked in there order of importance.




Figure 4.4: SHAP beeswarm plot for the final 13-feature Stage 1 XGBoost model, over 5,000 observations from the 2025–2026 test set.

From Figure 4.4, VTEC(t-1h) and VTEC(t-24h) were the dominant features with high red magnitude and a SHAP value of up to about 55 TECU. This implies that the lag features which are the persistence measure and diurnal measures are the greatest contribution to prediction of VTEC. This is followed by the Solar Zenith Angle (SZA), the blue in this indicate a low SZA which means that the sun is closer to the zenith and therefore highly considered for VTEC prediction. The cyclical time parameters, the solar and geomagnetic indices and the quasi-dipole coordinates have a small contribution which is still relevant. Among the solar and geomagnetic predictors, F10.7 ranks above Kp and Dst since Solar EUV radiation is the primary driver of

ionospheric plasma production on both diurnal and solar-cycle timescales (Hathaway, 2015). The Dst value has a slight tail which affirms its contribution during critical roles. This SHAP analysis reflects the final thirteen-feature model (Section 3.5); ap , geographic latitude/longitude, and the 81-day F10.7 running mean were already excluded prior to this analysis, having ranked last on an earlier seventeen-feature candidate set.

4.2.1 XGBoost Model

The XGBoost model is a machine learning algorithm based on ensemble of decision trees which builds decision trees sequentially while reducing the loss function. It was trained on 22 years of data from 2000 to 2022, tuned on data from 2023 to 2024 and evaluated on data from 2025 to 2026. Table 4.1 summarizes the performance of the XGBoost model during quiet and storm conditions.

Table 4.1: Stage 1 XGBoost 5-fold ensemble on the CODE GIM grid, 2025–2026 test set, by geomagnetic activity condition and RMSE and MAE in TECU.

Condition	RMSE	MAE	MBE	R^2	d	N
All conditions	2.484	1.847	+0.073	0.9892	0.9973	457,236
Quiet ($K_p < 4$)	2.372	1.773	+0.068	0.9899	0.9975	384,912
Storm ($K_p \geq 4$)	3.012	2.242	+0.096	0.9854	0.9963	72,324
OOB training: RMSE = 1.548 TECU, $R^2 = 0.9942$ ($N = 5,886,630$)						

The model achieved an overall Root Mean Squared Error (RMSE) of 2.484 TECU and $R^2 = 0.9892$. Storm-time RMSE was 3.012 TECU, 27% above the quiet-time value of 2.372 TECU.

An R^2 of 0.9892 implies that the gradient boosting architecture successfully captured the highly non-linear relationship between and complex plasma dynamics for the equatorial dataset. The storm and quiet time RMSE imply that there was increased variability due to geomagnetic field distortions which increase fluctuation in the TEC. This model prediction is also consistent with the disturbed days having higher TEC variability as compared to quiet days. This demonstrates the model's ability to effectively represent the equatorial ionosphere and reconstructs the CODE GIM VTEC field with high accuracy, even under the extreme Solar Cycle 25 conditions of the test period. The storm-time R^2 remains at 0.9854, adding to this benefit as it shows that the model continues tracking enhanced TEC excursions during disturbed conditions. Due to inclusion of K_p and Dst, the model is able to avoid inferring measurements from lag features.

To further assess the performance of the model, the model predictions were compared with the observed GIM values. Figure 4.5 left panel(a) shows the distribution of Observed GIM which is shown with the black plot against XGBoost predictions shown with the red colour plot. The right panel(b) shows a correlation between observed VTEC and predicted VTEC.




Figure 4.5: Figure 4.5 shows the comparison between the Observed TEC and the predicted TEC from the XGBoost model prediction.

The scatter plot (Figure 4.5) shows tight agreement across the full 0–146 TECU range, with a mean overestimation of 2–3 TECU only at the highest observed values (> 80 TECU). VTEC above 80 TECU is essentially exclusive to SC25 conditions, which are underrepresented in the 23-year SC23–SC24 training record. The model is therefore extrapolating slightly beyond its training range at those values, which explains this upper-tail.

From the figure, its evident that the scatter plot on the right panel(a) clearly shows the model prediction closely follow the Observed GIM for the five fold cross validation result and the perfect 1:1 correlation with a slope of 1 that is also observed on the left panel(b) implies that the XGBoost model reproduced the GIM observations well with very good correlation.

Table 4.1 and Figure 4.5 have demonstrated that the XGBoost model has successfully reproduced the CODE GIM observations. This is very instrumental in the next Stacked modeling step.

4.2.2 Residual Correction

After the XGBoost model, the residue in equation 3.4 was modelled using a Multi-layered perceptron Artificial Neural Network. Figure 4.6 shows a plot of the training curve together with the validation curve for the residue correcting ANN.




Figure 4.6: Stage 2 ANN training history: training loss (Huber, $\delta = 3.0$ TECU) and validation loss against epoch. The blue line representing the training curve decreases steadily over the 35 epochs completed and the validation curve which is red reached its minimum at epoch 5 then diverged as training continued. The dashed green line represents the model's best checkpoint.

The training curve continued decreasing over the 35 epochs completed before early stopping, while the validation curve reached its minimum at epoch 5 (training loss 0.095, validation loss 0.095, a negligible gap) before diverging as training continued: by epoch 35 the training loss had fallen to 0.059 while the validation loss had risen to 0.119. This divergence implies that the station-specific residual detail the network began fitting late in training did not carry over to the next month. Early stopping (patience of 30 epochs with no improvement) was triggered at epoch 35, retaining the best-performing weights from epoch 5.

To effectively quantify the distribution of the determined residues from the Artificial Neural Network modelling, per-station mean residue was determined. Figure 4.7 shows the distribution of the per-station change in mean residual between the training and test data.

The figure 4.6 shows a distribution which is more towards the positive side of the zero line with a mean shift of $+0.92$ TECU and a standard deviation of 0.48 TECU. The distribution means that the the Stage 1 residual increased at the CORS stations




Figure 4.7: Histogram of the per-station change in mean Stage 1 residual, $\Delta = \text{Test (Mar–Apr)} - \text{Train (Aug–Jan)}$, across the 40 Uganda CORS stations. The dashed vertical line at $\Delta = 0$ represents no change in the mean residual between the two periods.

during the March–April test period relative to the August–January training period. This intern indicates that the Stage 1 residual structure is not strictly stationary between the two periods. Nevertheless, this shift is substantially smaller than the original Stage 1 systematic error, which was approximately 18.2 TECU. The residual correction can therefore be transferred to the independent test period while allowing for a modest increase in the remaining error under the changed seasonal conditions.

A comparable residual-correction skill was reported by Okoh et al. (2019) for AfriTEC’s local bias correction over the African sector, suggesting the leftover Stage 1 structure captured here is consistent with what other regional GNSS-TEC correction efforts have found, not an artefact specific to this dataset.

4.2.2.1 Final Ensemble

To construct the stacked ensemble model, the residual prediction is combined with the Stage 1 background and the per-station bias term using equation. 3.10. Table 4.2 summarises the resulting ensemble performance on the validation and test datasets.

On the February 2026 validation set, the ensemble achieved a RMSE of 3.277 TECU, MAE of 2.413 TECU, MBE of -0.399 , and R^2 of 0.9691. A persistence skill score of $+0.477$ indicates the ensemble outperforms a simple 1-hour lag forecast at the stations

Table 4.2: Ensemble model performance on the validation and test CORS datasets.

Split	RMSE	MAE	MBE	R^2	SS _{pers}	N
Validation (Feb 2026)	3.277	2.413	−0.399	0.9691	+0.477	24,946
Test (Mar–Apr 2026)	3.970	3.016	−1.180	0.9581	+0.361	48,118

SS_{pers} = Skill Score relative to the 1-hour persistence baseline.

it was trained on. On the independent March–April 2026 test set, performance remained strong across the 48,118 CORS hours: RMSE of 3.970 TECU, MAE of 3.016 TECU, MBE of −1.180 TECU, and R^2 of 0.9581. The positive persistence skill score of +0.361 shows the ensemble still outperforms a 1-hour persistence forecast on a period the model never saw during training; the climatological skill score for this same test set is +0.724 (Table 4.3). The scale of this correction, an 80% RMSE reduction relative to the Stage 1 background is consistent with reported gains from ANN-based local corrections elsewhere on the continent (Habarulema & McKinnell, 2012; Okoh et al., 2019), and reflects the same physical mechanism.

4.3 Evaluation of the Two-Stage Ensemble Model Performance

After the development of the Stacked ensemble machine learning model, the following section details the evaluation of this developed model against IRI and CORS data. This section also presents the diurnal, storm-time and spatial performance of the model.

4.3.1 Model Comparison Against CORS Ground Truth

The model was evaluated by comparing the reference baselines, IRI-2020, the raw CODE GIM field, the Stage 1 XGBoost prediction at CORS stations, and the full ensemble prediction to the ground truth using the CORS data from March to April 2026. Table 4.3 compares these tiers against ground truth; the raw CODE GIM row is reported separately from the Stage 1 XGBoost row so that any error XGBoost adds on top of the GIM field it was trained to reproduce can be distinguished from error already present in GIM itself.

The raw CODE GIM field, evaluated directly at the CORS station locations, gives

Table 4.3: Model comparison against the March–April 2026 CORS GPS VTEC ground truth

Model	RMSE	MAE	MBE	R^2	SS_{clim}	SS_{IRI}	N
Climatological mean (reference)	7.561	5.589	0.390	0.848	0.000	—	48,118
Persistence 1 h (reference)	6.215	4.860	0.001	0.897	0.324	—	48,078
IRI-2020	32.093	28.392	27.948	-1.738	-17.016	0.000	48,118
CODE GIM (raw) at CORS	19.760	18.050	17.990	-0.038	-5.830	0.621	48,118
Stage 1 XGBoost at CORS	20.133	18.331	18.275	-0.078	-6.090	0.606	48,118
Stage 2 (full ensemble)	3.970	3.016	-1.180	0.958	+0.724	+0.985	48,118

RMSE = 19.760 TECU, MBE = +17.990 TECU, and $R^2 = -0.038$. The Stage 1 XGBoost prediction closely reproduces this same error profile (RMSE = 20.133 TECU, MBE = +18.275 TECU, $R^2 = -0.078$), confirming that XGBoost is faithfully learning the CODE GIM field rather than introducing substantial additional error of its own; the negative R^2 in both rows implies worse performance than the time mean of the stations. The shared positive bias is therefore inherited from the mismatch between GIM and the CORS point measurements, not from a property of the XGBoost model itself. IRI-2020 carries an even larger bias RMSE = 32.093 TECU and MBE = +27.948 TECU ($R^2 = -1.74$). The GIM’s $2.5^\circ \times 5^\circ$ grid cell spans about 280 km north-south at Ugandan latitudes, wide enough to straddle the equatorial trough and the EIA crest within a single cell; the resulting cell-average value need not match the column-integrated VTEC measured directly overhead by any single GPS receiver within that cell. The high RMSE of the IRI-2020 can best imply that IRI explains none of the variance in the observed GPS VTEC. This is due to the altitude-integration/coordinate mismatch between IRI-2020’s modelled profile and GPS slant-to-vertical VTEC, and also due to the sparsity of ionosonde stations feeding IRI’s CCIR/URSI coefficients over equatorial East Africa.

A further comparison of the model performances was done for a station TIDO (3.04° N, 34.13° E). Figure 4.8 shows the 55-day time series at station TIDO (3.04° N, 34.13° E) for the full test period, with all the model tiers plotted against the GPS VTEC ground truth.

There is a clear overestimation of GPS VTEC by CODE GIM (blue) with the largest overestimation during daytime peak hours (1:00 pm to 3:00 pm local time). Stage 1 XGBoost (red) follows the GIM closely pattern and inherits the same systematic bias which give it a similar appearance to that of CODE GIM graph. The Stage 2 ensemble (green) closely tracks the GPS observations. The day time peak VTEC is primarily due the Equatorial Ionization anomaly. the Stage 1 prediction follows the GIM closely because it was to inherit the same long term patterns observed in the GIM dataset which intern also comes with its offsets. The ensemble performance




Figure 4.8: Time series of GPS-observed VTEC, CODE GIM, Stage 1 XGBoost background, and Stage 2 ANN ensemble at station TIDO (3.04° N, 34.13° E), March–April 2026.

closeness to the ground truth CORS shows the models good performance.

The latitudinal performance of the ensemble model was also assessed, two stations were chosen and model performance was evaluated. Figure 4.9 compares two stations at opposite ends of the Uganda latitudinal range. .

It is observed that their midday maxima are nearly identical over March 2026 and the two stations separate in the around 17:00–20:00 UT. MBRA carries 30.6 TECU against 23.0 TECU at ABIM. This is due to proximity to the creat of the southern Equatorial Ionisation Anomaly described in section 4.1.1. In both cases the Stage 2 ensemble eliminates the GIM positive bias and tracks the diurnal GPS VTEC amplitude, including that 7–8 TECU post-sunset separation between the two stations.

4.3.2 Diurnal Performance

In order to understand the ensemble daily performance, an analysis was performed to detail the average diurnal variability in the RMSE of the ensemble. Figure 4.10 shows the hourly RMSE of the ensemble model for the March to April test period.

The diurnal RMSE profile (Figure 4.10) shows Stage 1 background generally has the higher values of RMSE across the day, peaking in the early afternoon when the GIM scale-mismatch bias is amplified by the daytime TEC maximum. Stage 2 reduces this sharply, bringing RMSE to 3–6 TECU at all hours. The most notable feature is the post-sunset rise in ensemble RMSE between 18:00 and 23:00 LT (shaded in Figure 4.10). The higher RMSE values of Stage 1 are inherent to the GIM data which generally has higher TEC values than the CORS GPS TEC. The post sunset




Figure 4.9: Time series of GPS-observed VTEC, CODE GIM, Stage 1 XGBoost background, and Stage 2 ANN ensemble for March–April 2026. *Top: ABIM Bottom: MBRA*




Figure 4.10: Hourly RMSE by local time (LT = UT + 3 h) for the Stage 2 ensemble, Stage 1 background, persistence 1-h baseline, and climatological mean baseline, March–April 2026.

rise can be attributed to the pre-reversal enhancement (PRE) of the $E \times B$ drift and with equatorial plasma bubble onset during the March–April equinox (Kelley, 2009). Plasma bubbles introduce kilometre-scale TEC depletions that evolve faster than the hourly input cadence and cannot be captured from smooth solar and geomagnetic index features.

4.3.3 Storm-Time Performance

A storm-time performance analysis was also performed to see how the ensemble performs during disturbed conditions. Table 4.4 shows the quiet/storm stratification against the March–April 2026 CORS GPS VTEC ground truth.

Table 4.4: Stage 1 background, climatological mean baseline, and Stage 2 full ensemble against March–April 2026 CORS GPS VTEC, stratified by geomagnetic activity (storm threshold $K_p \geq 4$). RMSE and MAE in TECU.

Condition	Model	RMSE	MAE	MBE	R^2	SS_{clim}	N
Quiet ($K_p < 4$)	Stage 1 background	20.310	18.567	+18.510	−0.118	−6.320	40,798
	Climatological mean	7.507	5.520	+0.691	0.847	0.000	40,798
	Stage 2 ensemble	3.844	2.925	−1.084	0.960	+0.738	40,798
Storm ($K_p \geq 4$)	Stage 1 background	19.120	17.017	+16.967	0.122	−4.922	7,320
	Climatological mean	7.857	5.976	−1.285	0.852	0.000	7,320
	Stage 2 ensemble	4.608	3.527	−1.713	0.949	+0.656	7,320

The disturbed hours account for 15.2% of the test set ($N_{\text{storm}} = 7,320$) and the remaining 40,798 hours are quiet. The ensemble keeps a positive Skill Score under both geomagnetic regimes beating the climatological mean in every geomagnetic state observed during the test period. One unexpected result in Table 4.4 is that the Stage 1 background RMSE decreases from 20.3 TECU (quiet) to 19.1 TECU (storm). The higher storm time RMSE reflects high TEC variability driven by prompt penetration electric fields and the disturbance dynamo during March–April equinoctial storms. The climatological baseline degrades by about 5% under the same conditions (7.507 to 7.857 TECU), shows that storm-time TEC variability is genuinely harder to predict for the baseline as well. The ensemble’s 20% degradation is proportional to the difficulty of the conditions. The unexpected RMSE change of Stage 1 background is attributed to the constant positive bias of ~ 18 TECU. During storms, the equatorial fountain drives observed VTEC upward, which partially closes the gap between the biased prediction and the higher observed value. Stage 1 does not track the storm dynamics. it just happens to benefit from the storm-time VTEC enhancement. On the ensemble, the storm-time MBE of -1.71 TECU (a shift of

−0.63 TECU from quiet-time) shows a small systematic under-prediction of about 7% of the storm-time mean VTEC.

4.3.4 Spatial Performance

The spacial performance was done using the CORS dataset particularly at the stations to see what the average RMSE would be. Figure 4.11 shows the spatial distribution of the RMSE from model prediction at the station. Panel (a) shows the distribution of RMSE while panel (b) shows the distribution of skill score against climatological mean.




Figure 4.11: Stage 2 ensemble over the 40 Uganda CORS stations during March–April 2026: (a) per-station RMSE (TECU), (b) per-station Skill Score against the climatological mean baseline.

A uniform performance is observed across the network. RMSE ranges narrowly from 3.550 TECU at the best station (ABIM) to 4.500 TECU at the worst (KABA). All 40 stations produced a positive Skill Score relative to the climatological baseline, which is consistent with the ANN generalising across the full network without over-fitting near the highest training-data density.

Genuine spatial generalisation – performance at a site with no local CORS receiver – was tested with a Leave-One-Station-Out (LOSO) cross-validation over the same eight stations. Figure 4.12 compares, for each held-out station, its regular test RMSE (station present in training), its lenient LOSO RMSE (station withheld from training but still supplied its own bias and residual-lag history at prediction time), and its strict LOSO RMSE (station withheld from training, bias, and residual-lag construction alike).




Figure 4.12: Leave-One-Station-Out cross-validation of the Stage 2 ANN for eight stations spanning the Uganda CORS network, March–April 2026, comparing regular (station in training), lenient LOSO, and strict LOSO RMSE.

Pooled across the eight held-out stations ($N = 9,278$), the strict LOSO RMSE is 6.316 TECU, compared with 4.701 TECU for the lenient variant and 4.440 TECU for the same eight stations with their own receiver in training – a penalty of +1.877 TECU (+42.3%) for removing the local receiver entirely, of which +1.615 TECU (+34.3%) is attributable to the bias/residual-lag leakage in the lenient procedure alone. The penalty is uneven across the network: ARUA (+3.11 TECU), MBAL (+2.08 TECU), KTGM (+2.02 TECU) and BUGI (+2.02 TECU) lose the most, while MRTO (+1.49 TECU), FPRT (+1.50 TECU) and MBRA (+1.41 TECU) are comparatively robust. Skill Score against 1-hour persistence turns negative under the strict variant at ARUA (−0.138), KTGM (−0.030), MBAL (−0.045), BUGI (−0.043) and KABA (−0.044), meaning the corrected model does not clearly beat a naive persistence forecast at these five locations once local receiver information is unavailable. Only MRTO, MBRA and FPRT retain positive skill (+0.05 to +0.09). This indicates that the ensemble’s spatial generalisation, while real, is markedly weaker than its same-network in-training performance and should not be claimed at the same accuracy level as the headline test metrics.

Chapter 5

Conclusions and Recommendations

5.1 Conclusions

This study modelled ionospheric Total Electron Content (TEC) over Uganda using a two-stage stacked ensemble. Stage 1, an XGBoost model, was trained on 26 years of CODE Global Ionosphere Map data. Stage 2, a feedforward ANN, used six months of GPS VTEC observations from the 40-station Uganda CORS network as a local correction model. The conclusions from this research are as follows.

The diurnal TEC showed maximum values from 1:00 pm to 5:00 pm local time, peaking at 3:00 pm, with a minimum occurring at about 6:00 am local time. Seasonal variability results showed higher TEC during equinox months than during solstice months. The latitudinal gradient reversed direction over the course of the day (Section 4.1.1): the network's northern station (ABON) led the southern station (KABA) in the morning, while KABA led ABON after sunset.

The stacked ensemble model reproduced the diurnal, seasonal and spatial TEC characteristics over Uganda. Comparison of the model with IRI-2020 (Table 4.3) showed that the ensemble reduced RMSE from 32.093 TECU (IRI-2020) to 3.970 TECU, corresponding to a skill score of $SS_{\text{IRI}} = +0.985$.

5.2 Recommendations

The stacked ensemble demonstrated in this study suggest that it is an appropriate TEC model over equatorial latitude region of Uganda and therefore can be adopted.

Extension of the CORS training archive to atleast 24 months to give the Stage 2 ANN coverage of both equinoctial peak and adding predictors like the southward component of the interplanetary magnetic field (B_z) and the Rate of TEC change Index (ROTI) to characterises the small-scale irregularities associated with post-sunset plasma bubbles.

References

- Abdu, M. A. (2016). Electrodynamics of ionospheric weather over low latitudes. *Geoscience Letters*, 3, 11. <https://doi.org/10.1186/s40562-016-0043-6> (cit. on pp. 26, 48).
- Akinyemi, G. A., Bello, S. A., Kazeem, A. K., Sanyaolu, M. E., & Olabode, O. P. (2026). Ionospheric response to the strong geomagnetic storm of January 2025 at African equatorial stations. *Ajayi Crowther Journal of Pure and Applied Sciences*, 5(2), 57–67. <https://doi.org/10.56534/acjpas.v5i2.204> (cit. on p. 2).
- Amabayo, E. B., Jurua, E., Cilliers, P. J., & Habarulema, J. B. (2014). Climatology of ionospheric scintillations and TEC trend over the Ugandan region. *Advances in Space Research*, 53(5), 734–743. <https://doi.org/10.1016/j.asr.2013.12.015> (cit. on pp. 11, 24, 26, 27).
- Andima, G., Amabayo, E. B., Jurua, E., & Cilliers, P. J. (2019). Modeling of GPS total electron content over the African low-latitude region using empirical orthogonal functions. *Annales Geophysicae*, 37(1), 65–76. <https://doi.org/10.5194/angeo-37-65-2019> (cit. on pp. 24, 28).
- Angelopoulos, A. N., & Bates, S. (2023). Conformal prediction: A gentle introduction. *Foundations and Trends in Machine Learning*, 16(4), 494–591. <https://doi.org/10.1561/22000000101>
- Ba, J. L., Kiros, J. R., & Hinton, G. E. (2016). *Layer normalization*. arXiv: [1607.06450](https://arxiv.org/abs/1607.06450) [stat.ML]. <https://arxiv.org/abs/1607.06450> (cit. on p. 40).
- Banks, P. M., & Kockarts, G. (1973). *Aeronomy* (Vol. Part A & Part B). Academic Press. (Cit. on p. 8).
- Bilitza, D., Pezzopane, M., Truhlik, V., Altadill, D., Reinisch, B. W., & Nava, B. (2022). The International Reference Ionosphere model: A review and description of an ionospheric benchmark. *Reviews of Geophysics*, 60(4), e2022RG000792. <https://doi.org/10.1029/2022RG000792> (cit. on pp. 16, 18).
- Blewitt, G. (1990). An automatic editing algorithm for GPS data. *Geophysical Research Letters*, 17(3), 199–202. <https://doi.org/10.1029/GL017i003p00199> (cit. on p. 11).
- Bolaji, O. S., Adekoya, B. J., Adebisi, S. J., Adebisin, B. O., & Ikubanni, S. O. (2022). The African equatorial ionization anomaly response to the St. Patrick's Day storms of March 2013 and 2015. *Astrophysics and Space Science*, 367(1), 1 (cit. on pp. 15, 16).
- Breiman, L. (1996). Stacked regressions. *Machine Learning*, 24(1), 49–64. <https://doi.org/10.1007/BF00117832> (cit. on p. 23).

- Cele, G., Andima, G., Habyarimana, V., & Jurua, E. (2023). Mapping of the ionospheric total electron content over the East African low-latitude region. *Advances in Space Research* (cit. on pp. 2, 24, 28).
- Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785–794. <https://doi.org/10.1145/2939672.2939785> (cit. on p. 23).
- Ciraolo, L., Azpilicueta, F., Brunini, C., Meza, A., & Radicella, S. M. (2007). Calibration errors on experimental slant total electron content (TEC) determined with GPS. *Journal of Geodesy*, 81, 111–120 (cit. on pp. 12, 31).
- Davies, K. (1990). *Ionospheric radio*. IET. (Cit. on pp. 1, 8, 11).
- D’ujanga, F. M., Mubiru, J., Twinamasiko, B. F., Basalirwa, C., & Ssenyonga, T. J. (2012). Total electron content variations in equatorial anomaly region. *Advances in Space Research*, 50(4), 441–449. <https://doi.org/10.1016/j.asr.2012.03.022> (cit. on pp. 24, 27).
- Fejer, B. G., & Kelley, M. (1980). Ionospheric irregularities. *Reviews of Geophysics*, 18(2), 401–454 (cit. on pp. 14, 15, 17).
- Friedman, J. H. (2001). Greedy function approximation: A gradient boosting machine. *Annals of Statistics*, 29(5), 1189–1232. <https://doi.org/10.1214/aos/1013203451> (cit. on p. 23).
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press. (Cit. on p. 41).
- Goodman, J. M. (2005). *Space weather and telecommunications*. Springer Science & Business Media. <https://doi.org/10.1007/b101673> (cit. on p. 6).
- Habarulema, J. B., & McKinnell, L.-A. (2012). Investigating the performance of neural network backpropagation algorithms for TEC estimations using South African GPS data. *Annales Geophysicae*, 30, 857–866. <https://doi.org/10.5194/angeo-30-857-2012> (cit. on pp. 21, 56).
- Habarulema, J. B., McKinnell, L.-A., Burešová, D., Zhang, Y., Seemala, G., Ngwira, C., Chum, J., & Opperman, B. (2013). A comparative study of TEC response for the African equatorial and mid-latitudes during storm conditions. *Journal of Atmospheric and Solar-Terrestrial Physics*, 102, 105–114. <https://doi.org/10.1016/j.jastp.2013.05.008> (cit. on pp. 28, 50).
- Habarulema, J. B., McKinnell, L.-A., & Cilliers, P. J. (2009). Prediction of global positioning system total electron content using neural networks over South Africa. *Journal of Atmospheric and Solar-Terrestrial Physics*, 71(15-16), 1842–1850. <https://doi.org/10.1016/j.jastp.2009.07.013> (cit. on p. 21).

- Hathaway, D. H. (2015). The solar cycle. *Living Reviews in Solar Physics*, 12, 1–87 (cit. on pp. 7, 52).
- Hernández-Pajares, M., Juan, J. M., Sanz, J., Orus, R., Garcia-Rigo, A., Feltens, J., Komjathy, A., Schaer, S. C., & Krankowski, A. (2009). The IGS VTEC maps: A reliable source of ionospheric information since 1998. *Journal of Geodesy*, 83(3–4), 263–275. <https://doi.org/10.1007/s00190-008-0266-1> (cit. on pp. 2, 3, 19, 20, 48).
- Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735> (cit. on p. 22).
- Hofmann-Wellenhof, B., Lichtenegger, H., & Wasle, E. (2007). *GNSS—global navigation satellite systems: GPS, GLONASS, Galileo, and more*. Springer Science & Business Media. (Cit. on pp. 9–12).
- Jeong, S.-H., Lee, W. K., Kil, H., Jang, S., Kim, J.-H., & Kwak, Y.-S. (2024). Deep learning-based regional ionospheric total electron content prediction: Long Short-Term Memory (LSTM) and convolutional LSTM approach. *Space Weather*, 22, e2023SW003763. <https://doi.org/10.1029/2023SW003763> (cit. on p. 22).
- Kelley, M. C. (2009). *The earth’s ionosphere: Plasma physics and electrodynamics*. Academic press. (Cit. on pp. 1, 2, 7, 8, 14, 15, 26, 48, 49, 60).
- Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. *International Conference on Learning Representations (ICLR)* (cit. on p. 41).
- Klobuchar, J. A. (1987). Ionospheric time-delay algorithm for single-frequency GPS users. *IEEE Transactions on Aerospace and Electronic Systems*, AES-23(3), 325–331. <https://doi.org/10.1109/TAES.1987.310829> (cit. on p. 17).
- Klobuchar, J. A. (1996). Ionospheric effects on GPS. In B. W. Parkinson & J. J. Spilker (Eds.), *Global positioning system: Theory and applications, Vol. I* (pp. 485–515). American Institute of Aeronautics; Astronautics. (Cit. on pp. 10, 18).
- Komjathy, A. (1997). *Global ionospheric total electron content mapping using the Global Positioning System* [Doctoral dissertation, Department of Geodesy and Geomatics Engineering, University of New Brunswick] [Technical Report No. 188]. (Cit. on pp. 2, 13).
- Laštovička, J. (2009). Lower ionosphere response to external forcing: A brief review. *Advances in Space Research*, 43(1), 1–14 (cit. on p. 7).
- Liu, L., Morton, Y. J., & Liu, Y. (2022). ML prediction of global ionospheric TEC maps. *Space Weather*, 20, e2022SW003135. <https://doi.org/10.1029/2022SW003135> (cit. on p. 22).

- Loshchilov, I., & Hutter, F. (2017). SGDR: Stochastic gradient descent with warm restarts. *International Conference on Learning Representations (ICLR)* (cit. on p. 41).
- Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30, 4765–4774 (cit. on p. 23).
- Luo, J., Yu, X., Fang, H., & Zhang, L. (2023). Prediction of global ionospheric total electron content (TEC) based on SAM-ConvLSTM model. *Space Weather*, 21(12), e2023SW003707. <https://doi.org/10.1029/2023SW003707> (cit. on p. 22).
- Mannucci, A. J., Iijima, B. A., Lindqwister, U. J., Pi, X., Sparks, L., & Wilson, B. D. (1999). GPS and ionosphere. In W. R. Stone (Ed.), *Review of radio science 1996–1999* (pp. 625–665). Oxford University Press / URSI. (Cit. on pp. 1, 12, 13, 31).
- Mao, S., Hernández-Pajares, M., & Soja, B. (2025). A review of machine learning-based ionospheric spatial and temporal modeling. *Journal of Geophysical Research: Machine Learning and Computation*, 2, e2024JH000555. <https://doi.org/10.1029/2024JH000555> (cit. on p. 22).
- Mao, S., Pan, Y., Klopotek, G., Schartner, M., Krásná, H., de Witt, A., & Soja, B. (2025). Enhanced global ionospheric mapping using deep ensemble neural networks with uncertainty quantification. *Space Weather*, 23, e2025SW004446. <https://doi.org/10.1029/2025SW004446> (cit. on p. 3).
- Mitra, S. (1946). Geomagnetic control of region f2 of the ionosphere. *Nature*, 158(4019), 668–669 (cit. on pp. 2, 15).
- Mungufeni, P., Habarulema, J. B., & Migoya-Orúé, Y. (2016). Statistical modeling of ionospheric TEC from GPS observations over an equatorial African region during a year of low solar activity. *Advances in Space Research*, 57(7), 1647–1659. <https://doi.org/10.1016/j.asr.2016.01.014> (cit. on p. 28).
- Murphy, A. H. (1988). Skill scores based on the mean square error and their relationships to the correlation coefficient. *Monthly Weather Review*, 116(12), 2417–2424. [https://doi.org/10.1175/1520-0493\(1988\)116<2417:SSBOTM>2.0.CO;2](https://doi.org/10.1175/1520-0493(1988)116<2417:SSBOTM>2.0.CO;2) (cit. on p. 45).
- Nava, B., Coisson, P., & Radicella, S. M. (2008). A new version of the NeQuick ionosphere electron density model. *Journal of Atmospheric and Solar-Terrestrial Physics*, 70(15), 1856–1862. <https://doi.org/10.1016/j.jastp.2008.01.015> (cit. on p. 19).
- Nibigira, J. d. D., Venkat, D. R., & Prasad, G. S. V. (2019). Analysis of ionospheric GPS-TEC variability over Uganda IGS station during descending phase of

- 24th solar cycle, 2018 year. *International Journal of Scientific & Technology Research*, 8(12), 3947–3950 (cit. on pp. 2, 5, 24, 28).
- Nigusie, A., Tebabal, A., & Galas, R. (2024). Modeling ionospheric TEC using gradient boosting based and stacking machine learning techniques. *Space Weather*, 22, e2023SW003821. <https://doi.org/10.1029/2023SW003821> (cit. on pp. 3, 23, 24, 37).
- Ogwang, J. B. (2020). Evaluation of GNSS CORS data quality in Uganda. *FIG Working Week 2020: Smart Surveyors for Land and Water Management* (cit. on pp. 24, 29).
- Okoh, D., Habarulema, J. B., Rabiou, B., Seemala, G. K., Pfaff, R., Stoneback, R., Nava, B., Ekelle, B., Shiokawa, K., & Otsuka, Y. (2019). AfriTEC: A neural network-based ionospheric total electron content model for Africa. *Space Weather*, 17(11), 1629–1644. <https://doi.org/10.1029/2019SW002241> (cit. on pp. 3, 21, 22, 55, 56).
- Olwendo, O. J., Yamazaki, Y., Cilliers, P. J., Baki, P., & Doherty, P. (2016). A study on the variability of ionospheric total electron content over the East African low-latitude region and storm time ionospheric variations. *Radio Science*, 51(9), 1503–1518. <https://doi.org/10.1002/2015RS005785> (cit. on pp. 2, 5, 15, 16, 26, 28, 49).
- Ren, X., Yang, P., Liu, H., Chen, J., & Liu, W. (2022). Deep learning for global ionospheric TEC forecasting: Different approaches and validation. *Space Weather*, 20, e2021SW003011. <https://doi.org/10.1029/2021SW003011> (cit. on p. 22).
- Richmond, A. D. (1995). Ionospheric electrodynamics using magnetic apex coordinates. *Journal of Geomagnetism and Geoelectricity*, 47(2), 191–212. <https://doi.org/10.5636/jgg.47.191> (cit. on pp. 14, 17, 33, 37).
- Rishbeth, H. (1997). The ionospheric e-layer and f-layer dynamos—a tutorial review. *Journal of Atmospheric and Solar-Terrestrial Physics*, 59(15), 1873–1880 (cit. on pp. 2, 14, 26).
- Rishbeth, H., & Garriott, O. K. (1969). *Introduction to ionospheric physics*. Academic Press. (Cit. on pp. 1, 6, 7, 14).
- Schaer, S., Gurtner, W., & Feltens, J. (1998). IONEX: The IONosphere Map EX-change format version 1. *Proceedings of the IGS AC Workshop*, 233–247 (cit. on pp. 2, 3, 19, 20).
- Schunk, R., & Nagy, A. (2009). *Ionospheres: Physics, plasma physics, and chemistry*. Cambridge university press. (Cit. on pp. 1, 6, 7, 14, 15, 17).
- Seemala, G. K. (2023). Estimation of ionospheric total electron content (TEC) from GNSS observations. In A. K. Singh & S. Tiwari (Eds.), *Earth observation*,

- atmospheric remote sensing* (pp. 63–84). Elsevier. <https://doi.org/10.1016/B978-0-323-99262-6.00022-5> (cit. on pp. 12, 29, 31).
- Sisay, A., Nigusie, A., Tsidu, G. M., & Mahrous, A. (2024). Machine learning based storm time modeling of ionospheric vertical total electron content over Ethiopia. *Scientific Reports*, *14*(1), 19311. <https://doi.org/10.1038/s41598-024-69738-0> (cit. on pp. 23, 24).
- Spogli, L., Aol, S., Buchert, S., Cesaroni, C., Cherniak, I., De Franceschi, G., Habarulema, J. B., Jin, Y., Liu, Y., Mitchell, C., Pignalberi, A., Tornow, C., Tzagouri, I., Urbar, J., & Zakharenkova, I. (2023). Stepping into an equatorial plasma bubble with a Swarm overfly. *Space Weather*, *21*(4), e2022SW003331. <https://doi.org/10.1029/2022SW003331> (cit. on p. 11).
- Ssessanga, N., Yoshikawa, A., Uwamahoro, J., & Habarulema, J. B. (2024). A neural network approach for the prediction of low-latitude ionospheric scintillation indices over the African region. *Space Weather*, *22*(1), e2023SW003715. <https://doi.org/10.1029/2023SW003715> (cit. on p. 21).
- Ssessanga, N., Kim, Y. H., Habarulema, J. B., & Kwak, Y.-S. (2019). On imaging South African regional ionosphere using 4D-var technique. *Space Weather*, *17*(11), 1584–1604 (cit. on p. 28).
- Sun, M., Liu, T., Ding, Z., Liu, J., Huang, Y., Zhang, K., Fang, S., Li, S., Kong, Q., & Chen, B. (2025). Real-time regional ionosphere modeling with RFR-Net over China. *Space Weather*, *23*(1), e2024SW004237. <https://doi.org/10.1029/2024SW004237> (cit. on p. 24).
- Tariku, Y. A. (2015). Patterns of GPS-TEC variation over low-latitude regions (African sector) during the deep solar minimum (2008 to 2009) and solar maximum (2012 to 2013) phases. *Earth, Planets and Space*, *67*, 35. <https://doi.org/10.1186/s40623-015-0202-6> (cit. on p. 28).
- Twinomugisha, F., Ssebiyonga, N., & D’ujanga, F. M. (2017). TEC derived from some GPS stations in the East African equatorial region and comparison with the TEC from the NeQuick2 model. *Advances in Space Research*, *60*(9), 1905–1920. <https://doi.org/10.1016/j.asr.2017.06.007> (cit. on pp. 19, 24, 27).
- Uwamahoro, J. C., McKinnell, L.-A., & Habarulema, J. B. (2018). Reconstruction of storm-time total electron content using ionospheric tomography and artificial neural networks: A comparative study over the African region. *Radio Science*, *53*(12), 1361–1373. <https://doi.org/10.1029/2017RS006499> (cit. on pp. 21, 28).
- Willmott, C. J. (1981). On the validation of models. *Physical Geography*, *2*(2), 184–194. <https://doi.org/10.1080/02723646.1981.10642213> (cit. on p. 45).

- Wolpert, D. H. (1992). Stacked generalization. *Neural Networks*, 5(2), 241–259. [https://doi.org/10.1016/S0893-6080\(05\)80023-1](https://doi.org/10.1016/S0893-6080(05)80023-1) (cit. on pp. 23, 43).
- Zhou, Y., Liu, J., Li, S., & Li, Q. (2024). Ionospheric TEC prediction based on ensemble learning models. *Space Weather*, 22, e2023SW003790. <https://doi.org/10.1029/2023SW003790> (cit. on pp. 3, 24).